



NSP

Network Services Platform

Network Functions Manager - Packet (NFM-P)

Release 25.4

Statistics Management Guide

3HE-21474-AAAA-TQZZA

Issue 1

May 2025

© 2025 Nokia.

Use subject to Terms available at: www.nokia.com/terms

Legal notice

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Contents

About this document	7
Part I: Statistics overview	9
1 Statistics management	11
1.1 Statistics management in the NFM-P	11
1.2 NFM-P statistics types	13
1.3 NFM-P statistics scalability	15
1.4 NFM-P statistics and OSS applications	16
Part II: Statistics collection configuration	17
2 Statistics collection	19
2.1 Statistics collection in the NFM-P	19
2.2 To create and manage custom statistics aggregators	23
2.3 Statistics policies	24
3 Performance statistics collection	29
3.1 Collecting performance statistics	29
3.2 Pathway: performance statistics collection	30
3.3 To create or modify an NE MIB statistics policy using a top-down method	31
3.4 To modify an NE MIB statistics policy using a bottom-up method	32
3.5 To create or modify a specific MIB statistics policy using a top-down method	33
3.6 To create or modify a specific MIB statistics policy using a bottom-up method	35
3.7 To assign the default 1830 VWM OSU performance management policy to 1830 devices	36
3.8 To configure polling for a MIB statistics class	37
3.9 To configure a statistics policy for MIB statistics	38
3.10 To delete performance statistics records	38
4 Accounting statistics	41
4.1 Accounting statistics collection	41
4.2 To create and manage periodic accounting statistics calculations	44
4.3 Pathway: accounting statistics collection	45
4.4 To configure a file policy	47
4.5 To configure an accounting policy	48
4.6 To configure a statistics policy for accounting statistics on a SAP or an SDP	50
4.7 To configure a statistics policy for accounting statistics on a network interface	52
4.8 To configure a statistics policy for accounting statistics on a subscriber	53

4.9	To configure a statistics policy for AA accounting statistics on a subscriber	54
4.10	To configure a statistics policy for an AA accounting statistics application.....	55
4.11	To configure a statistics policy for an AA accounting statistics application group	56
4.12	To configure a statistics policy for an AA accounting statistics protocol	57
4.13	To configure ingress and egress accounting policies for an LSP	58
4.14	To delete accounting statistics records	60
5	Flow statistics collection.....	61
5.1	Flow statistics collection in the NFM-P	61
5.2	Pathway: configure flow statistics collection	62
5.3	Pathway: configure AA Cflowd special study statistics collection.....	64
5.4	To configure flow statistics aggregation	64
5.5	To set upload targets and policies.....	69
5.6	To configure the transfer of results files	71
5.7	To configure AA filters.....	73
5.8	To configure an AA Cflowd special-study policy	75
5.9	To configure AA subnet summarization filters	78
5.10	To configure the AA flow data persistence	80
5.11	To specify the NEs for flow statistics collection	81
5.12	To extract an NFM-P network data snapshot	82
5.13	To export IPDR statistics files to a different format.....	83
6	Server performance statistics.....	87
6.1	Server performance statistics collection.....	87
6.2	Pathway: server performance statistics collection	87
6.3	To configure a statistics policy for server performance statistics	87
6.4	To configure a statistics collection policy for server performance statistics	88
6.5	To delete server performance statistics records.....	89
	Part III: Statistics presentation	91
7	Statistics presentation.....	93
7.1	Statistics presentation in the NFM-P	93
8	Viewing statistics	97
8.1	Viewing statistics in the NFM-P	97
8.2	Pathway: viewing statistics.....	97
8.3	To view on-demand statistics	98
8.4	To view statistics for a set of objects	99
8.5	To view network accounting statistics	100

8.6	To view performance statistics	101
8.7	To view server performance statistics	102
8.8	To view service accounting statistics.....	103
8.9	To view subscriber accounting statistics	104
8.10	To view per-subscriber AA accounting statistics.....	105
8.11	To view AA accounting statistics for an application	106
8.12	To view AA accounting statistics for an application group	107
8.13	To view AA accounting statistics for a protocol.....	109
9	Graphing statistics	111
9.1	Graphing statistics overview	111
9.2	Pathway: graphing statistics.....	116
9.3	To create a statistics graph.....	116
9.4	To modify a statistics graph	121
9.5	To plot utilization statistics	122
9.6	To create a customized plotter profile	123
9.7	To plot statistics using a plotter profile	124
9.8	To manage plotter profiles	124
	Part IV: Statistics content	127
10	Statistics record format.....	129
10.1	Statistics record format in the NFM-P	129
10.2	Performance statistics	129
10.3	Accounting statistics.....	131
10.4	Server performance statistics.....	136
10.5	Nokia internal statistics	147

About this document

Purpose

The *NSP NFM-P Statistics Management Guide* describes statistics management using the NFM-P. Supported performance and accounting statistics are searchable in the Statistics Search Tool.

Scope

The scope of this document is limited to the NFM-P. Many configuration, monitoring, and assurance functions that can be accomplished from the NFM-P Java GUI are also delivered in NSP. Readers of this NFM-P guide should familiarize themselves with the capabilities of the NSP, which often offer more efficient and sophisticated features for network and service management.

Safety information

For your safety, this document contains safety statements. Safety statements are given at points where risks of damage to personnel, equipment, and operation may exist. Failure to follow the directions in a safety statement may result in serious consequences.

Document support

Customer documentation and product support URLs:

- [Documentation Center](#)
- [Technical support](#)

How to comment

Please send your feedback to documentation.feedback@nokia.com.

Part I: Statistics overview

Overview

Purpose

This volume contains an overview of NFM-P statistics management.

Contents

Chapter 1, Statistics management	11
--	----

1 Statistics management

1.1 Statistics management in the NFM-P

1.1.1 Overview

The NFM-P provides a scalable platform for reliable statistics collection from managed NEs, and also provides statistics for monitoring NFM-P system processes and functions.

Depending on the statistics type, you can do the following with the collected data:

- Use an NFM-P GUI client to:
 - Schedule or perform on-demand collection.
 - View raw counter values in tabular format on the Statistics tab of an object properties form.
 - View multiple series of historical or real-time counter values in graphical format using the Statistics Plotter.
 - Save the tabular data or graphical representation to a file.
- Use an OSS client to retrieve the data for processing and analysis.
- Use the NSP Analytics to generate comprehensive, customized reports for traffic monitoring and trend analysis.



Note: Before you configure statistics collection, you must consider the NFM-P resource constraints; see [1.3 “NFM-P statistics scalability” \(p. 15\)](#) for general information, and the *NSP Planning Guide* for release-specific scalability guidelines.

The NFM-P can collect the following statistics types:

- performance statistics—collected by polling NE MIBs and transferred to the NFM-P using SNMP, or, for wireless PM counters, transferred from NEs to the NFM-P using FTP
- accounting statistics—collected in files on NEs and transferred to the NFM-P using FTP or SCP
- flow statistics—forwarded by NEs to a target file server for retrieval by an OSS, or to the NFM-P for use by the NSP Analytics.
- server performance statistics—collected by NFM-P system processes and stored in the NFM-P database

See [1.2 “NFM-P statistics types” \(p. 13\)](#) for more information about each statistic type.

Policy-based collection

To collect statistics such as NE performance or accounting statistics, the NFM-P uses policies that specify the following:

- the network or service objects to collect statistics from
- the statistics counters to collect
- the collection rate
- how long the NFM-P is to retain the collected statistics data

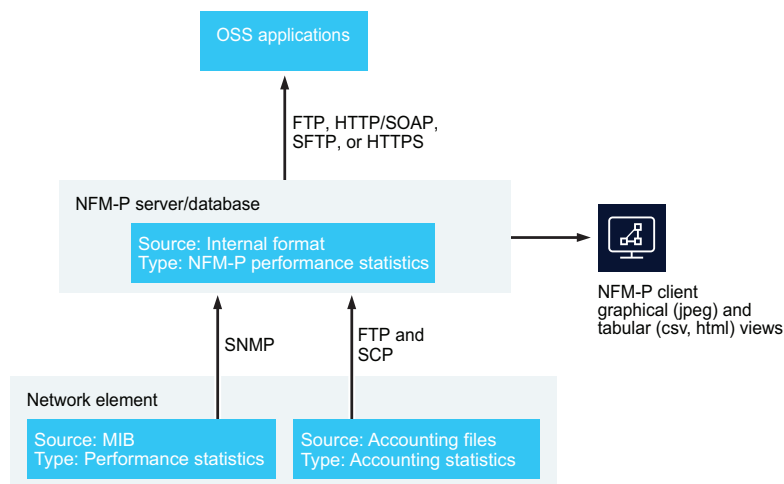
From the Statistics tab of an object properties form, you can use the following buttons to perform an on-demand collection of network or server performance statistics:

- Collect—returns one statistics record
- Collect All—returns one statistics record for each statistic type that the object supports

Note: On-demand statistics collection may yield inconsistent results the first time it is run against an object. To obtain valid results, you must run the on-demand statistics collection more than once.

Statistics are stored for a configurable retention period that is defined in an NFM-P statistics policy. When the retention period elapses, the statistics are removed. However, statistics that are collected for real-time display are not stored; the statistics are available only for the duration of the current session, and only for the operator that initiates the session.

Figure 1-1 Basic statistics-collection architecture



19719

1.2 NFM-P statistics types

1.2.1 Overview

Table 1-1 Characteristics of statistics types

Characteristic	Statistics type			
	Performance	Accounting • service • network • subscriber • AA	Server performance	Flow • AA Cflowd
Typical uses	NE monitoring NE troubleshooting	Billing SLA compliance Trend analysis	NFM-P system monitoring NFM-P system troubleshooting	Traffic monitoring Data analytics
Source	NE MIBs	All but AA: NE accounting files and NE MIBs AA: ISA-AA MIBs	Internal NFM-P system metrics	Traffic samples
On-demand collection	Yes	No	Yes	No
Scheduled collection	Yes Requires Statistics or Specific MIB policy	Yes Requires File and Accounting policies	Yes Requires Server Performance Statistics policy	Yes Requires NSP Flow Collector
Real-time graphing	Yes	Yes, if MIB-based	Yes	No
Historical graphing	Yes	Yes	Yes	No
Available via the XML API	Yes	Yes	Yes	No
Collection default	Off	Off	On	Off

1.2.2 Performance statistics

Performance statistics counters record NE data, for example, physical equipment status and routing throughput, for monitoring and troubleshooting. Use the Performance Statistics Search Tool to view device-specific lists of the MIB-based performance statistics that the NFM-P supports.

Performance statistics, except for AA network performance statistics, are collected from NE MIBs using SNMP. The collection can be scheduled, or performed on demand.

You can graphically display performance statistics in real time, or export the data using the XML API. Historical performance statistics can be viewed in tabular or graphical form in the NFM-P client GUI.

1.2.3 Accounting statistics

Accounting statistics counters typically record service or subscriber usage data for billing or to ensure SLA compliance. Use the Accounting Statistics Search Tool to view device-specific lists of the accounting statistics that the NFM-P supports.

Accounting policies define the schedules that NEs use for the regular collection of accounting statistics. File policies specify how NEs store the accounting statistics data. An NE can perform scheduled accounting statistics collection only when a file policy and an accounting policy are deployed to the NE.

The NFM-P uses FTP or SCP to collect accounting statistics files from NEs. See the device documentation for information about accounting file creation and storage on a device.

You can view historical and real-time accounting statistics in tabular or graphical form in the NFM-P client GUI.

The supported accounting statistics types are:

- service accounting statistics—collected on each queue of each SAP that is associated with an accounting policy; provide queue throughput and drop information, and can be used for billing and SLA verification
- network accounting statistics—collected on each SDP queue or network port associated with an accounting policy; measure FC queue usage, which is of use for monitoring link utilization, identifying traffic patterns and trends, capacity planning, and traffic engineering

Note:

SDP statistics collection is supported only on devices that are in chassis mode B, C, or D.

- subscriber accounting statistics—collected on a subscriber profile for residential subscriber instances; used for billing and SLA verification
- application assurance, or AA, accounting statistics—collected from applications, application groups, and protocols on ISA-AA MDAs; AA network performance statistics are also available for monitoring the AA processing load on an ISA-AA MDA

See the *IPDR Reference* for comprehensive information about AA accounting and AA network performance statistics.

You can customize the statistics record in a service, subscriber, or AA accounting policy by specifying the counters that are collected and the data thresholds for collection. See the *NSP NFM-P Classic Management User Guide* for information about creating custom accounting records.

1.2.4 Flow statistics

Flow statistics measure the activity associated with traffic flows. A traffic flow is essentially a series of packets that have the same source address, destination address, and payload type, for example, a specific VoIP or social-networking application.

AA Cflowd statistics

ISA-AA groups support Cflowd sampling and TCP performance data collection for AA applications and application groups. NEs use the IPFIX protocol to send the statistics data in XDR-encoded binary records to one or more NSP Flow Collectors. An NSP Flow Collector can forward the statistics data to an OSS or third-party system for analysis or reporting purposes, and store the data in an NSP auxiliary database for use by NSP Analytics.

See the *NSP System Architecture Guide* for information about NSP Flow Collector and auxiliary database communication, and the *NSP Planning Guide* for deployment and scaling guidelines.

The *IPDR Reference* has comprehensive information about AA Cflowd and AA network performance statistics types and the associated counters.

1.2.5 Server performance statistics

You can collect server performance statistics such as the following from main and auxiliary servers in an NFM-P system:

- server statistics, which include memory usage and alarm counters
- network activity statistics, which include SNMP trap counters, accounting and SNMP polled statistic record counters, statistics collection counters, application assurance statistics collection counters, and NE resynchronization counters

You can collect server performance statistics on demand, or schedule regular collection using a policy. Server performance statistics can also be collected and graphically displayed in real time, or exported using the XML API. Historical server performance statistics can be viewed in tabular or graphical form in the NFM-P client GUI.

1.3 NFM-P statistics scalability

1.3.1 Overview

The overall volume of collected statistics is limited by the available storage resources. Therefore, the number of objects and collection frequency must be specified to meet operational requirements and remain within the storage constraints.

When a statistics collection policy or an accounting policy applies to a large number of objects, the collection interval must be sufficient to collect all of the data during the collection period. If the statistics collection time exceeds the collection interval, the NFM-P raises an alarm.

The maximum number of statistics the NFM-P can collect is specified in the *NSP Planning Guide*. To prevent statistics loss and NFM-P performance degradation, it is recommended that you stay within the specified maximum guidelines and ensure that each collection interval is sufficient for the number of statistics to be collected during the interval.

You can use NFM-P auxiliary servers to reduce the statistics collection load on a main server. Auxiliary server deployment is supported only in a distributed NFM-P system. See the *NSP Planning Guide* and *NSP System Architecture Guide* for information about auxiliary servers.

Real-time statistics collection and graphical display are supported for many network and service objects. An NFM-P client can open up to five real-time plotter windows, and each window can display up to eight counters. If multiple open statistics plotters point to the same NE or the same object on an NE, the NE receives a large number of SNMP requests. It is recommended that you limit the number of open statistics plotters for the same NE. You can use a span of control to limit operator access to specific NEs and thereby reduce the number of simultaneous real-time collections from a specific NE.



Note: The NFM-P limits the number of active MIB-based accounting statistics plots per NE to four among all NFM-P clients. For example, when two NFM-P clients each have two active


MIB-based accounting statistics plots associated with the same NE, no other client can open a plot of the same type for the NE.

1.3.2 Accounting statistics

If you need to collect large numbers of accounting statistics, for example, more than one million, it is recommended that you enable the statistics collection in a staggered manner. This helps to prevent the first-time collection delays associated with building the initial cache. For example, if you need to collect 6 million statistics in a 15-min interval, enable 500 000 statistics and wait until the interval completes before you add the next 500 000 statistics, and so on.

You can use custom accounting records to limit the processing requirements of service, subscriber, or AA accounting statistics collection. See [Chapter 4, “Accounting statistics”](#) for information about custom accounting records.

The NFM-P uses FTP or SCP to retrieve an NE statistics file when the NE notifies the NFM-P of the new file. The NFM-P retrieves and processes the files in the order that it receives the notifications. An NE retains files only for the period specified in the file policy; after this period, the NE deletes the files. The statistics data in the files is lost if the files are not retrieved from the NE during this time.

 **Note:** After an extended loss of connectivity to an NE that has an active accounting policy, the NFM-P resumes accounting statistics file retrieval by starting with the most recent statistics file.

When an accounting statistics counter is not supported by an NE or is excluded from collection using a custom record, the counter value is not included in a statistics record or displayed in the NFM-P client GUI.

1.3.3 Performance statistics

When a performance statistics collection interval is too short, statistics data may be lost. Performance statistics that remain to be collected at the end of a collection interval are skipped, and the next collection begins.

1.4 NFM-P statistics and OSS applications

1.4.1 Overview

The NFM-P can provide statistics data to OSS applications using the XML API. See the [Network Developer Portal](#) for information about using the XML API to retrieve statistics from the NFM-P.

Part II: Statistics collection configuration

Overview

Purpose

This volume describes statistics collection in the NFM-P.

Contents

Chapter 2, Statistics collection	19
Chapter 3, Performance statistics collection	29
Chapter 4, Accounting statistics	41
Chapter 5, Flow statistics collection	61
Chapter 6, Server performance statistics	87

2 Statistics collection

2.1 Statistics collection in the NFM-P

2.1.1 Overview

The NFM-P can be configured to collect statistics counters from managed Nokia NEs and NFM-P servers. Statistics collection requires the configuration and deployment of various policies.

Statistics policies can be configured in the following ways:

- for an entire network, or top-down, using the forms available through the Tools→Statistics menu option
- for a specific object, or bottom-up, using the Statistics tab of an object configuration form

An NFM-P operator can create statistics collection policies only for NEs that are within the span of control of the current NFM-P user. However, an operator can view statistics from NEs that are not within the current span of control.

By default, the NFM-P stores performance and accounting statistics data in the NFM-P database. If required, you can disable the NFM-P database storage of accounting statistics, performance statistics, or both, to prevent unnecessary database growth. For example, if you use an OSS client to retrieve statistics from the NFM-P, and do not require the NFM-P statistics presentation functions such as historical plotting, you may want to disable database storage for one or both statistics types. See the *NSP System Administrator Guide* for information about configuring NFM-P system preferences.

2.1.2 Using the NFM-P XML OSS interface to collect statistics

An OSS application can use the XML API to collect statistics. See the [Network Developer Portal](#) for information about using the XML API to transfer statistics records to an OSS client.

2.1.3 Dynamic custom statistics aggregators

Some functions, for example NSP Analytics reporting, may require aggregated statistics data. You can define and create custom aggregators for performance, accounting, AA accounting, and AA Cflowd statistics data, and apply them dynamically.

Note: The aggregated statistics data persists through NFM-P system upgrades.

To create one or more custom aggregators, you must define the aggregators in a file and apply the definitions. An aggregator definition file has the format shown in [Figure 2-1, “Custom aggregator definition format”](#) (p. 19).

Figure 2-1 Custom aggregator definition format

```
<aggregators>
  <aggregator aggregationName="AggregationName"
    description="Aggregation description"
```

```

        displayName="Aggregation name to display"
        aggregationKeys="key_1, key_2, ..."
        counters="counter1, counter2, ..."
        sumFields="counter2, counter4, ..."
        maxFields="counter2, counter4, ..."
        minFields="counter2, counter4, ..."
        varianceFields="counter6, counter7, ..."
        avgFields="counter2, counter7, ..."
        absAvgFields="counter4, counter4, ..."
        stddevFields="counter2, counter3, ..."
        absoluteValue="true or false"
        applyFunctions="function1,function2,..."
    <dataSource sourceName="Aggregation source table" sourceType="table
| class"/>
    </aggregator>
</aggregators>

```

Table 2-1, “Custom aggregator attributes” (p. 19) lists and describes the attributes in a custom aggregator definition.


 **Note:** The values in a list attribute are comma-separated.

Table 2-1 Custom aggregator attributes

Attribute	Description	Mandatory
aggregationName	Unique aggregator name; must not contain a space character The hourly, daily, weekly, and monthly aggregation table names have the following format: <i>aggregationName_timescale</i> where <i>timescale</i> is hour, day, week, or month	Yes
description	Aggregator description shown by the Aggregation Manager in the client GUI	Yes
displayName	Aggregator displayed name shown by the Aggregation Manager in the client GUI	Yes
aggregationKeys	Composite attribute that acts as a primary key for data aggregation; analogous to a 'group by' query function	Yes
counters	List of aggregation columns; the aggregation functions specified in the applyFunctions attribute are each applied to the counter values	No
sumFields	List of sum aggregation fields; sum aggregation is performed against each specified field	No
maxFields	List of max aggregation fields; max aggregation is performed against each specified field	No
minFields	List of min aggregation fields; min aggregation is performed against each specified field	No
varianceFields	List of variance aggregation fields; variance aggregation is performed against each specified field	No

Table 2-1 Custom aggregator attributes (continued)

Attribute	Description	Mandatory
avgFields	List of average aggregation fields; average aggregation is performed against each specified field	No
absAvgFields	List of absolute aggregation fields; absolute average aggregation is performed against each specified field	No
stddevFields	List of standard deviation aggregation fields; standard deviation aggregation is performed against each specified field	No
absoluteValue	Whether aggregation is to use a periodic data table as the data source; the value is true or false Setting the attribute to true indicates that the specified source table contains only absolute values. The aggregation is run against the associated <i>sourceName</i> _PERIODIC tables instead of against the raw data source specified in the XML file.	No; default is false if not specified
applyFunctions	List of aggregation functions to apply against the specified counters The supported values are sum, max, min, avg, absavg, stddev, and variance	No; default is sum if not specified
sourceName	Name of aggregation data source; statistics class name or auxiliary database table name	Yes
sourceType	Defines the type of data source; the value is table or class	Yes

Figure 2-2, “Table-sourced aggregator definition” (p. 21) shows the contents of a single-aggregator definition file in which an auxiliary database table is the data source.

Figure 2-2 Table-sourced aggregator definition

```
<aggregators>
  <aggregator aggregationName="analytics_cflowd_aa_wifi_tcp_ag"
    description="Measures WiFi TCP performance per app group"
    displayName="WiFi TCP Performance per Application Group
Aggregator"
    aggregationKeys="siteId, ipVers, subType, obsSubId,
observationSide, groupId, partId, appGrpName"
    counters="flowCount, avgFlowDuration"
    maxFields="maxCtd, maxStd"
    minFields="minCtd,maxStd, maxFlowBytesC2s" />
    <dataSource sourceName="analytics_cflowd_aa_wifi_tcp_ag"
sourceType="table"/>
  </aggregator>
</aggregators>
```

Figure 2-3, “Class-sourced aggregator definitions” (p. 22) shows the contents of a file that contains two aggregator definitions; each definition specifies a statistics class as the data source.

Figure 2-3 Class-sourced aggregator definitions

```
<aggregators>
  <aggregator aggregationName="acct_statistics_egress_ag"
    description="Complete Service Egress Packets"
    displayName="CompServEgPktAgg"
    aggregationKeys="monitoredobjectpointer,queueid,svcid,
sapid"
    counters="alloctetsforwarded, inprofilepktsoffered,
inprofilepktsforwarded, inprofilepktsdropped, allpktsdropped"
    absoluteValue="true"
    applyFunctions="max,min,avg" >
    <dataSource sourceName="service.CompleteServiceEgressPack-
etOctets" sourceType="class"/>
  </aggregator>
  <aggregator aggregationName="acct_stats_ingress_ag"
    description="Complete Service Ingress Packets"
    displayName="CompServIngPktAgg"
    aggregationKeys="monitoredobjectpointer,queueid,svcid,
sapid"
    counters="alloctetsforwarded, inprofilepktsoffered,
inprofilepktsforwarded, inprofilepktsdropped, allpktsdropped"
    absoluteValue="true"
    applyFunctions="max,min,avg" >
    <dataSource sourceName="service.CompleteServiceIngressPack-
etOctets" sourceType="class"/>
  </aggregator>
</aggregators>
```

Creating and managing aggregator definitions

After you create an aggregator definition file, you must use a main server configuration utility to apply the definitions in the file. Subsequently, you can modify and re-apply the file to update the existing definitions, or remove definitions. See [2.2 “To create and manage custom statistics aggregators” \(p. 23\)](#) for information.

When you apply a definition and an aggregator with the same name does not exist, a new aggregator is created. If an aggregator with the same name exists, any changes to the aggregator are applied to the existing aggregator.

After you apply a definition, you can view and manage the aggregator using the Aggregation Manager in the NFM-P client GUI.

2.2 To create and manage custom statistics aggregators

2.2.1 Steps

- 1

Create an aggregator definition file using the format described in [2.1.3 "Dynamic custom statistics aggregators" \(p. 19\)](#).
- 2

Log in to the standalone or primary main server as the nsp user.
- 3

Open a console window.
- 4

Navigate to the `/opt/nsp/nfmp/server/nms/bin` directory.
- 5

To apply new or modified definitions, enter the following:

```
bash$ ./nmserver.bash dynamic_aggregation create definition_file ↵
```


where *definition_file* is the absolute path and name of the aggregator definition file

The aggregator definition is applied.
- 6

To remove a definition, enter the following:

```
bash$ ./nmserver.bash dynamic_aggregation remove aggregationName ↵
```

where *aggregationName* is the aggregationName value in the aggregator definition

 **Note:** Only dynamically created aggregators can be removed.

The aggregator definition is removed.
- 7

Close the console window.
- 8

To view and manage the aggregators using the client GUI, choose Tools→Analytics→Aggregation Manager from the NFM-P main menu to open the Aggregation Manager form.

END OF STEPS

2.3 Statistics policies

2.3.1 Overview

You can configure the following policy types for statistics collection:

- accounting policy—specifies the accounting record type and collection interval
- file policy—specifies the storage criteria for accounting statistics files on NEs
- statistics policy—specifies the storage criteria for statistics in the NFM-P
- MIB statistics policy—specifies the collection of MIB-based statistics counters from managed NEs, and is one of the following:
 - NE MIB statistics policy—applies to all NE objects of the specified type
 - specific MIB statistics policy—applies to a specific NE object
- server performance statistics policy—specifies the collection criteria for statistics related to NFM-P server performance

Table 2-2 Statistics policies per statistic type

Statistics type	Policy				
	Accounting	File	Statistics	MIB statistics	Server performance statistics
Accounting <ul style="list-style-type: none"> • service • network • subscriber • AA 	✓	✓	✓	—	—
Performance	—	—	✓	✓	—
Server performance	—	—	✓	—	✓

2.3.2 Accounting policies

NEs collect accounting statistics using an accounting policy and an associated file policy that are assigned to a SAP, SDP, network port, or subscriber profile.

An accounting policy specifies an accounting statistics record type, a collection interval, an administrative state, and a file policy.

The NFM-P supports the following accounting policy types:

- service accounting policies—apply to SAPs, and specify the accounting records to collect for services
- network accounting policies—apply to network ports and SDPs, and specify the accounting records to collect for network resources

- subscriber accounting policies—apply to subscriber profiles and specify the accounting records to collect for residential subscribers
- AA accounting policies—apply to applications, application groups, and protocols, and specify the accounting records to collect for flows

An NE collects accounting statistics based on a specified collection interval and writes the statistics data in XML format to a file on the NE. After the rollover period specified in a file policy, the NE closes and compresses the file. The NE notifies the NFM-P that a new file is ready for processing, the NFM-P uses FTP or SCP to obtain the file from the NE, and adds the file contents to the NFM-P database. A third-party application can gather the statistics data and process it according to your specifications.

i Note: To conserve NE and NFM-P resources, Nokia recommends that you disable an accounting policy on an object when statistics for the object are not required.

When an accounting policy is administratively disabled, accounting statistics data is not written to a file on the NE. When the accounting policy is re-enabled, the new accounting data represents traffic activity since the re-enabling of the policy. Because the statistics data accumulates in rolling counters, no information is lost.

You can customize the record in a service, subscriber, or AA accounting policy by specifying the counters that are collected and the data thresholds for collection. See [Chapter 4, “Accounting statistics”](#) for more information.

Changes to an accounting policy that is disabled apply immediately to all objects to which the policy is applied. Changes to an accounting policy that is enabled take effect at the beginning of the next collection period.

Consider the following before you configure an accounting policy.

- An accounting policy requires a file policy. If you deploy an accounting policy in the absence of a file policy, a default file policy is automatically created.
- For service and network accounting policies, there can be one default policy. Default accounting policies are not in effect until they are distributed to the NEs by the NFM-P operator.
- There is a one-to-many relationship between accounting policies and accounting objects. For example, one service accounting policy can apply to many SAPs, but a SAP can have only one service accounting policy.

i Note: When you configure an accounting policy, Nokia recommends that you specify the same value for the Collection Interval parameter in the accounting policy and the Rollover parameter in the file policy. Failure to align these values may result in resource contention when a file rollover occurs.

2.3.3 File policies

A file policy specifies the relative size, storage location, and backup location of the files on the NE that contain accounting statistics data. An NE collects accounting statistics based on the collection interval specified for a SAP, network port, or subscriber in an accounting policy, and writes the statistics data in XML format to a file on the NE. After the rollover period specified in a file policy, the NE closes and compresses the file. The NE then notifies the NFM-P that a new file is ready for processing.

One file policy can be defined as a default policy and is automatically associated with an accounting policy if no file policy is specified.

i **Note:** When configuring both an Accounting Policy and a File Policy, Nokia recommends that the intervals for both policies are aligned. Failure to align these intervals may result in a resource contention at the file Rollover time.

2.3.4 Statistics policies

A statistics policy specifies a retention period and alarm thresholds for a statistics record. The retention period defines how long the NFM-P database retains the statistics record after collection, which affects the database storage requirements.

When statistics collection is enabled, the oldest statistics records are periodically removed from the database. When the collection interval is short and the retention period is long, more disk space is required to store the statistics data.

i **Note:** You can configure the length of time that the NFM-P database globally retains accounting statistics data. See the *NSP System Administrator Guide* for information about configuring the accounting statistics data retention period.

2.3.5 MIB statistics policies

The collection of performance statistics from NEs is controlled by MIB statistics policies that specify an administrative state, polling synchronization start time, and collection interval. There are two types of MIB statistics policies:

- NE MIB statistics policies
- specific MIB statistics policies

Each type of MIB statistics policy contains a list of MIB entry policies. A MIB entry policy defines the collection criteria for a specific MIB row. In an NE MIB policy, a MIB entry policy applies to all objects on the NE that use the MIB entry. In a specific MIB policy, the MIB entry policy applies only to the specified objects on the NE. For example, the `sapEgrQosPlcyQueueStatsEntry` MIB entry applies to L2 and L3 SAPs. In an NE MIB statistics policy, selecting this MIB entry enables collection of the entry on all L2 and L3 SAPs on the NE. In a specific MIB statistics policy, selecting this MIB entry enables collection only on the selected L2 or L3 SAPs.

If the statistics collection time required for a MIB entry exceeds the collection interval specified in the MIB statistics policy, the NFM-P raises an alarm. An NFM-P operator can change the polling interval for a statistics class to prevent this. See [Chapter 3, “Performance statistics collection”](#) for information about modifying polling criteria.

The collection of MIB statistics counters is disabled by default. See [Chapter 3, “Performance statistics collection”](#) for information about enabling performance statistics collection.

i **Note:** To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB policy to UP and the Polling Admin State parameter must be set to UP.

NE MIB statistics policies

An NE MIB statistics policy defines the global collection of specific statistics on specific NEs, for example, the port statistics from all ports in a group of NEs. In such a collection scenario,

configuring NE MIB policies rather than specific MIB policies is more efficient and uses fewer collection resources. Nokia recommends using NE MIB policies for general performance statistics collection on groups of NEs, and specific policies to enable or disable collection for specific NE objects.

After an NE MIB statistics policy is applied to an NE, statistics are collected for all objects on the NE, except for objects that have a specific MIB statistics policy. This prevents statistics from being collected twice—once by the NE MIB policy and once by the specific MIB policy.

Each NE requires an NE MIB statistics policy. The NFM-P has a default policy that it applies to an NE automatically when no NE MIB statistics policy is specified. The collection interval for each counter in the default policy is 15 min, and collection is disabled by default to conserve NE resources.

Specific MIB statistics policies

A specific MIB statistics policy defines the collection of selected performance statistics from specific objects on specific NEs to achieve a high statistics-collection granularity. For example, to collect port statistics at one rate for access ports and another rate for network ports, you can configure two policies that specify different collection intervals, and explicitly specify the access or network ports to which each policy applies. When a new access or network port is enabled, an NFM-P operator can add the port to the specific MIB statistics policy to enable statistics collection on the port.

The settings in a specific MIB statistics policy override the settings in an NE MIB statistics policy, and can be used to disable statistics collection for specific objects. For example, you enable statistics collection globally for an NE using an NE MIB statistics policy, and then disable the collection of specific statistics using a specific MIB statistics policy.

Distributing the collection of performance statistics

You can distribute the polling load on the NFM-P by creating multiple MIB statistics policies to define the collection of performance statistics across the network. This allows you to configure the Polling Synchronization Time parameter in each MIB statistic policy to have a different start time and to have the policy only poll for statistics from one NE type. For example, you can create two policies: one policy polling 7750 SR NEs at 00:00, and the second policy polling 7705 SAR NEs at 01:00.

2.3.6 Server performance statistics policies

Server performance statistics provide information about NFM-P server performance. You can specify the type of server statistics that are collected and the collection interval for each type.

Each type of server performance statistic has the following:

- a collection policy that specifies a synchronization time and a collection interval
- a statistics policy that specifies the NFM-P database retention period

3 Performance statistics collection

3.1 Collecting performance statistics

3.1.1 Overview

Performance statistics provide information about physical equipment, routing, and other NE properties for monitoring and troubleshooting purposes. See the Statistics Search Tool for the MIB-based performance statistics that the NFM-P supports.

Performance statistics collection is enabled using a MIB statistics collection policy and associating the policy with one or more NEs or specific objects within the NEs, for example, ports. For greater efficiency and collection granularity, the following types of MIB statistics policies are available:

- NE MIB statistics policies—define collection at the NE level
An NE MIB statistics policy contains a list of the MIB entry policies that are invoked for all objects on the NE.
- specific MIB statistics policies—define collection at the object level
A specific MIB statistics policy contains the same list of MIB entry policies as an NE MIB statistics policy, but the MIB entry policies are applied only to the objects specified in the specific MIB statistics policy. A specific MIB statistics policy uses the concept of a monitored class, which is the type of object on which to collect statistics, for example, a port or service site. The monitored object instances must be specified in the specific MIB statistics policy. For example, if port is selected as the monitored object, then the actual ports in the network to which the policy applies must be specified, and become part of the policy.

Before the NFM-P performs a statistics collection based on a MIB statistics policy, it checks for duplicate collection requests, such as when the same object is included in an NE MIB policy and a specific MIB policy. If a duplicate is found, the NFM-P performs the collection on the object based on the specific policy rather than on the NE policy.

For performance reasons, Nokia recommends using NE MIB policies to collect statistics from all instances of an object on an NE, and specific policies to enable or disable collection for specific NE objects. This is a much more efficient use of NE resources than using a specific policy in which each object instance is specified. For example, you can configure an NE MIB policy to collect OSPF routing statistics on all the routing instances of the NE, and use a specific MIB policy to collect statistics on a subset of the routing instances.

Specific and NE MIB policies can operate together to streamline statistics collection. For example, to collect network port statistics at 5-min intervals and access port statistics at 15-min intervals, you can create an NE MIB policy for all ports with a 15-min collection interval and create a specific policy for network ports with a 5-min collection interval. At every third interval, when the two policy activations coincide, the NE policy is used and duplicate collection is prevented.

NE MIB policies can be specified in NE discovery rules to ensure that statistics collection starts immediately after the NE is discovered. Specific MIB policies must be updated manually for specific objects on the new NE after it is discovered or new objects are created.

Specific MIB policy settings override NE MIB policy settings.

Note: To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB policy to UP and the Polling Admin State parameter must be set to UP.

Table 3-1 MIB statistics policy overrides

For a given object		
NE MIB policy	Specific MIB policy	Action
Collect	Collect	Statistic is collected once
Do not collect	Do not collect	Statistic is not collected
Collect	Do not collect	Statistic is not collected
Do not collect	Collect	Statistic is collected

For top-down performance statistics configuration, choose Tools→Statistics→MIB Policies from the NFM-P main menu to systematically configure statistics for the network. For bottom-up performance statistics configuration, configure the MIB entry policy from the Statistics tab of the properties form for a specific object.

3.1.2 GNE performance statistics support

The NFM-P supports the collection of a limited set of statistics counters from standard system, interface, and routing MIBs on GNEs. These statistics are processed and presented in the same manner as statistics from other devices. You can view GNE statistics on the Statistics tab of a GNE interface properties form, retrieve them using the XML API, and display them graphically using the NFM-P Statistics Plotter.

Note: If persistent SNMP indexes are not enabled on a GNE, one or more GNE interface indexes may change after a GNE reboots. This can cause a mismatch between the statistics records collected before the reboot and the current interface indexes. The NFM-P takes no action to identify or correct such a mismatch.

3.2 Pathway: performance statistics collection

3.2.1 Process

- Configure the MIB statistics policy for NEs. See [3.3 “To create or modify an NE MIB statistics policy using a top-down method” \(p. 31\)](#) for the top-down method. See [3.4 “To modify an NE MIB statistics policy using a bottom-up method” \(p. 32\)](#) for the bottom-up method.
- Configure the MIB statistics policy for specific objects. See [3.5 “To create or modify a specific MIB statistics policy using a top-down method” \(p. 33\)](#) for the top-down method. See [3.6 “To create or modify a specific MIB statistics policy using a bottom-up method” \(p. 35\)](#) for the bottom-up method.

-
- 3

Specify the MIB statistics policy polling interval. See [3.8 “To configure polling for a MIB statistics class” \(p. 37\)](#) for more information.
 - 4

Configure the statistics policy for an object. See [3.9 “To configure a statistics policy for MIB statistics” \(p. 38\)](#) for more information.
 - 5

If required, use an NFM-P client to view on-demand, scheduled, and real-time performance statistics. See [Chapter 7, “Statistics presentation”](#) for information about viewing statistics.
 - 6

Use the XML API to retrieve the performance statistics records from the NFM-P for processing by a third-party application. See the [Network Developer Portal](#) for information about using the XML API to transfer statistics records from the NFM-P database to an OSS client application.

3.3 To create or modify an NE MIB statistics policy using a top-down method

3.3.1 Steps


- 1

Choose Tools→Statistics→MIB Policies from the NFM-P main menu. The Manage MIB Statistics Policies form opens.
- 2

Choose NE MIB Statistics Policy (SNMP).
- 3

Specify a filter to create a filtered list of MIB statistics policies, if required, and click Search. A list of MIB statistics policies is displayed.
- 4

Click Create, or choose a MIB statistics policy from the list and click Properties. The NE MIB Statistics Policy (Create|Edit) form opens.

 **Note:** When you change a MIB statistics policy for an NE, the change affects all of the NEs to which the policy is assigned.
When you change a MIB statistics policy for a statistics class of an object, the change applies to all objects that use the same statistics class.
- 5

Configure the required parameters.



Note: To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB statistics policy to UP and the Polling Admin State parameter must be set to UP.

6

Click Apply. The form refreshes to display additional tabs.

7

Click on the Network Elements tab.

8

Click Assign Sites. A filter form opens.

9

Configure the filter criteria and click OK. The Assign *policy_name* form opens with a list of NEs displayed.

10

Select one or more NEs from the Unassigned Sites list, and click the right arrow. The selected NEs move to the Assigned Sites list.



Note: You can also assign an NE MIB policy to NEs using a discovery rule. Doing this does not affect the previously discovered NEs. See the *NSP NFM-P Classic Management User Guide* for information about device discovery.

11

Click OK. The Assign *policy_name* form closes and the NE MIB Statistics Policy form reappears.

12

Save your changes and close the forms.

END OF STEPS

3.4 To modify an NE MIB statistics policy using a bottom-up method

3.4.1 Steps

1

Choose Administration→Discovery Manager from the NFM-P main menu. The Discovery Manager form opens.

2

Click on the Managed State tab.

-
- 3

Choose an NE from the list and click Properties. The Node Discovery Control (Edit) form opens.
 - 4

Click on the MIB Statistics Policy tab.
 - 5

Click Select. The Configure MIB Statistics Policy form opens.
 - 6

Select a policy in the list and click OK. The Configure MIB Statistics Policy form closes and the Node Discovery Control form reappears.
 - 7

Save your changes and close the forms.

END OF STEPS

3.5 To create or modify a specific MIB statistics policy using a top-down method

3.5.1 Steps


- 1

Choose Tools→Statistics→MIB Policies from the NFM-P main menu. The Manage MIB Statistics Policies form opens.
- 2

Choose Specific MIB Statistics Policy (SNMP).
- 3

Specify a filter to create a filtered list of MIB statistics policies, if required, and click Search. A list of MIB statistics policies is displayed.
- 4

Click Create, or choose a MIB statistics policy from the list and click Properties. The Specific MIB Statistics Policy form (Create|Edit) opens.

 **Note:** When you change a MIB statistics policy for an object statistics class, the change applies to each object that uses the statistics class.
- 5

Configure the required parameters.



Note: To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB statistics policy to UP and the Polling Admin State parameter must be set to UP.

6

Click Select. The Specific Stats Polling Policy form opens.

7

Choose an object type from the list and click OK. The Specific Stats Polling Policy form closes and the Specific MIB Statistics Policy form reappears with the object type displayed in the Monitored Class Name field.

8

Click Apply. The Specific MIB Statistics Policy form refreshes to display additional tabs.

9

Click on the Monitored Objects tab.

10

Click Add. The Select *monitored_object* for Specific MIB Statistics Policy form opens.

11

Configure the filter criteria and click Search. A list of monitored objects is displayed.

12

Select one or more objects in the list and click OK. The Select *monitored_object* for Specific MIB Statistics Policy form closes and the Specific MIB Statistics Policy form reappears with the selected objects listed.

13

Click OK.

14


Save your changes and close the forms.


END OF STEPS

3.6 To create or modify a specific MIB statistics policy using a bottom-up method

3.6.1 Steps

- 1 _____
Choose an object on which to configure the MIB statistics policy.
- 2 _____
Right-click on the object and choose Properties. The properties form for the object opens.
- 3 _____
Click on the Statistics tab.
- 4 _____
Click Statistics Policies and choose Manage Specific MIB Policy. The Manage Specific MIB Policy form opens.
- 5 _____
Choose Specific MIB Statistics Policy (SNMP) from the object drop-down menu.
- 6 _____
Click Create or choose a MIB statistics policy and click Properties. The Specific MIB Statistics Policy (Create|Edit) form opens.

 **Note:** When you change a MIB statistics policy for an object statistics class, the change applies to each object that uses the statistics class.
- 7 _____
Configure the required parameters.

 **Note:** To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB statistics policy to UP and the Polling Admin State parameter must be set to UP.
- 8 _____
Click Select. The Specific Stats Polling Policy form opens.
- 9 _____
Choose an object type from the list and click OK. The Specific Stats Polling Policy form closes and the Specific MIB Statistics Policy form reappears with the object type displayed in the Monitored Class Name field.

-
- 10

Click Apply. The Specific MIB Statistics Policy form refreshes to display additional tabs.
 - 11

Click on the Monitored Objects tab.
 - 12

Click Add. The Select *monitored_object* for Specific MIB Statistics Policy form opens.
 - 13

Configure the filter criteria and click Search. A list of monitored objects is displayed.
 - 14

Select one or more objects in the list and click OK. The Select *monitored_object* for Specific MIB Statistics Policy form closes and the Specific MIB Statistics Policy form reappears with the selected objects displayed.
 - 15

Save your changes and close the forms.

END OF STEPS

3.7 To assign the default 1830 VWM OSU performance management policy to 1830 devices

3.7.1 Steps

- 1

Choose Tools→Statistics→VWM OSU Performance Management Policies from the NFM-P main menu. The Performance Management Policies form opens.
- 2

Click Search, choose the default policy, and click Properties. The Performance Management Policy - 1 (Edit) form opens.
- 3

Configure the required parameters.
- 4

Click on the Elements tab and click Assign 1830 VWM OSUs. The Assign Filter form open.
- 5

Configure the filter criteria as required in the Assign form.

-
- 6

Click Apply to deploy the performance management policy to the assigned devices.
 - 7

Save your changes and close the forms.

END OF STEPS

3.8 To configure polling for a MIB statistics class

3.8.1 Steps


- 1


Choose an object on which to configure the MIB statistics polling interval.
- 2

Right-click on the object and choose Properties. The properties form for the object opens.
- 3

Click on the Statistics tab.
- 4

Choose a statistics class from the object drop-down menu.
- 5

Click Statistics Policies and choose MIB Entry Policy. The MIB Entry Policy form opens.
 **Note:** The MIB Entry Policy menu item is dimmed if you choose an invalid performance statistics class.
- 6

Configure the required parameters.
 **Note:** To view statistics on the NFM-P GUI, you must set the Administrative State parameter of the MIB statistics policy to UP and the Polling Admin State parameter must be set to UP.
- 7

Save your changes and close the forms.

END OF STEPS

3.9 To configure a statistics policy for MIB statistics

3.9.1 Steps

- 1 _____
Choose the object on which to configure the statistics policy.
- 2 _____
Right-click on the object and choose Properties. The properties form for the object opens.
- 3 _____
Click on the Statistics tab.
- 4 _____
Choose a statistics class from the object drop-down menu.
- 5 _____
Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.
- 6 _____
Configure the required parameters.
- 7 _____
Click on the Thresholds tab. The Thresholds tab contains a threshold parameter for each counter in the statistics class. When a statistics counter threshold is exceeded, the NFM-P raises a threshold-crossing alarm.
- 8 _____
Configure one or more threshold parameters by selecting the check box beside a parameter and then entering a threshold value.
- 9 _____
Save your changes and close the forms.

END OF STEPS _____

3.10 To delete performance statistics records

3.10.1 Steps

- 1 _____
Open the properties form of the object that has the statistics records to delete.

-
- 2

Click on the Statistics tab.
 - 3

Choose a statistics class from the object drop-down menu.
 - 4

Click Statistics Policy. The Statistics Policy form opens.
 - 5

Click Purge Statistics Records. The Statistics Policy filter form opens.
 - 6

Configure the filter and click OK. A confirmation message is displayed.
 - 7

Click Yes. The statistics records are deleted.
 - 8

Close the open forms.

END OF STEPS

4 Accounting statistics

4.1 Accounting statistics collection

4.1.1 Description

Accounting statistics provide packet and octet throughput information for queues that are associated with the following objects:

- SAPs or SDPs, which provide service accounting statistics
- network ports, which provide network accounting statistics
- subscriber profiles, which provide subscriber accounting statistics
- subscribers, SAPs, and spoke SDP bindings, which provide application assurance, or AA, accounting statistics

See the *NSP NFM-P Classic Management User Guide* for information about AA accounting statistics. See the Accounting statistics volume in this guide for lists of the MIB-based accounting statistics that the NFM-P supports.

To collect accounting statistics, you need to create and apply the following:

- an accounting policy
- a file policy
- a statistics policy

See [Chapter 2, “Statistics collection”](#) for a description of each policy type.



Note: Not all managed NEs collect accounting statistics. See the Accounting statistics volume of this guide for information about NEs that collect accounting statistics.

OSS retrieval of accounting statistics by policy

Multiple OSS applications may collect accounting statistics from the same object using different policies. To enable the OSS filtering and retrieval of statistics associated with a specific policy, you can configure the NFM-P to include the accounting policy ID in the header of each accounting statistics file.

When the accounting policy ID is included in a statistics file, an OSS client can use XML API registerLogToFile and JMS filters to retrieve only the statistics files associated with the policy.

See the system preferences configuration procedure in the *NSP System Administrator Guide* for information.

4.1.2 Queue filters

NOTICE

Service Disruption

Modifying the server configuration can have serious consequences that can include service disruption.

Contact technical support before you attempt to modify a server configuration file.

You can use the NFM-P to configure queue filters for accounting statistics so that NEs use queues that do not have to be monitored. The NEs generate accounting data files that include all of the configured queues. When you configure a queue filter, the queue of each statistics record is checked; the queues that do not match the filter are discarded and are not processed.

You can specify a queue filter in the NFM-P main server configuration file. For example, you can enable the processing of queues 1 and 2 for combinedServiceEgressOctets and completeSubscriberIngressPacketOctets, and queues 1, 2, and 3 for completeSubscriberEgressPacketOctets:

```
<accountingStatsFilter  
combinedServiceEgressOctets="1,2"  
completeSubscriberEgressPacketOctets="1,2,3"  
completeSubscriberIngressPacketOctets="1,2" />
```

4.1.3 Custom accounting records

You can customize the record in an accounting policy to contain only the accounting data that you require. This reduces the statistics-collection processing load and the volume of data that collection generates. Custom accounting records allow you to specify the following in an accounting policy:

- the counters to include in each record, for example, packet but not octet counters
- the queues or policers to monitor for collection
- significant-change criteria that define when an NE saves the data to a file:
 - significant change threshold
 - reference counters
 - reference queues and policers
- HSMDA scheduler override counters and queues

The following accounting policy types support custom records:

- Custom Record Service
- Custom Record Subscriber
- Custom Record AA Subscriber



Note: In a Custom Record AA Subscriber accounting policy, you can specify optional attributes to include in each record; for example, the application profile or ASO associated with an AA subscriber.

4.1.4 Dynamic periodic accounting statistics calculations

You can define periodic value calculations for accounting statistics data and dynamically apply them. Periodic values, which represent the difference between the current and previous counter values, are available for functions such as reporting by the NSP Analytics.

i Note: The periodic data values persist through NFM-P system upgrades.

i Note: The periodic data values are not retrievable using the NFM-P XML API.

The dynamic calculation definitions for one or more statistics classes must be specified in an XML file that has the format shown in [Figure 4-1, “Dynamic periodic calculation definition format” \(p. 42\)](#).

Figure 4-1 Dynamic periodic calculation definition format

```
<classes>
  <className name="accounting statistics class name"
    periodicFields="counters for which to calculate periodic
values"
    dataFields="list of additional fields to include in
periodic tables"/>
</classes>
```

i Note: The dataFields entry is optional, and specifies data fields to include other than counters that are required for reports or data aggregations; for example, a SAP or port ID.

[Figure 4-2, “Dynamic periodic definition example” \(p. 42\)](#) shows a periodic calculation definition example that defines periodic values to calculate for four statistics classes.

Figure 4-2 Dynamic periodic definition example

```
<classes>
  <className name="service.CompleteServiceEgressPacketOctets"
    periodicFields="inProfilePktsOffered,inProfilePktsForwarded,inProfileP-
ktsDropped"
    dataFields="svcId,sapId,portId"/>
  <className name="service.CompleteServiceIngressPacketOctets"
    periodicFields="inProfilePktsOffered,inProfilePktsForwarded,inProfileP-
ktsDropped "
    dataFields="svcId,sapId,portId"/>
  <className name="service.ServiceEgressOctets"
    periodicFields="allOctetsForwarded,allOctetsDropped,allOctetsOf-
fered"
    dataFields="svcId,sapId"/>
  <className name="service.ServiceIngressOctets"
    periodicFields="allOctetsForwarded,allOctetsDropped"
    dataFields="svcId,sapId"/>
</classes>
```

The following sample periodic calculation file is available on an NFM-P main server:

/opt/nsp/nfmp/server/nms/config/accountingPeriodicSample.xml

Creating and managing definitions

After you create a definition file, you must use a main server configuration utility to apply the definitions in the file. Subsequently, you can modify and re-apply the file to update the existing definitions, or remove definitions. See [4.2 “To create and manage periodic accounting statistics calculations” \(p. 43\)](#) for information.



Note: If you re-apply a definition for a statistics class that has a statically defined periodic calculation, any additional columns in the definition are added, but no statically defined columns are removed.

You can view and manage static and dynamically defined periodic calculations using the Periodic Counters Manager in the client GUI.

4.2 To create and manage periodic accounting statistics calculations

4.2.1 Steps

1

Create a periodic counter definition file using the format described in [4.1.4 “Dynamic periodic accounting statistics calculations” \(p. 43\)](#).

2

Log in to the standalone or primary main server as the nsp user.

3

Open a console window.

4

Enter the following:

```
bash$ cd /opt/nsp/nfmp/server/nms/bin ↵
```

5

To apply new or modified definitions, enter the following:

```
bash$ ./nmserver.bash accountingPeriodic create definition_file ↵
```

where *definition_file* is the absolute path and name of the periodic counter definition file

The periodic counter definition is applied.

6

To remove a definition, enter the following:

```
bash$ ./nmserver.bash accountingPeriodic remove class_name ↵
```

where *class_name* is the class of a periodic counter definition

The periodic counter definition is removed.

7

Close the console window.

8

To view and manage the periodic counter definitions using the client GUI, choose Tools→Analytics→Periodic Counters Manager from the NFM-P main menu.

END OF STEPS

4.3 Pathway: accounting statistics collection

4.3.1 Process

1

Ensure that FTP is enabled on the NEs from which you want to collect accounting statistics. See the *NSP NFM-P Classic Management User Guide* for information about device commissioning and enabling FTP.

2

Configure polling policies to enable NFM-P FTP or SCP access to the NEs, as required. See the *NSP NFM-P Classic Management User Guide* for information about configuring polling policies.

3

If required by OSS applications

4

Enable statistics collection on the object. See the *NSP NFM-P Classic Management User Guide* for information about enabling statistics collection on an object.

5

Configure a file policy with rollover and retention values that are appropriate for the traffic volume, number of objects, and available file storage space. See [4.4 “To configure a file policy” \(p. 47\)](#).

6

Configure an accounting policy. See [4.5 “To configure an accounting policy” \(p. 48\)](#).

7



Note: It is recommended to apply an accounting policy to an object, rather than distribute the policy to an NE. The NFM-P distributes the policy to the NE when the policy is applied to the object.

Apply the accounting policy. You can apply one accounting policy to a SAP, SDP, network object such as a port, or a subscriber profile.

- a. For service accounting statistics, apply the service accounting policy to a SAP or SDP.
 1. Administratively disable accounting on the SAP.
 2. Choose the newly created accounting policy.
- b. For network accounting statistics, apply the accounting policy to a network object such as a port. You can specify a network accounting policy when you configure a port in network mode.
 1. Administratively disable accounting for the network object.
 2. Choose the newly created accounting policy.
- c. For subscriber accounting statistics, apply the accounting policy to a subscriber profile.
 1. Administratively disable accounting on the subscriber.
 2. Choose the newly created accounting policy.
- d. For AA accounting statistics, apply the accounting policy to an AA group or partition.

See the *NSP NFM-P Classic Management User Guide* for information about applying an accounting policy to an object.

8

Configure a statistics policy to define the retention requirements for the accounting statistics records.

- a. For service accounting statistics, see [4.6 “To configure a statistics policy for accounting statistics on a SAP or an SDP” \(p. 50\)](#).
- b. For network accounting statistics, see [4.7 “To configure a statistics policy for accounting statistics on a network interface” \(p. 52\)](#).
- c. For subscriber accounting statistics, see [4.8 “To configure a statistics policy for accounting statistics on a subscriber” \(p. 53\)](#).
- d. For AA accounting statistics per subscriber, see [4.9 “To configure a statistics policy for AA accounting statistics on a subscriber” \(p. 54\)](#).
- e. For AA accounting statistics for an application, see [4.10 “To configure a statistics policy for an AA accounting statistics application” \(p. 55\)](#).
- f. For AA accounting statistics for an application group, see [4.10 “To configure a statistics policy for an AA accounting statistics application” \(p. 55\)](#).
- g. For AA accounting statistics for a protocol, see [4.12 “To configure a statistics policy for an AA accounting statistics protocol” \(p. 57\)](#).

9

View the status of accounting statistics collection for NEs. See the *NSP System Administrator Guide* for more information.

10

If required, use the NFM-P client to view accounting statistics. See [Chapter 7, “Statistics presentation”](#) for more information.

11

Use the XML API to retrieve the accounting statistics records from the NFM-P for processing by a third-party application.

See the [Network Developer Portal](#) for information about using the XML API to transfer statistics records from the NFM-P database to an OSS client application.

12

As required, create or manage periodic accounting statistics calculations; see [4.2 “To create and manage periodic accounting statistics calculations”](#) (p. 44).

4.4 To configure a file policy

4.4.1 Purpose

Perform this procedure to configure a file policy that controls the management of accounting statistics files on one or more NEs.



Note: You cannot modify or delete a file policy when the policy is associated with an accounting policy that is administratively enabled. You must set the Administrative State parameter of the corresponding accounting policy to Down before you can modify or delete the file policy. See [4.5 “To configure an accounting policy”](#) (p. 48) for more information.

4.4.2 Steps

1

Choose Tools→Statistics→File Policies from the NFM-P main menu. The File Policies form opens.

2

Click Create, or choose an existing file policy and click Properties. The File Policy (Create) form opens.

3

Configure the parameters.

i **Note:** Ensure that the NE resources are sufficient to support the file policy and associated accounting policy specifications. The collection, retention, and rollover intervals must be appropriate, and the statistics must be regularly retrieved from the NEs.

i **Note:** Some 7210 SAS devices are equipped with USB flash memory (uf1). In the equipment tree and on Alarm Info forms, the NFM-P displays only Flash Memory. The display does not distinguish between compact flash memory (cf) and USB flash memory (uf).

4 _____
Click Apply.

5 _____
Distribute the policy to NEs, as required.

6 _____
Close the open forms.

END OF STEPS _____

4.5 To configure an accounting policy

4.5.1 Steps

1 _____
Choose Tools→Statistics→Accounting Policies from the NFM-P main menu. The Manage Accounting Policies form opens.

2 _____
Click Create, or choose an existing accounting policy and click Properties. The Accounting Policy (Create) form opens.

3 _____
Specify a file policy.

i **Note:** If you do not specify a file policy, a file policy with default values is automatically created and associated with the accounting policy. When a file policy is automatically created, you must verify that the file policy settings are appropriate for the accounting policy.

4 _____
Configure the required parameters.

5 Click Apply. The form displays additional tabs.

6 If the policy type is not a Custom Record policy type, go to [Step 14](#) .

7 Click on the Custom Record tab. The Significant Change Criteria tab is displayed.

8 Configure the significant change criteria, if required.

1. Configure the Significant Change Delta parameter.
2. In the Reference Queue panel, select the All Queues parameter to monitor all queues for significant change, or click Select to choose one queue.
3. In the Reference Queue panel, select the Ingress Counters and Egress Counters to monitor.
4. In the Reference Override panel, select the All Overrides parameter to monitor all override counters for significant change, or click Select to choose one override counter.
5. In the Reference Override panel, select the Ingress Counters and Egress Counters to monitor.
6. In the Reference Policers panel, select the All Policers parameter to monitor all policer counters for significant change, or click Select to choose one policer counter.
7. In the Reference Policers panel, select the Ingress Counters and Egress Counters to monitor.

9 If the policy type is Custom Record AA Subscriber, specify the AA counters to monitor.

1. Click on the Application Assurance tab.
2. Select one or more Application Assurance, To Subscriber, and From Subscriber counters to monitor.
3. Specify the Optional Attributes to include in the record.

10 Configure queue counters to monitor, if required.

1. Click on the Queue Counter Config tab.
2. Click Create. The CustomQueue Config (Create) form opens.
3. Configure the ID parameter to specify a queue.
4. Specify one or more Ingress and Egress counters to monitor.
5. Click OK to save your changes and close the form.

11

Configure override counters to monitor, if required.

1. Click on the Override Counter Config tab.
2. Click Create. The Custom Override Config (Create) form opens.
3. Configure the ID parameter to specify an override.
4. Specify one or more Ingress and Egress counters to monitor.
5. Click OK to save your changes and close the form.

12

Configure policer counters to monitor, if required.

1. Click on the Policer Counter Config tab.
2. Click Create. The Policer (Create) form opens.
3. Configure the ID parameter to specify a policer.
4. Specify one or more Ingress and Egress counters to monitor.
5. Click OK to save your changes and close the form.

13

Click Apply.

14

Distribute the policy to NEs, as required.

15

Close the open forms.

END OF STEPS

4.6 To configure a statistics policy for accounting statistics on a SAP or an SDP

4.6.1 Steps

1

Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.

2

Specify a filter to create a filtered list of services and click Search. A list of services is displayed.

3

Select a service in the list and click Properties. The service properties form opens.

4

Click on one of the following tabs:



Note: The tabs that are displayed depend on the type of service that is chosen.

- L2 Access Interfaces
- L3 Access Interfaces
- Mesh SDP Bindings
- Spoke SDP Bindings

The properties form for the object opens.

5

Click on the Statistics tab.

6

Choose a statistics class from the object drop-down menu.

7

Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.

8

Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected statistics class type are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.

9

Click on the Thresholds tab. The Thresholds tab contains a threshold parameter for each counter in the statistics class. When a statistics counter threshold is exceeded, the NFM-P raises a threshold-crossing alarm.

10

Configure one or more threshold parameters by selecting the check box beside a parameter and then entering a threshold value.

11

Click OK.

12

Save your changes and close the forms.

END OF STEPS

4.7 To configure a statistics policy for accounting statistics on a network interface

4.7.1 Steps

- 1 _____
In the navigation tree routing view, navigate to the required network interface.
- 2 _____
Right-click on the interface icon and choose Properties. The Network Interface (Edit) form opens.
- 3 _____
Click on the Statistics tab.
- 4 _____
Choose a statistics class from the object drop-down menu.
- 5 _____
Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.
- 6 _____
Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected statistics class type are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.
- 7 _____
Click on the Thresholds tab. The Thresholds tab contains a threshold parameter for each counter in the statistics class.
- 8 _____
Configure the threshold parameters. When a statistics counter threshold is exceeded, the NFM-P raises a threshold-crossing alarm.
- 9 _____
Save your changes and close the forms.

END OF STEPS _____

4.8 To configure a statistics policy for accounting statistics on a subscriber

4.8.1 Steps

- 1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.
- 2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search. A list of subscriber instances is displayed.
- 3

Select a subscriber instance in the list and click Properties. The Residential Subscriber Instance form opens.
- 4

Click on the Statistics tab.
- 5

Choose a statistics class from the object drop-down menu.
- 6

Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.
- 7

Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected statistics class type are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.
- 8

Click on the Thresholds tab. The Thresholds tab contains a threshold parameter for each counter in the statistics class.
- 9

Configure the threshold parameters. When a statistics counter threshold is exceeded, the NFM-P raises a threshold-crossing alarm.

10

Click OK.

END OF STEPS

4.9 To configure a statistics policy for AA accounting statistics on a subscriber

4.9.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search. A list of subscriber instances is displayed.

3

Select a subscriber instance in the list and click Properties. The Residential Subscriber Instance form opens.

4

Click on the Statistics tab.

5

Choose an AA statistics class from the object drop-down menu.

6

Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.

7

Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected statistics class type are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.

8

Click OK.

-
- 9 _____
Save your changes and close the forms.

END OF STEPS _____

4.10 To configure a statistics policy for an AA accounting statistics application

4.10.1 Steps

- 1 _____
Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.
- 2 _____
Choose AA Group Policy and click Search. A list of application group policies is displayed.
- 3 _____
Select a policy in the list and click Properties. The global AA Group Policy (Edit) form opens.
- 4 _____
Click on the Local Definitions tab.
- 5 _____
Select a local policy definition in the list and click Properties. The local AA Group Policy (Edit) form opens.
- 6 _____
Click on the Applications tab.
- 7 _____
Select an application in the list and click Properties. The Application (Edit) form opens.
- 8 _____
Click on the Statistics tab.
- 9 _____
Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.
- 10 _____
Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected application are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.

11

Click OK.

12

Save your changes and close the forms.

END OF STEPS

4.11 To configure a statistics policy for an AA accounting statistics application group

4.11.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose AA Group Policy and click Search. A list of application group policies is displayed.

3

Select a policy in the list and click Properties. The global AA Group Policy (Edit) form opens.

4

Click on the Local Definitions tab.

5

Select a local policy definition in the list and click Properties. The local AA Group Policy (Edit) form opens.

6

Click on the Application Groups tab.

7

Select an application group in the list and click Properties. The Application Group (Edit) form opens.

8 _____
Click on the Statistics tab.

9 _____
Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.

10 _____
Configure the required parameters.
If the Administrative State parameter is set to Down, accounting statistics for the selected application group are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.

11 _____
Click OK.

12 _____
Save your changes and close the forms.

END OF STEPS _____

4.12 To configure a statistics policy for an AA accounting statistics protocol

4.12.1 Steps

1 _____
In the navigation tree equipment view, expand *device*→ISA-AA Groups→*ISA_group*.

2 _____
Right-click on the ISA-AA Group icon and choose Properties. The ISA-AA Group (Edit) form opens.

3 _____
Click on the Statistics tab.

4 _____
Choose AA Protocol Stats (Application Assurance) from the object drop-down menu. A list of AA protocols is displayed.

5 _____
Select a protocol in the list.

6 Click Statistics Policies and choose Statistics Policy. The Statistics Policy form opens.

7 Configure the required parameters.

If the Administrative State parameter is set to Down, accounting statistics for the selected protocol are not stored in the NFM-P database. The statistics cannot be viewed in the NFM-P GUI or exported to an NFM-P XML API client using find or findToFile operations. However, if an NFM-P XML API client registers using registerLogToFile, the accounting statistics data is exported to files.

8 Click OK.

9 Save your changes and close the forms.

END OF STEPS

4.13 To configure ingress and egress accounting policies for an LSP

4.13.1 Purpose

Perform this procedure to configure accounting policies for an LSP. Collecting statistics on an LSP requires configuring an egress accounting policy on the LSP and a separate ingress accounting policy on the MPLS instance at the LSP destination.

4.13.2 Steps

1 Configure a file policy, as described in [4.4 "To configure a file policy" \(p. 47\)](#).

2 Create an accounting policy for the ingress statistics, as described in [4.5 "To configure an accounting policy" \(p. 48\)](#). Configure the Type parameter of the policy to Combined MPLS LSP Ingress, and use the file policy you created in [Step 1](#).

3 Create a second accounting policy for the egress statistics. Configure the Type parameter to Combined MPLS LSP Egress, and use the file policy you created in [Step 1](#).

4 From the routing view, expand NE→Routing Instance→MPLS.

-
- 5

Right-click on the MPLS instance and choose Properties. The MPLS (edit) form opens.
 - 6

Click on the Accounting tab.
 - 7

Click Create. The IngStatsPolicy, Routing Instance (Create) form opens.
 - 8

Configure the required parameters in the LSP Information panel.
 - 9

Configure the required parameters in the Ingress Accounting Statistics panel.
 - 10

Click OK. Close the MPLS form.
 - 11

Navigate to the LSP.
 - 12

Right-click on the LSP and choose Properties. The LSP (edit) form opens.
 - 13

Click on the Accounting tab.
 - 14

Configure the required parameters in the Egress Accounting Statistics panel.
The Ingress Accounting Statistics panel displays a read-only view of the ingress accounting policy you configured on the MPLS router.
 - 15

Click OK.
 - 16

Save your changes and close the forms.

END OF STEPS

4.14 To delete accounting statistics records

4.14.1 Steps

- 1 _____
Open the properties form of the object that has the statistics records to delete.
- 2 _____
Click on the Statistics tab.
- 3 _____
Choose a statistics class from the object drop-down menu.
- 4 _____
Click Statistics Policy. The Statistics Policy form opens.
- 5 _____
Click Purge Statistics Records. The Statistics Policy filter form opens.
- 6 _____
Configure the filter and click OK. A confirmation message is displayed.
- 7 _____
Click Yes. The statistics records are deleted.
- 8 _____
Close the open forms.

END OF STEPS _____

5 Flow statistics collection

5.1 Flow statistics collection in the NFM-P

5.1.1 Introduction

The NFM-P can collect flow statistics from managed NEs. Although other criteria may apply, a flow is basically a series of IP packets that share a common source, destination, and type of payload, for example, traffic that is specific to an application.

When an NE is configured to collect flow statistics, the NE monitors the traffic on one or more interfaces to identify the flows, aggregates the flow data, and regularly exports flow data records to an external system.

The NFM-P uses NSP Flow Collectors to collect the IPFIX v10 (RFC 7011) flow statistics with the AA Cflowd Nokia variant from managed NEs.



Note: For information about using NSP to collect IPFIX v10 statistics with the System Cflowd or CGNAT Nokia variants, or Netflow v5 statistics, see the *NSP Data Collection and Analysis Guide*.

5.1.2 Functional description

IPFIX data is sent from routers, packaged in IPFIX messages that contain IPFIX records. Each record is a collection of fields determined by an IPFIX template. These templates are periodically sent by the routers to the collectors. IPFIX templates can be modified (fields can be removed from or added to the records). A flow collector uses these templates to decode the IPFIX records.

The fields contained in a record may include - but are not limited to - the following:

- counters
- timestamps
- delays
- quantities
- labels
- IDs

The Flow Collector generates "data types" using sets of counters that are found in IPFIX records and/or values computed based on IPFIX records. The Flow Collector collects these data types at configured intervals, then exports the collected data in various formats (IPDR binary files, CSV text files, Kafka notifications, Vertica database records).

Each data type can be aggregated in one or more ways by applying an "aggregation type" to it. Each aggregation type is defined by a set of fields from the IPFIX record (usually IDs) and, potentially, correlated information from the NFM-P database. When no aggregation is applied, the

Flow Collector refers to the data as "raw" (raw aggregation). During each interval, aggregated data is collected in a memory buffer, and is exported in one of the supported formats when the interval ends.

Collected data can be exported in the following formats:

- AA Cflowd
 - Vertica database
 - IPDR binary files
 - CSV text files (for the Comprehensive Raw data types)

5.1.3 IPDR Reference

The NFM-P IPDR Reference lists and describes the fields in an IPDR-formatted record for each collection or aggregation domain based on user-specified criteria. The IPDR Reference also lists the IPDR schema changes between NFM-P releases.

5.1.4 Exporting IPDR file contents to other formats

For compatibility with an OSS application, you can use an NFM-P tool to export the contents of XDR-encoded IPDR statistics files to files in a format such as text or XML. Because the export process is resource-intensive, you must observe the following requirements before you attempt to use the tool.

- It is strongly recommended that you use the tool only on a standby NFM-P main server, whether collocated with a main database or in a distributed deployment.
- The system that receives the exported files must meet the following requirements.
 - The system must be dedicated to the export operation.
 - The output location must have sufficient space for the output files; a large network-attached storage system is recommended. An exported file is approximately six times the size of the associated input file.
- The user of the tool requires full access permissions on the input and output directories.

CLI- and GUI-based versions of the tool are available; see [5.13 "To export IPDR statistics files to a different format" \(p. 83\)](#) for information about how to export IPDR files.

5.2 Pathway: configure flow statistics collection

5.2.1 Stages

1

Specify the data and aggregation types that must be turned on in the aggregation policy. Turn on only what is needed to minimize resource consumption (CPU, memory, network bandwidth, disk IO, and database storage space) and specify the collection interval length for each data type. Note that the size of the interval has an impact on memory resources (longer intervals consume more memory). See [5.4 "To configure flow statistics aggregation" \(p. 64\)](#).

2

Specify the destination of the collected data. For some data types, the destination is

predetermined. For others it can be selected. Specify the upload target for each file format in the upload policy, and for IPDR and CSV file formats, configure FTP or SFTP server upload targets. See [5.5 “To set upload targets and policies”](#) (p. 69).

3

For AA Cflowd collection, specify the extraction and transfer of correlation data from NFM-P. See [5.6 “To configure the transfer of results files”](#) (p. 71).

4

Enable or disable AA collection for specific subscriber types, and specify Application and Protocol filters for UAPPS data type. See [5.7 “To configure AA filters”](#) (p. 73).

5

Configure special-study filters for Volume, TCP Performance, and Comprehensive AA data types. See [5.8 “To configure an AA Cflowd special-study policy”](#) (p. 75).

6

Configure subnet summarization for Volume and TCP Performance Business data types. See [5.9 “To configure AA subnet summarization filters”](#) (p. 78).

7

Specify the export formats for each AA data domain. See [5.10 “To configure the AA flow data persistence”](#) (p. 80).

8

Add routers to the collection policy, specifying the management IP (source of UDP data), the system IP (the collector translates the management IP to the system ID, which must be IPv4), and the data protocol. See [5.11 “To specify the NEs for flow statistics collection”](#) (p. 81).

9

Configure the routers (NEs) to collect desired flow information (according to their capabilities) and send the collected flow information to the NSP Flow Forwarder Service. The Flow Forwarder exposes a NSP VIP that has to be used by the routers as the target for the IPFIX data. The Flow Forwarder uses specific UDP ports for each data type and the routers have to be configured to send the data to the proper UDP port. The UDP port use for the AA Cflowd data type is 4740.

In DR deployments, NEs must be configured to send flows to both the primary and standby sites.

10

As required, refresh the NFM-P network data snapshot for AA Cflowd reporting. See [5.12 “To extract an NFM-P network data snapshot”](#) (p. 82).

5.3 Pathway: configure AA Cflowd special study statistics collection

5.3.1 Purpose

Use this pathway along with [5.2 “Pathway: configure flow statistics collection” \(p. 62\)](#) to configure collection of Special Study AA Cflowd statistics. Special study configuration is required for DNS Performance, VoIP OTT, or OTT Video reports in the NSP Analytics. See the *NSP Analytics Report Catalog* for more information about Analytics reports.

5.3.2 Stages

- 1

Select the Enable Application Performance Reporting for Residential Groups parameter in each AA group policy associated with the special-study subscribers.

Note: The parameter must be selected if special study statistics are to be collected for any subscriber type.
- 2

Configure a special-study policy; see [5.8 “To configure an AA Cflowd special-study policy” \(p. 75\)](#).

For NSP Analytics reports, the following considerations apply:

 - AA applications and profile filters must be set to All Traffic.
 - To collect DNS statistics, add applications that have been created to model DNS server types.
 - To collect OTT VoIP statistics, enter the names of the VoIP OTT applications, for example, Skype Calls, you want to report on.
- 3

Specify the NEs from which to collect AA Cflowd statistics; see [5.11 “To specify the NEs for flow statistics collection” \(p. 81\)](#).
- 4

Proceed with statistics collection configuration, for example, configure AA application or protocol filters.

5.4 To configure flow statistics aggregation

5.4.1 Steps

- 1

Log in to a station that has access to the NSP system.

2 —————
Open a console window.

3 —————
Issue the following RESTCONF API call to specify what data/aggregation types are enabled for collection:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH https://address:443/restconf/data/flowconf:config/policies/aggregation

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "aggregation" : [
    {
      "enabled" : "true",
    }
  ]
}
```

The following table shows the data and aggregation types that can be enabled, as well as their default state. All of these data and aggregation types use the FLOW_INFO_10_NOKIAAA protocol.

Table 5-1 Data types and aggregation types

Data type	Aggregation type	Default
COMPREHENSIVE_BUS	AGG_COMP_BUS_APP_GRP_SUB_DEV	true
COMPREHENSIVE_BUS	AGG_COMP_BUS_APP_GRP_SUB_HOST	true
COMPREHENSIVE_BUS	AGG_COMP_BUS_APP_SUB_DEV	true
COMPREHENSIVE_BUS	AGG_COMP_BUS_APP_SUB_HOST	true
COMPREHENSIVE_RAW_BUS	AGG_BUS_RAW	false
COMPREHENSIVE_RAW_RES	AGG_RES_RAW	false

Table 5-1 Data types and aggregation types (continued)

Data type	Aggregation type	Default
COMPREHENSIVE_RES	AGG_COMP_RES_APP_DEV	true
COMPREHENSIVE_RES	AGG_COMP_RES_APP_GRP_DEV	true
COMPREHENSIVE_RES	AGG_COMP_RES_APP_GRP_HOST	true
COMPREHENSIVE_RES	AGG_COMP_RES_APP_GRP_SUB_DEV	true
COMPREHENSIVE_RES	AGG_COMP_RES_APP_HOST	true
COMPREHENSIVE_RES	AGG_COMP_RES_APP_SUB_DEV	true
COMPREHENSIVE_RES	AGG_COMP_RES_GRP_APP_DEV	true
COMPREHENSIVE_RES	AGG_COMP_RES_GRP_APP_GRP_DEV	true
COMPREHENSIVE_SS_BUS	AGG_BUS_RAW	false
COMPREHENSIVE_SS_RES	AGG_RES_RAW	false
FLOW_CONGESTION	AGG_FLOW_RES_ANL_APP	false
FLOW_CONGESTION	AGG_FLOW_RES_ANL_APP_GRP	false
FLOW_CONGESTION	AGG_FLOW_RES_ANL	false
RTP_AUDIO_BUS	AGG_BUS_RAW	true
RTP_AUDIO_RES	AGG_RES_RAW	true
RTP_VIDEO_BUS	AGG_BUS_RAW	true
RTP_VIDEO_RES	AGG_RES_RAW	true
RTP_VOICE_BUS	AGG_BUS_RAW	true
RTP_VOICE_RES	AGG_RES_RAW	true
SESSION_CONGESTION	AGG_SESSION_RES_ANL_APP	false

Table 5-1 Data types and aggregation types (continued)

Data type	Aggregation type	Default
SESSION_CONGESTION	AGG_SESSION_RES_ANL_APP_GRP	false
SESSION_CONGESTION	AGG_SESSION_RES_ANL	false
TCP_PERF_BUS	AGG_TCPPERF_BUS_APP_GRP	true
TCP_PERF_BUS	AGG_TCPPERF_BUS_APP	true
TCP_PERF_BUS	AGG_TCPPERF_BUS_SUBNET_APP_GRP	true
TCP_PERF_BUS	AGG_TCPPERF_BUS_SUBNET_APP	true
TCP_PERF_RES	AGG_TCPPERF_RES_APP_GRP	true
TCP_PERF_RES	AGG_TCPPERF_RES_APP	true
TCP_PERF_RES	AGG_TCPPERF_RES_GRP_APP_GRP	true
TCP_PERF_RES	AGG_TCPPERF_RES_GRP_APP	true
TCP_PERF_SS_BUS	AGG_BUS_RAW	false
TCP_PERF_SS_RES	AGG_RES_RAW	false
UAPPS	AGG_UNKNOWN_APP	true
VOLUME_BUS	AGG_VOL_BUS_APP_GRP	true
VOLUME_BUS	AGG_VOL_BUS_APP	true
VOLUME_BUS	AGG_VOL_BUS_SUBNET_APP_GRP	true
VOLUME_BUS	AGG_VOL_BUS_SUBNET_APP	true
VOLUME_RES	AGG_VOL_RES_APP_GRP	true
VOLUME_RES	AGG_VOL_RES_APP	true
VOLUME_RES	AGG_VOL_RES_GRP_APP_GRP	true
VOLUME_RES	AGG_VOL_RES_GRP_APP	true
VOLUME_SS_BUS	AGG_BUS_RAW	false

Table 5-1 Data types and aggregation types (continued)

Data type	Aggregation type	Default
VOLUME_SS_RES	AGG_RES_RAW	false

4

Issue the following RESTCONF API call to specify the length of the aggregation intervals for each data type:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH `https://address:443/restconf/data/flowconf:config/policies/agg-interval`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "agg-interval" : [
    {
      "duration" : "3",
    }
  ]
}
```

i **Note:** The settings in the above example represent the default values.

i **Note:** The statistics collection interval affects NSP Flow Collector performance. A larger interval results in proportionally larger files, which take longer to store and transfer.

The following table shows the aggregation intervals that can be configured, as well as their default state. All of these entries use the FLOW_INFO_10_NOKIAAA protocol.

Table 5-2 Aggregation intervals

Entry	Default (minutes)
COMPREHENSIVE_BUS	60
COMPREHENSIVE_RAW_BUS	5
COMPREHENSIVE_RAW_RES	5
COMPREHENSIVE_RES	60
COMPREHENSIVE_SS_BUS	15
COMPREHENSIVE_SS_RES	15

Table 5-2 Aggregation intervals (continued)

Entry	Default (minutes)
FLOW_CONGESTION	15
PGW_EDR	5
RTP_AUDIO_BUS	15
RTP_AUDIO_RES	15
RTP_VIDEO_BUS	15
RTP_VIDEO_RES	15
RTP_VOICE_BUS	15
RTP_VOICE_RES	15
SESSION_CONGESTION	15
TCP_PERF_BUS	15
TCP_PERF_RES	15
TCP_PERF_SS_BUS	15
TCP_PERF_SS_RES	15
UAPPS	15
VOLUME_BUS	15
VOLUME_RES	15
VOLUME_SS_BUS	15
VOLUME_SS_RES	15

- 5 _____
Close the console window.

END OF STEPS _____

5.5 To set upload targets and policies

5.5.1 Steps

- 1 _____
Log in to a station that has access to the NSP system.
- 2 _____
Open a console window.

3

Issue the following RESTCONF API call to configure a policy that defines the FTP/SFTP servers that can be used to upload the data collected in files:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**POST https://address:443/restconf/data/flowconf:
config/policies/upload-target**

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "upload-target" : [
    {
      "name" : "name",
      "protocol" : "protocol",
      "ip-address" : "IP address",
      "port" : "port"
      "username" : "username"
      "password" : "password"
      "descr" : "description",
    }
  ]
}
```

where

name is the server name

protocol is one of the following:

- FTP
- SFTP

IP address is the server IP address

port is the server port

username is the FTP or SFTP username

password is the FTP or SFTP password

description is an optional description

4

Issue the following RESTCONF API call to configure a policy that defines the FTP/SFTP servers that can be used to upload the data collected in files:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST https://address:443/restconf/data/flowconf:config/policies/upload

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "upload" : [
    {
      "file-type" : "file type",
      "server" : "server",
      "enabled" : "true",
    }
  ]
}
```

where

file type is NAT, IPDR, or CSV

server is the server IP address

i **Note:** The settings in the above example represent the default values.

5 _____
Close the console window.

END OF STEPS _____

5.6 To configure the transfer of results files

5.6.1 Purpose

You can enable the transfer of result files to a file server in the following formats.

- NAT
- IPDR
- CSV

i **Note:** A minimum 1-Gbyte/s link is required between the NSP Flow Collector and the file server.

i **Note:** SFTP transfers are considerably slower than FTP transfers.

5.6.2 Steps

1 _____
Log in to a station that has access to the NSP system.

2 —————
Open a console window.

3 —————
Issue the following RESTCONF API call to configure the extraction of AA correlation data from NFM-P using the XML API:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH https://address:443/restconf/data/flowconf:config/extraction/nfmp/xml-api

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "xml-api" : [
    {
      "https-port" : "port",
      "request-timeout" : "request timeout",
      "username" : "username"
      "password" : "password",
    }
  ]
}
```

where

port is the port

request timeout is amount of time after which an unanswered request will expire

username is the username

password is the password

i **Note:** This configuration has only one entry, which is pre-configured with default values. This entry can be updated (PATCH), but cannot be replaced or deleted.

4 —————
Issue the following RESTCONF API call to configure the transfer of AA correlation data from NFM-P after an extraction:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH https://address:443/restconf/data/flowconf:config/extraction/nfmp/transfer

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "transfer" : [
    {
      "protocol" : "protocol",
      "port" : "port",
      "username" : "username"
      "password" : "password",
    }
  ]
}
```

where

port is the file server port

username is the FTP or SFTP username

password is the FTP or SFTP password

protocol is one of the following:

- FTP
- SFTP



Note: This configuration has only one entry, which is pre-configured with default values. This entry can be updated (PATCH), but cannot be replaced or deleted.

5

Close the console window.

END OF STEPS

5.7 To configure AA filters

5.7.1 Steps

1

Log in to a station that has access to the NSP system.

2

Open a console window.

3

As required, issue the following RESTCONF API call to configure an AA filter that specifies a subscriber type:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH `https://address:443/restconf/data/flowconf:config/filters/aa/sub-type`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "sub-type" : [
    {
      "sub-type" : "sub",
      "enabled" : "true",
    }
  ]
}
```

where *sub* is the subscriber type

i **Note:** The settings in the above example represent the default values.

i **Note:** This configuration has a fixed number of entries (8), which are pre-configured with default values. These entries can be updated (PATCH), but cannot be replaced or deleted.

4

As required, issue the following RESTCONF API call to specify an AA application defined on the router:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST `https://address:443/restconf/data/flowconf:config/filters/aa/uapps/app`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "app" : [
    {
      "name" : "name",
    }
  ]
}
```

where *name* is the name of the AA application defined on the router

5

As required, issue the following RESTCONF API call to specify an AA protocol defined on the router:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST `https://address:443/restconf/data/flowconf:config/filters/aa/uapps/protocol`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "protocol" : [
    {
      "name" : "name",
    }
  ]
}
```

where *name* is the name of the AA protocol defined on the router

6

Close the console window.

END OF STEPS

5.8 To configure an AA Cflowd special-study policy

5.8.1 Steps

1

Log in to a station that has access to the NSP system.

2

Open a console window.

3

As required, issue the following RESTCONF API call to specify an AA application defined on the router:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST `https://address:443/restconf/data/flowconf:config/filters/aa/special-study/app`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "app" : [
    {
      "name" : "name",
    }
  ]
}
```

where *name* is the name of the AA application defined on the router

4

As required, issue the following RESTCONF API call to specify an AA application group defined on the router:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**POST https://address:443/restconf/data/flowconf:
config/filters/aa/special-study/app-grp**

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "app-grp" : [
    {
      "name" : "name",
    }
  ]
}
```

where *name* is the name of the AA application group defined on the router

5

As required, issue the following RESTCONF API call to specify an AA subscriber defined on the router:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**POST https://address:443/restconf/data/flowconf:
config/filters/aa/special-study/sub**

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "sub" : [
    {
      "sys-ip" : "sys-ip"
      "sub-type" : "sub-type"
      "sub-id" : "sub-id",
    }
  ]
}
```

where

sys-ip is the router IP address

sub-type is the subscriber type

sub-id is the subscriber ID

6

As required, issue the following RESTCONF API call to specify an AA subscriber and application defined on the router:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST <https://address:443/restconf/data/flowconf:config/filters/aa/special-study/sub-app>

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "sub-app" : [
    {
      "sys-ip" : "sys-ip"
      "sub-type" : "sub-type"
      "sub-id" : "sub-id"
      "app" : "app",
    }
  ]
}
```

where

sys-ip is the NE system IP address

sub-type is the subscriber type

sub-id is the subscriber ID

app is the AA application defined on the router

7

Close the console window.

END OF STEPS

5.9 To configure AA subnet summarization filters

5.9.1 Steps

1

Log in to a station that has access to the NSP system.

2

Open a console window.

3

As required, issue the following RESTCONF API call to specify default values for AA subnet summarization filtering:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

PATCH `https://address:443/restconf/data/flowconf:config/filters/aa/subnet-sum/defaults`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "defaults" : [
    {
      "local-ipv4-mask-len" : "local IPv4 mask length"
      "remote-ipv4-mask-len" : "remote IPv4 mask length"
      "local-ipv6-mask-len" : "local IPv6 mask length"
      "remote-ipv6-mask-len" : "remote IPv6 mask length"
      "enabled" : "true",
    }
  ]
}
```

where

local IPv4 mask length is the local IPv4 mask length

remote IPv4 mask length is the remote IPv4 mask length

local IPv6 mask length is the local IPv6 mask length
remote IPv6 mask length is the remote IPv6 mask length

i **Note:** The settings in the above example represent the default values.

i **Note:** If no other matching filters are defined, these values are applied to AA Subnet Summarization data types.

4

As required, issue the following RESTCONF API call to specify AA subnet summarization filters:

i **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**POST https://address:443/restconf/data/flowconf:
config/filters/aa/subnet-sum/entry**

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "entry" : [
    {
      "sys-ip" : "sys-ip"
      "sub-type" : "sub-type"
      "sub-id" : "sub-id",
      "local-ipv4-mask-len" : "local IPv4 mask length"
      "remote-ipv4-mask-len" : "remote IPv4 mask length"
      "local-ipv6-mask-len" : "local IPv6 mask length"
      "remote-ipv6-mask-len" : "remote IPv6 mask length"
      "enabled" : "true",
    }
  ]
}
```

where

sys-ip is the router IP address

sub-type is the subscriber type

sub-id is the subscriber ID

local IPv4 mask length is the local IPv4 mask length

remote IPv4 mask length is the remote IPv4 mask length

local IPv6 mask length is the local IPv6 mask length

remote IPv6 mask length is the remote IPv6 mask length



Note: The settings in the above example represent the default values.

5

Close the console window.

END OF STEPS

5.10 To configure the AA flow data persistence

5.10.1 Purpose

Perform this procedure to specify, for one or more AA statistics domains, whether the NFM-P retains the collected flow data for NSP Analytics reporting, forwards the data to a target file server, or both.

5.10.2 Steps

1

Log in to a station that has access to the NSP system.

2

Open a console window.

3

Issue the following RESTCONF API call to specify export formats for AA data types:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**PATCH https://address:443/restconf/data/flowconf:
config/policies/aa/data-export**

where *address* is the NSP cluster advertised address


The message body is the following:

```
{
  "data-export" : [
    {
      "ipdr" : "true",
      "csv" : "false",
      "vertica" : "false",
    }
  ]
}
```



Note: The settings in the above example represent the default values.

The IPDR option enables data encoding in XDR format.
The CSV option enables data encoding in CSV format.
The Vertica option enables data storage for use by the NSP auxiliary database.

 **Note:** Supported data domains (seen in the response body) include BUSINESS, RESIDENTIAL, and WIFI.

-
- 4 _____
Close the console window.

END OF STEPS _____


5.11 To specify the NEs for flow statistics collection

5.11.1 Steps

-
- 1 _____
Log in to a station that has access to the NSP system.

-
- 2 _____
Open a console window.

-
- 3 _____
Issue the following RESTCONF API call for each NE:

 **Note:** In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

**POST https://address:443/restconf/data/flowconf:
config/policies/collection**

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "collection" : [
    {
      "mgmt-ip" : "mgmt_IP",
      "sys-ip" : "sys_IP",
      "descr" : "description",
      "protocol" : "FLOW_INFO_10_NOKIAAA"
    }
  ]
}
```



Note: The `sys_IP` value must match the System ID that the NFM-P associates with the NE, for example, as shown on the NE properties form in the GUI. System ID must be IPv4.

where

`mgmt_IP` is the management IP address

`sys_IP` is the system IP address

`description` is an optional description

4

Close the console window.

END OF STEPS

5.12 To extract an NFM-P network data snapshot

5.12.1 Purpose

An NSP Flow Collector Controller requires an image, called a snapshot, of current NFM-P data that is subsequently distributed to each NSP Flow Collector that it controls.

The following steps describe how to force an NSP Flow Collector Controller to extract the AA Cflowd provisioned-object snapshot from the NFM-P.



CAUTION

Service Disruption

Performing the procedure consumes NFM-P main server resources, and is typically required only when recommended by technical support.

Perform the procedure only if required, and only under the guidance of technical support during a period of low NFM-P system activity.

5.12.2 Steps

1

Log in to a station that has access to the NSP system.

2

Open a console window.

3

Issue the following RESTCONF API call to configure the extraction of NFM-P correlation data for the AA use case:



Note: In order to issue a RESTCONF API call, you require a token; see the My First NSP API Client tutorial on the [Network Developer Portal](#) for information.

POST `https://address:443/restconf/operations/flowconf:extract-po`

where *address* is the NSP cluster advertised address

The message body is the following:

```
{
  "flowconf:input" : [
    {
      "po-category" : "AA",
    }
  ]
}
```

4

Close the console window.

END OF STEPS

5.13 To export IPDR statistics files to a different format

5.13.1 Purpose



CAUTION

Data Loss

The tool used in the procedure deletes all .xdr files in the output directory before starting the export operation.

If the output directory contains files that you want to retain, copy the files to another location before you use the tool.

Use this procedure to convert XDR-encoded IPDR statistics files to a different format for compatibility with an OSS application.



Note: If you use the GUI version of the tool over a remote connection, you must first enable X.11 or native X display redirection on the main server as the root user.



Note: You require root user privileges and full access permissions on the input and output directories that you specify.



Note: Command lines use the # symbol to represent the RHEL CLI prompt for the root user. Do not type the leading # symbol when you enter a command.

5.13.2 Steps

- 1

Log in to the NFM-P main server station as the root user.
- 2

Open a console window.
- 3

To export files using the CLI-based tool, perform the following steps.
 1. Enter the following:

```
# /opt/nsp/nfmp/server/nms/bin/ipdrDecode.bash input_dir output_dir format ↵
```

where
input_dir is the absolute path of a directory that contains the IPDR files to export, or the absolute path of a file to export
output_dir is the directory that is to contain the exported files
format is the output file format, either txt or xml; if omitted, the default is xml
The following message and prompt are displayed:

```
INFO: This utility decodes all *.xdr files in the specified directory.  
Do you wish to proceed <y/n> ?
```
 2. Enter y ↵.
The tool begins to convert and export files. As each file is processed, messages like the following are displayed:

```
Decoding file...  
...finished successfully!
```

The export of all files is complete when the following is displayed:

```
...finished decoding!
```
- 4

To export files using the GUI-based tool, perform the following steps.
 1. Enter the following:

```
# /opt/nsp/nfmp/server/nms/bin/ipdrDecodeGui.bash ↵
```

The following message and prompt are displayed:

```
INFO: This utility decodes all *.xdr files in the specified directory.  
Do you wish to proceed <y/n> ?
```
 2. Enter y ↵.
The IPDR/XDR File Decoder window opens.

-
3. Use the Output Format drop-down to specify the required output file format.
 4. To convert and export one file, click Decode a File... and use the file browser form that opens to choose an input file.
 5. To convert and export all files in a directory, click Decode XDR Files in Dir... and use the file browser form that opens to choose an input directory.
 6. The tool begins to convert and export files. As each file is processed, messages like the following are displayed:

`Decoding file...`

`...finished successfully!`

The export of all files is complete when the following is displayed:

`...finished decoding!`

7. Close the IPDR/XDR File Decoder window.

5

Close the console window.

END OF STEPS

6 Server performance statistics

6.1 Server performance statistics collection

6.1.1 Overview

Server performance statistics are collected from the NFM-P main servers in an NFM-P system. To collect server performance statistics, you need to create and apply the following policies:

- statistics policy
- server performance collection policy

See [Chapter 2, “Statistics collection”](#) for more information about the policies.

6.2 Pathway: server performance statistics collection

6.2.1 Process

1

Configure the statistics policy associated with the statistics data. See [6.3 “To configure a statistics policy for server performance statistics” \(p. 87\)](#) for more information.

2

Configure the server performance collection policy that is associated with the statistics data. See [6.4 “To configure a statistics collection policy for server performance statistics” \(p. 88\)](#) for more information.

3

Delete statistics records, as required. See [6.5 “To delete server performance statistics records” \(p. 89\)](#) for more information.

6.3 To configure a statistics policy for server performance statistics

6.3.1 Steps

1

Choose Tools→Statistics→Server Performance Statistics from the NFM-P main menu. The Server Performance Statistics form opens.

2

Choose one of the statistics classes from the object drop-down menu.

-
- 3

Click Statistics Policy. The Statistics Policy form opens.
 - 4

Configure the required parameters.
 - 5

Click on the Thresholds tab. The Thresholds tab contains a threshold parameter for each counter in the statistics class. When a statistics counter threshold is exceeded, the NFM-P raises a threshold-crossing alarm.
 - 6

Configure one or more threshold parameters by selecting the check box beside a parameter and then entering a threshold value.
 - 7

Save your changes and close the forms.

END OF STEPS

6.4 To configure a statistics collection policy for server performance statistics

6.4.1 Steps

- 1

Choose Tools→Statistics→Server Performance Statistics from the NFM-P main menu. The Server Performance Statistics form opens.
- 2

Choose a statistics class from the object drop-down menu.
- 3

Click Collection Policy. The Collection Policy form opens.
- 4

Configure the required parameters.
- 5

Save your changes and close the forms.

END OF STEPS

6.5 To delete server performance statistics records

6.5.1 Steps

- 1 _____
Choose Tools→Statistics→Server Performance Statistics from the NFM-P main menu. The Server Performance Statistics form opens.
- 2 _____
Choose a statistics class from the object drop-down menu.
- 3 _____
Click Statistics Policy. The Statistics Policy form opens.
- 4 _____
Click Purge Statistics Records. The Statistics Policy filter form opens.
- 5 _____
Configure the filter and click OK. A confirmation message is displayed.
- 6 _____
Click Yes. The statistics records are deleted.
- 7 _____
Close the open forms.

END OF STEPS _____

Part III: Statistics presentation

Overview

Purpose

This volume describes the presentation of statistics in the NFM-P.

Contents

Chapter 7, Statistics presentation	93
Chapter 8, Viewing statistics	97
Chapter 9, Graphing statistics	111

7 Statistics presentation

7.1 Statistics presentation in the NFM-P

7.1.1 Overview

You can view all NFM-P-supported statistics types in a table or graph. A table lists specific values that you can filter and sort. A graph can display multiple statistics counters simultaneously, and helps to identify trends. A graph can also use dual, independently scaled axes that allow an NFM-P operator to visually compare statistics with high numbers, such as throughput statistics, to statistics with very low numbers, such as dropped packet statistics.

You can display statistics for multiple objects in one list, and can export table data or a graph image to a file in multiple formats.


There are two ways to choose the statistics to view in a table or graph.

- bottom-up method—uses the Statistics tab of an object properties form; applies to performance and accounting statistics
- top-down method—uses the NFM-P Statistics Manager; applies to all statistics types

7.1.2 Tabular statistics view

Statistics records for an object are listed on the Statistics tab of the object properties form. Each statistics class in the object drop-down menu corresponds to an NE MIB table or accounting file.

From the Statistics tab, you can use the Collect button to collect on-demand network or server performance statistics, and the Collect All button to collect one on-demand statistics record for each statistic type that the object supports.

 **Note:** On-demand collection is not available for accounting statistics, because accounting statistics are available only after an NE transfers an accounting statistics file to the NFM-P.

You can sort and filter the listed records, for example, to display only the statistics collected during a specific period, and can export the list to a file.

A Properties button on the Statistics tab enables you to view a selected record. Depending on the number and types of counters that a record contains, the counters are displayed on separate tabs of the record properties form. The General tab displays information about the record, for example, the record type and collection time, and the other tabs each display a group of counters, for example, octet-based or packet-based counters, to facilitate finding a specific counter.

7.1.3 Graphical statistics view



NOTICE

Service-disruption hazard

When multiple NFM-P clients each have multiple open plotters that display multiple counters, a large volume of statistics is collected from the NEs. If the data for multiple plotters is from one NE, the independent polls of the NE by each plotter may affect performance.

You can use a scope of command role to limit plotter access to specific NFM-P user groups. See the NSP System Administrator Guide for scope of command role information.

The NFM-P Statistics Plotter creates graphs using statistics data. The plotter can display multiple performance, accounting, and server performance statistics simultaneously, using dual Y axes to represent different value scales. The plotter also displays the numerical value of each point on the graph in a table, and can display minimum, maximum, and average values for a plot area. Multiple Statistics Plotter windows can be simultaneously active in one client GUI.

You can use the NFM-P main menu, contextual menu options, or the Plotter button on a form to open the Statistics Plotter and create graphs of real-time or historical statistics. Historical plots use data from previous scheduled collections. Real-time plots collect statistics data while the plotter window is open and plot the data as it is collected.



Note: Historical plots use only statistics data from scheduled collections, not data from on-demand collections.

User-defined plotter profiles specify the plot attributes for a type of object. Plotter profiles are a convenient way to open the same type of plot for different objects of the same type. Right-click menu options for objects in navigation trees and maps open a plotter window for the object and optionally display the plots defined in one or more plotter profiles.

Graphical presentation options

You can plot raw values, and periodic counter values, which the NFM-P calculates by subtracting the previous counter value from the current counter value. In graphical form, periodic data is typically more informative than raw counter data for troubleshooting and trend analysis. You can also plot values such as the following, which provide a more comprehensive view of network activity:

- link utilization

The NFM-P uses ingress and egress port and channel statistics to calculate utilization values for plotting. Plotted utilization values provide a convenient view of the bandwidth usage.

- throughput data expressed as a flow rate

When you need to monitor a data flow rate, you can choose to plot normalized counter data, such as the number of octets per second received on a port, using a scale that is appropriate for the flow rate.

- event-based data expressed as a rate

You can plot normalized event-based counter data, such as the number of received DHCP requests or dropped packets per second.

- matching counter data for the opposite endpoint of a link

When you create a plot of endpoint data from an object such as a physical or optical link that terminates on a physical port, you can use a button in the plotter window to create a plot of the same data for the opposite endpoint, and create a profile that you can use to automatically plot both endpoints of other links. This function can help to isolate throughput issues.

See [Chapter 9, “Graphing statistics”](#) for information about using the NFM-P Statistics Plotter.

8 Viewing statistics

8.1 Viewing statistics in the NFM-P

8.1.1 Overview

Statistics can be viewed as numerical data in a table, which is useful for obtaining raw counter values. The tabular data can be sorted, filtered, and exported to files in different formats.

AA statistics can be viewed from the Statistics tab of a local AA policy object, such as an application, application group, or protocol.

8.2 Pathway: viewing statistics

8.2.1 Process

1

Monitor on-demand performance statistics. See [8.3 “To view on-demand statistics” \(p. 98\)](#) for more information.

2

Monitor statistics for sets of objects. See [8.4 “To view statistics for a set of objects” \(p. 99\)](#) for more information.

3

Monitor scheduled collection for the following statistics types, as required:

- performance—see [8.6 “To view performance statistics” \(p. 101\)](#) for more information
- service accounting—see [8.8 “To view service accounting statistics” \(p. 103\)](#) for more information
- specific-object—see [8.5 “To view network accounting statistics” \(p. 100\)](#) for more information
- subscriber accounting—see [8.9 “To view subscriber accounting statistics” \(p. 104\)](#) for more information
- AA accounting—see [8.9 “To view subscriber accounting statistics” \(p. 104\)](#) to [8.13 “To view AA accounting statistics for a protocol” \(p. 109\)](#) for more information


4

Monitor NFM-P server performance statistics using the GUI. See [8.7 “To view server performance statistics” \(p. 102\)](#) for more information.

8.3 To view on-demand statistics

8.3.1 Steps

- 1 _____
Open the properties form of the object for which you want to view statistics. The General tab is displayed.
- 2 _____
Click on the Statistics tab.
- 3 _____
Choose a statistics class from the object drop-down menu.

 **Note:** The Collect and Collect All buttons are not displayed when you choose a statistics class that does not support on-demand statistics collection.
- 4 _____
Specify a filter to create a filtered list of statistics records, if required.
- 5 _____
Perform one of the following:
 - a. Click Collect to collect the statistics for only the specified class.
 - b. Click Collect All to collect statistics for all classes.The statistics records are displayed in a list.
- 6 _____
Identify a statistics record to view.
- 7 _____
Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.
- 8 _____
Close the object properties form.

END OF STEPS

8.4 To view statistics for a set of objects

8.4.1 Steps

- 1 _____
Choose Tools→Statistics→Statistics Manager from the NFM-P main menu. The Statistics Manager form opens.
- 2 _____
Configure the Statistics Type parameter.
- 3 _____
Do one of the following:
 - a. Choose a statistics class from the object drop-down tree.
 - b. Choose a statistics class by clicking Filter for Object Type.
- 4 _____
Click Search. A list of entries is displayed.
- 5 _____
Identify an entry to view.
- 6 _____
Do one of the following:
 - a. Scroll horizontally to view the values that the entry contains.
 - b. Open the entry to view it.
 1. Select the entry and click Properties. The appropriate form opens, based on the Statistics Type parameter setting.
 2. View the form contents.
 3. Close the form.
- 7 _____
Close the Statistics Manager form.

END OF STEPS _____

8.5 To view network accounting statistics

8.5.1 Steps

- 1 _____
Choose Manage→Service Tunnels from the NFM-P main menu. The Manage Service Tunnels form opens.
- 2 _____
Specify a filter to create a filtered list of service tunnels and click Search. A list of service tunnels is displayed.
- 3 _____
Select a service tunnel in the list and click Properties. The service tunnel properties form opens.
- 4 _____
Click on the Statistics tab.
- 5 _____
Choose a statistics class from the object drop-down menu.
- 6 _____
Specify a filter to create a filtered list of statistics classes, if required.
- 7 _____
Click Search. A list of statistics records is displayed.
- 8 _____
Identify a statistics record to view.
- 9 _____
Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.
- 10 _____
Close the service tunnel properties form.


-
- 11 _____
Close the Manage Service Tunnels form.

END OF STEPS _____

8.6 To view performance statistics

8.6.1 Steps

- 1 _____
Right-click on an object in the navigation tree and choose Properties. The properties form for the object opens.
- 2 _____
Click on the Statistics tab.
- 3 _____
Choose a statistics class from the object drop-down menu.

 **Note:** The Collect and Collect All buttons are not displayed when you choose a statistics class that does not support on-demand statistics collection.
- 4 _____
Specify a filter to create a filtered list of statistics classes.
- 5 _____
Perform one of the following:
 - a. Click Collect to collect the statistics for only the specified class.
 - b. Click Collect All to collect statistics for all classes.The statistics records are displayed in a list.
- 6 _____
Identify a statistics record to view.
- 7 _____
Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.

-
- 8 _____
- Close the object properties form.

END OF STEPS _____

8.7 To view server performance statistics

8.7.1 Steps

- 1 _____
- Choose Tools→Statistics→Server Performance Statistics from the NFM-P main menu. The Server Performance Statistics form opens.
- 2 _____
- Choose a statistics class from the object drop-down menu.
- 3 _____
- Specify a filter to create a filtered list of statistics records, if required.
- 4 _____
- Perform one of the following:
- a. Click Search to list statistics records for scheduled collections of the statistics class.
 - b. Click Collect to perform an on-demand collection for the statistics class.
- 5 _____
- Identify a statistics record to view.
- 6 _____
- Perform one of the following:
- a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 - 1. Select the statistics record and click Properties. The Statistics Record form opens.
 - 2. View the statistics record.
 - 3. Click Close to close the Statistics Record form.
- 7 _____
- Close the Server Performance Statistics form.

END OF STEPS _____

8.8 To view service accounting statistics

8.8.1 Steps

- 1 _____
Choose Manage→Service→Services from the NFM-P main menu. The Manage Services form opens.
- 2 _____
Specify a filter to create a filtered list of services and click Search. A list of services is displayed.
- 3 _____
Select a service in the list and click Properties. The service properties form opens.
- 4 _____
Click on the L2 Access Interfaces or L3 Access Interfaces tab, as required, to display a list of access interfaces.
- 5 _____
Select an interface in the list and click Properties. The appropriate access interface properties form opens.
- 6 _____
Click on the Statistics tab.
- 7 _____
Choose a statistics class from the object drop-down menu.
- 8 _____
Specify a filter to create a filtered list of statistics classes, if required.
- 9 _____
Click Search. A list of statistics records is displayed.
- 10 _____
Identify a statistics record to view.
- 11 _____
Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.

-
2. View the statistics record.
 3. Click Close to close the Statistics Record form.

12

Save your changes and close the forms.

END OF STEPS

8.9 To view subscriber accounting statistics

8.9.1 Steps

- 1**

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.
- 2**

Choose Residential Subscriber Instance (Residential Subscriber) and click Search. A list of subscriber instances is displayed.
- 3**

Select a subscriber instance in the list and click Properties. The Residential Subscriber Instance properties form opens.
- 4**

Click on the Statistics tab.
- 5**

Choose a statistics class from the object drop-down menu.
- 6**

Specify a filter to create a filtered list of statistics classes, if required.
- 7**

Click Search. A list of statistics records is displayed.
- 8**

Identify a statistics record to view.
- 9**

Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.

-
- b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.

10

Save your changes and close the forms.

END OF STEPS

8.10 To view per-subscriber AA accounting statistics

8.10.1 Steps

1

Choose Manage→Residential Subscribers from the NFM-P main menu. The Manage Residential Subscribers form opens.

2

Choose Residential Subscriber Instance (Residential Subscriber) and click Search. A list of subscriber instances is displayed.

3

Select a subscriber instance in the list and click Properties. The Residential Subscriber Instance properties form opens.

4

Click on the Statistics tab.

5

Choose an AA statistics class from the object drop-down menu.

6

Specify a filter to create a filtered list of statistics classes, if required.

7

Click Search. A list of statistics records is displayed.

8

Identify a statistics record to view.

9

Perform one of the following:

- a. Scroll horizontally to view the statistics counter values for the statistics record.
- b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.

10

Save your changes and close the forms.

END OF STEPS

8.11 To view AA accounting statistics for an application

8.11.1 Steps

1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.

2

Choose AA Group Policy and click Search. A list of AA group policies is displayed.

3

Select a policy in the list and click Properties. The AA Group Policy (Global) form opens.

4

Click on the Local Definitions tab.

5

Select a local policy definition in the list and click Properties. The local AA Group Policy form opens.

6

Click on the Applications tab.

7

Choose an application in the list and click Properties. The Application properties form opens.

-
- 8

Click on the Statistics tab.
 - 9

Specify a filter to create a filtered list of statistics classes, if required.
 - 10

Click Search. A list of statistics records is displayed.
 - 11

Identify a statistics record to view.
 - 12

Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record form.
 - 13

Close the forms.

END OF STEPS

8.12 To view AA accounting statistics for an application group

8.12.1 Steps

- 1

Choose Policies→ISA Policies→Application Assurance from the NFM-P main menu. The Application Assurance Policies form opens.
- 2

Choose AA Group Policy and click Search. A list of AA group policies is displayed.
- 3

Select a policy in the list and click Properties. The AA Group Policy (Global) form opens.
- 4

Click on the Local Definitions tab.

5 _____
Select a local policy definition in the list and click Properties. The local AA Group Policy form opens.

6 _____
Click on the Application Groups tab.

7 _____
Choose an application group in the list and click Properties. The Application Group properties form opens.

8 _____
Click on the Statistics tab.

9 _____
Specify a filter to create a filtered list of statistics classes, if required.

10 _____
Click Search. A list of statistics records is displayed.

11 _____
Identify a statistics record to view.

12 _____
Perform one of the following:

- a. Scroll horizontally to view the statistics counter values for the statistics record.
- b. Open the statistics record to view it.
 - 1. Select the statistics record and click Properties. The Statistics Record form opens.
 - 2. View the statistics record.
 - 3. Click Close to close the Statistics Record form.

13 _____
Close the forms.

END OF STEPS _____

8.13 To view AA accounting statistics for a protocol

8.13.1 Steps

- 1

In the navigation tree equipment view, expand *device*→ISA-AA Groups→*ISA_group*.
- 2

Right-click on the ISA-AA Group icon and choose Properties. The ISA-AA Group (Edit) form opens.
- 3

If the ISA-AA group is partitioned, click on the ISA-AA Partitions tab. A list of ISA-AA partitions is displayed.
- 4

Select an ISA-AA partition and click Properties. The ISA-AA Group Partition (Edit) form opens.
- 5

Click on the Statistics tab.
- 6

Perform one of the following:
 - a. Scroll horizontally to view the statistics counter values for the statistics record.
 - b. Open the statistics record to view it.
 1. Select the statistics record and click Properties. The Statistics Record - AA Protocol Stats form opens.
 2. View the statistics record.
 3. Click Close to close the Statistics Record - AA Protocol Stats form.
- 7

Close the ISA-AA Group form.

END OF STEPS

9 Graphing statistics


9.1 Graphing statistics overview

9.1.1 Overview

The NFM-P Statistics Plotter displays statistics data in graphical form. [Figure 9-1, “NFM-P Statistics Plotter” \(p. 111\)](#) shows the Statistics Plotter window, which has the following panels:

- configuration—allows you to specify the counters to plot
- detail—contains the plotted graph points, axes, and grid lines
- overview—provides an overview of the entire data range
- data list—displays the graph data in tabular form

Figure 9-1 NFM-P Statistics Plotter



23546

The Statistics Manager form, and the Statistics tab of an object properties form, each have a statistics class drop-down menu. An icon beside each statistics class in the menu indicates the plot types that the class supports; the icons are shown in [Figure 9-2, “Plotting support indicators”](#) (p. 111) .

Figure 9-2 Plotting support indicators



23771

9.1.2 About the NFM-P Statistics Plotter

The NFM-P Statistics Plotter allows you to plot counter values for real-time or historical statistics. An NFM-P GUI client can open up to five Statistics Plotter windows, and each window can simultaneously plot up to eight counters. You can maximize or minimize a panel in the Statistics Plotter by clicking on the up or down arrow on the left side of the divider bar between it and the neighboring panel.

You can open a new plot for a specific object using the following:

- contextual menu option in network navigation tree or map
- NFM-P clipboard paste function
- Plotter button on the Statistics tab of the object properties form

Note: Objects that are not stored in the NFM-P database, such as LSN subscriber hosts, have no Statistics tab on the object properties form; the Plotter button for such an object is on the home tab of the properties form.

See [9.3 “To create a statistics graph”](#) (p. 116) and [9.4 “To modify a statistics graph”](#) (p. 121) for information about creating and modifying plots.

Configuration panel

You use the configuration panel to specify the objects and statistics counters in a graph. You can enter up to eight counters in the panel. If more than one counter from the same monitored object is required, you can click Duplicate and choose the new counter and class. After you choose the counters, you can close the panel to increase the space available for the graph panels. The configuration panel contains the following columns:

- Marker—shows the symbol that marks each data point when markers are enabled
- Plot Name—identifies the plot
- Y Axis—allows you to choose the Y axis for a plot; see [“Choosing the Y axis”](#) (p. 113) in this section for information

- **Monitored Object**—identifies the object of a plot; if an object has child objects that support plotting, the Monitored Object in a new plot is a drop-down menu of the child objects, as shown for plot 4 in [Figure 9-1, “NFM-P Statistics Plotter”](#) (p. 111)
- **Statistics Group**—allows you to choose the statistics class; an icon beside each class indicates the supported plot types, as described in [Figure 9-2, “Plotting support indicators”](#) (p. 112)
- **Statistics Counter**—allows you to choose the statistics counter
- **Data Format**—allows you to choose how the statistics data is quantified; see [“Statistics data formats”](#) (p. 114) in this section for information

Detail panel

The detail panel contains the plotted statistics line graphs. The panel can contain historical or real-time graphs, but not both.

Real-time statistics are plotted as they are collected. Historical plotting uses the data stored in the NFM-P database. For real-time statistics, you can enable or disable automatic scrolling. See [“Overview panel”](#) (p. 113) in this section for more information.

You can zoom in and out on a graph, and can enable or disable the display of markers and a legend that includes the monitored object name. Markers are useful for identifying different lines, for example, when a graph is printed in black and white. You can save a graph as an image, with or without the legend information or markers.

You can enable or disable the X, Y1, and Y2 grid lines. The Y1 and Y2 grid lines are displayed only when the corresponding Y axis is enabled. See [“Choosing the Y axis”](#) (p. 113) for more information about the Y1 and Y2 axes.

Overview panel

The overview panel displays a high-level view of the plot data. The shaded box in the panel, which is called the slider, contains the plot area that is displayed in the detail panel. When you double-click on the slider, you can set automatic scrolling to on or off. The slider is green when automatic scrolling is enabled, and red when it is disabled. You can click on the slider and drag it horizontally to change the display in the detail panel. When the slider is not selected and real-time statistics are being plotted, the slider automatically scrolls to show the latest statistics. When automatic scrolling is disabled, the detail panel statically displays the plot area shown in the slider.

Data list panel

The data list panel displays the plotted statistics values in a tabular format. You can save the table data to a file in CSV or HTML format. See the *NSP NFM-P Classic Management User Guide* for information about using the contextual menu in the plot value list heading.

Choosing the Y axis

A statistics graph has one X axis, but can have one or two Y axes. The X axis is a time scale, and a Y axis is a scale for data values. The second Y axis facilitates the viewing of plots whose data scales differ widely, for example, when one set of plotted values is in the range of 900 000 to 1 000 000, and another set of values is in the range of 100 to 500. [Figure 9-1, “NFM-P Statistics Plotter”](#) (p. 111) shows the two Y axes; Y axis 1 is on the left side of the detail panel, and Y axis 2 is on the right side.

You specify the Y axis for a plot using the Y Axis drop-down menu in the configuration panel. The graph legend indicates which Y axis is associated with each plot, as shown in [Figure 9-1, “NFM-P Statistics Plotter” \(p. 111\)](#).

Statistics data formats

Using the Data Format column in the configuration panel, you can choose how the statistics data is represented in a plot. Depending on the type of object and statistics counter, the following options are available:

- Raw—displays the current counter value, which is cumulative
- Periodic—displays the periodic counter value, which is the difference between the current and previous raw counter values
- Rate (*units*)—displays a normalized throughput value; you can choose from rates such as bps, kbps, Mbps, or Gbps

9.1.3 Plotting real-time statistics data

The Statistics Plotter has a configurable polling interval and retention period for real-time statistics collection. The retention period specifies how long the NFM-P collects real-time statistics before it deletes the oldest sample. For example, at a sampling rate of 10 seconds, the plotter stores 8640 samples per counter in one day. The retention period and a default polling interval are configured on the User Preferences form; see the *NSP NFM-P Classic Management User Guide*.

The detail and data list panels indicate when a real-time statistic is missing from a plot. The detail panel plot markers in the plot list panel Markers column change from solid to dark grey with the marker color as the border. The data list panel displays N/A or Data Missing instead of a data value.

- N/A indicates that events are skewed. For example, there is data event plot for plot 1 and plot 2 at time T1, and a data event plot for plot 3 and plot 4 at time T2. Accordingly, the table row displays N/A for plot 3 and plot 4 at time T1, and for plot 1 and plot 2 at time T2. N/A is displayed whenever the same graph plots counters from different classes or different NEs.
- Data Missing indicates an error during statistics collection. Data Missing is displayed when an NE reports an error or when the polling interval elapses before the response to a previous poll arrives.

You can remove a plot at any time during real-time statistics collection; the plot disappears immediately from the detail and data list panels.

9.1.4 Plotting utilization statistics

The NFM-P can plot ingress and egress utilization statistics using calculated values. Plotted utilization statistics provide a quickly accessible view of the bandwidth usage, on a specified port or channel, in graphical and tabular form.

Ingress utilization statistics represent the received total octets during a collection interval, expressed as a percentage of the port bandwidth. Egress utilization statistics represent the transmitted total octets during a collection interval, expressed as a percentage of the port bandwidth.

You can plot the utilization statistics for all physical ports. Utilization statistics are available for members of the following objects, but not for the objects themselves:

- LAGs
- bundles
- TDM channel groups
- CCAGs
- SONET TUG3s
- SONET tributaries

See [9.5 “To plot utilization statistics” \(p. 122\)](#) for information about how to plot utilization statistics.

9.1.5 Plotter profiles

The NFM-P provides default statistics plotter profiles that you can use to generate customized statistics plots. This enables you to save a specific plotter configuration, for example, a set of queues to monitor, and later apply it to the same type of object. After you save a customized profile, you can use a right-click menu option on an object of the same type to open a plotter window. The window is automatically populated with the statistics counter information in the profile.

The right-click contextual menu for plotter profiles is available for objects in the network navigation tree and map, depending on the object type. The menu is available to all users, but a user with the Administrator or Statistics Plotter Profile Management role can specify whether a specific profile is included in the menu for other users.



Note: Only an Administrator or Statistics Plotter Profile Management user can create and delete customized profiles.

You cannot delete a default profile.

To list and view profiles, use the Statistics Plotter Profiles form, which lists all profiles, or a contextual menu option that lists the profiles that are specific to a type of object. Default profiles are available for objects that include the following:

- dynamic LSPs
- L3 SAPs
- MPLS interfaces
- physical and virtual ports

You can create a customized profile directly from the plotter window, or use a default profile as a template to create a customized profile. See [9.6 “To create a customized plotter profile” \(p. 123\)](#) for information about creating customized plotter profiles. See [9.7 “To plot statistics using a plotter profile” \(p. 124\)](#) for information about how to create a plot using a plotter profile.

Link endpoint profiles

When you create a plot for a port that is a link endpoint, you can enable the automatic plotting of the same counters for the port at the other end of the link. If the function is enabled in a plotter profile, the profile is accessible from links of the same type in topology maps. Using a right-click menu option, you can quickly plot and compare counter data for both endpoints of a link to monitor or troubleshoot the link.

9.2 Pathway: graphing statistics

9.2.1 Process

- 1 _____
Configure statistics collection for the equipment or logical object that is the source of the statistics data to be plotted.
- 2 _____
Create a real-time or historical statistics graph; see [9.3 “To create a statistics graph” \(p. 115\)](#) .
- 3 _____
Modify a real-time or historical statistics graph; see [9.4 “To modify a statistics graph” \(p. 121\)](#) .
- 4 _____
Create a customized plotter profile, if required; see [9.6 “To create a customized plotter profile” \(p. 123\)](#) .
- 5 _____
Use a default or customized plotter profile to plot statistics; see [9.7 “To plot statistics using a plotter profile” \(p. 124\)](#) .

9.3 To create a statistics graph

9.3.1 Steps

- 1 _____
Perform one of the following:
 - a. Use a contextual menu option. Right-click on a map or network navigation tree object and choose Plotter→New Plot or Plotter→*profile_name*. The Statistics Plotter form opens with a plot entry for the object listed in the configuration panel.
 - b. Use the Copy to Clipboard function:
 1. Right-click on a GUI object and choose Copy to Clipboard from the contextual menu, or select an object on a list form and click Copy to Clipboard.
 2. Choose Tools→Statistics→Statistics Plotter→New Plot from the NFM-P main menu. The Statistics Plotter form opens.
 3. Click Paste from Clipboard or right-click in the configuration panel and choose Paste from Clipboard. A plot entry for the object is listed in the configuration panel.
 - c. Click Plotter:
 1. Open the properties form for an object.
 2. Click on the Statistics tab.

Note:

Objects that are not stored in the NFM-P database, such as LSN or ESM subscriber hosts, have no Statistics tab on the object properties form; the Plotter button for such an object is on the home tab of the properties form.

3. Choose a statistics class from the object drop-down menu, if required.

Note:

The icon beside each class indicates the plot types that the class supports. See [Figure 9-2, “Plotting support indicators” \(p. 112\)](#) for the icon descriptions.

4. Click Search to view statistics record entries for scheduled collections of the selected statistics class.
5. Select a statistics record.
6. Click Plotter and choose New Plot. The Statistics Plotter form opens with a plot entry for the object listed in the configuration panel.

2

If the Monitored Object field for the plot entry displays a drop-down arrow, as shown for plot 4 in [Figure 9-1, “NFM-P Statistics Plotter” \(p. 111\)](#), the object has child objects. Use the Monitored Object drop-down menu to specify a child object, if required.

3

Use the Statistics Group drop-down menu to specify a statistics group.

Additional configurable columns are displayed in the configuration panel if the chosen statistics group has filterable attributes. The columns allow you to filter the statistics. Configure a filter, as required:

1. Click on an up or down arrow in a column heading to specify an enumerated entry such as a queue ID.
2. Double-click in a text-based column heading and type in a value.

Note:

The NFM-P does not validate a typed entry. You must type an entire entry correctly to generate the expected filter output.

4

Use the Statistics Counter drop-down menu to specify a statistics counter.

5

Use the Data Format drop-down menu to specify how the counter data is quantified.

6

Use the Y Axis drop-down menu to specify the Y axis that the plot is to use.

7

Add a plot to the graph, if required. Perform one of the following:

- a. Use an existing plot as a template:
 1. Select a plot entry in the list and click Duplicate. A duplicate plot entry is created.
 2. Customize the plot parameters, as described in [Step 2](#) to [Step 6](#).
- b. Add a plot using a new object. Repeat [Step 1](#) to [Step 6](#).

8

Repeat [Step 7](#) to add another plot, as required.



Note: A statistics graph can contain up to eight plots. Each plotted Statistics Counter and Managed Object combination must be unique.

9

Perform one of the following:



Note: Real-time statistics plotting is supported for network performance, server performance, and MIB-based accounting statistics. Historical statistics plotting is supported for network performance, server performance, and accounting statistics. Each plot must be unique. The NFM-P deletes the older plot if you attempt to create a duplicate a plot in this step.

- a. To create a real-time statistics graph:
 1. Choose a polling interval from the Real-time Polling Interval drop-down menu, or enter a value between 10 and 3600.

The default polling interval is configured on the User Preferences form; see the *NSP NFM-P Classic Management User Guide*.
 2. Click Real-time Plot. The detail panel displays the plotted statistics using the configured polling interval.
- Note:**
- You do not have to stop real-time statistics collection to add or remove plots.
- b. To create a graph using historical statistics:
 1. Configure the historical plot start time using the left-hand time field in the configuration panel.
 2. Configure the historical plot end time using the right-hand time field in the configuration panel.
 3. Click Historical Plot. The statistics are retrieved from the database and plotted in the detail panel.

10

If you are plotting statistics data for an object that has physical endpoints, such as a physical link or an optical link, you can plot statistics data for the other endpoint:

1. Select the plot in the configuration panel that you want to duplicate for the opposite endpoint.

2. Click Plot Other Endpoint(s). A plot entry for the other link endpoint is added to the configuration panel, and the plot is added to the graph.

11

Perform one or more of the following to modify the detail panel display.

- a. To display a plot tool tip, move the mouse pointer over the data points in the detail panel. A tool tip identifies the plot number, the statistics collection interval, and the data value at that point.
- b. To change the view displayed in the detail panel, click on the slider in the overview panel and drag it horizontally.



Note: When the slider is not selected and real-time statistics are being collected, the slider autoscrolls to display the latest statistics.

- c. To turn off autoscrolling, double-click on the slider. The slider color changes to red.
- d. To resize the objects in the detail panel, click Zoom in Tool or Zoom out Tool; click Reset Zoom tool to return to the default zoom level.
- e. To display the data points for each plot using markers, select the Markers check box.
- f. To display the legend elements:
 1. Select the Legend check box.
 2. To include the minimum, maximum, and average plot values in the legend, select the MinMaxAvg check box.

Note:

The source of the values in the legend is the plot range displayed in the detail panel. If you move the slider to display a different plot range, the values are recalculated.

- g. To display the X-axis grid lines, select the Grid X check box.
- h. To display the Y1 grid lines, select the Grid Y1 check box. The grid lines are displayed only when the Y1 axis is used by a plot.
- i. To display the Y2 grid lines, select the Grid Y2 check box. The grid lines are displayed only when the Y2 axis is used by a plot.
- j. You may need to hide a plot when plots in the detail panel overlap. To hide a plot:



Note: Statistics data collection for a plot does not stop when the plot is hidden; only the display of the plotted data is suppressed.

1. Right-click on the data list panel header row and choose Column Display. The Column Display form opens.
2. Select the plot to hide in the Displayed on Table list and click the left arrow. The selected plot moves to the Available for Table list.
3. Click OK. The plot information is removed from the detail and data list panels.
4. To show the plot in the detail panel, right-click in the column heading of the data list panel

and choose Plot n from the contextual menu, where n is the plot to show. A check mark is displayed beside the plot in the contextual menu and the plot is displayed in the detail and data list panels.

k. To show a previously hidden plot:

1. Right-click on the data list panel header row and choose Column Display. The Column Display form opens.
2. Select the plot to show in the Available for Table list and click the right arrow. The selected plot moves to the Displayed on Table list.
3. Click OK. The plot information is restored in the detail and data list panels.

12

To change the parameters of a plot:

1. Select the plot entry in the configuration panel and click Clear. The plot is removed from the detail and data list panels.
2. Customize the plot parameters, as described in [Step 2](#) to [Step 6](#).

13

To remove a plot entirely:

1. Select the plot entry in the configuration panel and click Clear. The plot is removed from the detail and data list panels.
2. Click Remove. The plot is removed from the Statistics Plotter window.

14

To stop all plotting, click Stop. The plotting of each counter stops.

15

To switch from real-time to historical plotting:

1. Click Stop.
2. Configure the historical plot start and end times to specify the time range of the plot.
3. Click Historical Plot. The detail panel clears, and the statistics are retrieved from the database and plotted.

16

To switch from historical to real-time plotting:

1. Choose a polling interval from the real-time polling interval drop-down menu, or type a value between 10 and 3600 seconds.
2. Click Real-time Plot. The detail panel clears and real-time statistics plotting begins.

17

To save the plotter view:



Note: Only the graphics information in the detail panel is saved. To change the view in the detail panel before you save the results, use the slider or the zoom buttons.

If the Legend check box is selected, the legend is saved with the statistics graph results.

1. Click Save Current View. The Save As form opens.
2. Use the form to specify a file location and file type.
3. Click Save. The Save As form closes, and the graph is saved in the specified file.

18

To save the statistics table results, if required:

1. Right-click on the plot value list heading and choose Save To File. The Save form opens.
2. Use the form to specify a file location and file type.
3. Click Save. The Save As form closes, and the tabulated plot information is saved in the specified file.

19

Close the Statistics Plotter window.



Note: The Statistics Plotter window cannot be saved. When you close a Statistics Plotter window, the data in the window is deleted.

END OF STEPS

9.4 To modify a statistics graph

9.4.1 Steps

1

Perform one of the following:

- a. Use a contextual menu option. Right-click on a map or network navigation tree object and choose Plotter→# Statistics Plotter (#). The Statistics Plotter form opens.
- b. Open a Statistics Plotter form. Choose Tools→Statistics→Statistics Plotter→# Statistics Plotter (#) from the NFM-P main menu. The Statistics Plotter form opens.
- c. Add an object to the Statistics Plotter form:
 1. Choose the object on which to modify a statistics graph.
 2. Locate the object in the network navigation tree.
 3. Right-click on the object in the navigation tree and choose Properties. The object properties form opens.
 4. Click on the Statistics tab.
 5. Choose a statistics class from the object drop-down menu.

Note:

The icon beside each class indicates the plot types that the class supports. See [Figure 9-2, “Plotting support indicators”](#) (p. 112) for the icon descriptions.

6. Click Plotter and choose Statistics Plotter (#). The Statistics Plotter form opens with the current plot information displayed.

Note:

A statistics graph can plot up to eight statistics counters. The statistics counters can be same only for different objects. If only one object is the source of the statistics, each counter must be unique.

If the Statistics Plotter (#) menu item is not available, the Statistics Plotter does not support the selected object type.

2

Perform [Step 3 to Step 19](#) of [9.3 “To create a statistics graph”](#) (p. 116) to configure the plot parameters.

END OF STEPS

9.5 To plot utilization statistics

9.5.1 Steps

1

Open a new plot for the object, as described in [Step 1](#) of [9.3 “To create a statistics graph”](#) (p. 116) .

2

Use the Statistics Group drop-down menu to specify Interface Additional Stats.

3

Perform one of the following:

- a. Use the Statistics Counter drop-down menu to specify Ingress Utilization.
- b. Use the Statistics Counter drop-down menu to specify Egress Utilization.

4

Start the plotting of real-time or historical utilization statistics, as described in [Step 9](#) of [9.3 “To create a statistics graph”](#) (p. 116) .

END OF STEPS

9.6 To create a customized plotter profile

9.6.1 Purpose

Perform this procedure to create a customized plotter profile.



Note: Only a user with the Admin or Statistics Plotter Profile Management scope of command role can create a custom profile.

9.6.2 Steps

1

Perform one of the following:

- a. If the object of the profile is a dynamic LSP:
 1. Open the dynamic LSP properties form.
 2. Click on the Statistics tab.
 3. Click Plotter button, or click More Actions and choose Plotter if the Plotter button is dimmed.
- b. If the object of the profile is not a dynamic LSP, right-click on the object, for example, a physical port in the network navigation tree, and choose Plotter.

2

Perform one of the following:

- a. Choose New Plot from the Plotter contextual menu. The Statistics Plotter window opens.
- b. Choose *profile_name*→New Plot.

3

Customize the plot, as described in [Step 3 to Step 19 of 9.3 “To create a statistics graph” \(p. 116\)](#).

4

Click Save Profile(s). The Plotter Profile (Create) form opens.

5

Configure the required parameters.

6

To change the Y axis of a plot in the profile:

1. Click on the Plots tab.
2. Select a plot in the list and click Properties. The Plot Profile Entry (Edit) form opens.
3. Configure the Y Axis parameter.
4. Click OK. The Plot Profile Entry (Edit) form closes.

7

Click OK. The Plotter Profile (Create) form closes. The new plotter profile is available for objects of the same type.

END OF STEPS

9.7 To plot statistics using a plotter profile

9.7.1 Steps

1

Perform [9.6 "To create a customized plotter profile" \(p. 123\)](#) to create a customized plotter profile, if the required profile does not exist.

2

Perform one of the following:

- a. If the object of the plot is a dynamic LSP:
 1. Open the dynamic LSP properties form.
 2. Click on the Statistics tab.
 3. Click Plotter button, or click More Actions and choose Plotter if the Plotter button is dimmed.
- b. If the object of the plot is a link in a topology map:
 1. Right-click on the link and choose Expand Group.
 2. Right-click on the link and choose Plotter.
- c. Right-click on a map or network navigation tree object, and choose Plotter.

3

Choose *profile_name*→New Plot from the Plotter contextual menu. The Statistics Plotter opens and begins to plot the statistics data for the counters in the profile.

4

When you are finished viewing the plot, close the Statistics Plotter window.

END OF STEPS

9.8 To manage plotter profiles

9.8.1 Steps

1

To list all plotter profiles that apply to a specific type of object, perform one of the following:

-
- a. For a dynamic LSP:
 1. Open the dynamic LSP properties form.
 2. Click on the Statistics tab.
 3. Click on the Plotter button, or click More Actions and choose Plotter if the Plotter button is dimmed, and choose View All Profiles.
 - b. For an object in the network navigation tree, right-click on the object and choose Plotter→View All Profiles.

The View All Profiles form opens. The form lists the default and customized plotter profiles that are selectable for the object type.

2

To list all plotter profiles, choose Tools→Statistics→Statistics Plotter Profiles from the NFM-P main menu. The Statistics Plotter Profiles form opens.

3

To delete a profile:



Note: Only a user with the Admin or Statistics Plotter Profile Management scope of command role can delete a custom profile.
You cannot delete a default profile.

1. Select the profile in the list and click Delete. A dialog box appears.
2. Click Yes. The profile is deleted and removed from the list.

4

Close the View All Profiles form.

END OF STEPS

Part IV: Statistics content

Overview

Purpose

This volume contains a description of the statistics record format for supported statistics types in the NFM-P.

Contents

Chapter 10, Statistics record format
--

129

10 Statistics record format

10.1 Statistics record format in the NFM-P

10.1.1 Supported statistics types

Table 10-1 NFM-P supported statistics types

Statistics type	Description
Performance	Performance statistics provide categorized information about network throughput. Performance statistics are SNMP-based, and acquired by sending SNMP queries to network elements.
Accounting	There are three main classes of accounting policy: service, network, and subscriber. Accounting statistics are aggregated in files on network elements, and then transferred to the NFM-P on request.
Server performance	Server performance statistics include NFM-P component performance indicators regarding memory usage, JMS activity, alarm activity, SNMP trap receipt, statistics collection and processing, and NE resynchronization.
Internal statistics	Internal statistics are NFM-P process performance counters that Nokia technical support use for system troubleshooting.

10.2 Performance statistics

10.2.1 Overview

Performance statistics provide categorized information about network throughput, and are obtained using SNMP. The information in a performance statistics record includes the following:

- collection timestamp
- collection record type
- source device and object identifiers
- Suspect indicator that indicates a problem with the collected statistics; for example, when an NE is unresponsive to SNMP requests
- statistics counters that contain the raw data

When the Suspect indicator in a record is set, the counters in the record contain the following values:

- periodic counters—zero
- non-periodic counters—the value of the same counter in the most recent non-suspect record

After a suspect record collection, the next non-suspect record may contain a large value that is seen as a dramatic increase. Such an increase skews the data value for the interval and may result in misinterpretation. To alleviate this, an operator can configure a System Preferences parameter to increase, rather than reset, the Periodic Time value after the collection of a suspect record. For example, for a 15-minute collection interval, if one record is suspect and the next is not, the periodic time of the second record is recorded as 30 minutes. In this way, the NFM-P presents the first non-suspect record collected after a collection failure using a more realistic average value. See the *NSP System Administrator Guide* for information about setting system preferences.

Table 10-2 Performance statistics record properties

Property	Description
Monitored Object	Unique identifier of the object from which the statistics are collected
Monitored Object Name	Name of the object from which the statistics are collected
Periodic Time	Number of milliseconds since previous statistics collection
Record Type	The type of record, based on the following collection types: <ul style="list-style-type: none"> Scheduled Full Performed according to a user-configurable schedule On-demand Performed immediately at the request of an NFM-P operator
Site ID	NE identifier
Site Name	NE name
Suspect	Indicates a statistics collection failure.
Time Captured	Unix Epoch time at which the statistics were collected and the NFM-P adds the log to the database
Time Logged	Time at which the NFM-P receives the SNMP get-response

Table 10-3 Performance statistics data types and examples

NE data type	NFM-P data type	Example
Counter32	long	32-bit counter such as Received Octets in the Interface Stats group
Gauge32		
Unsigned32		
TimeStamp		
Counter64	UINT128	64-bit counter such as Received Broadcast Packets in the Interface Additional Stats group
Integer32	double	Radio output power, in tenths of dBm; NFM-P converts to dBm

Table 10-3 Performance statistics data types and examples (continued)

NE data type	NFM-P data type	Example
Integer32	float	Port output power, in tenths of μ W
Integer32	int	Small value, such as index, enumerator, or identifier
TruthValue	Boolean	State indicator
various	String	MAC address, object name

10.3 Accounting statistics

10.3.1 Overview

Accounting statistics are available for network ports, SAPs, SDPs, and subscribers. Service accounting statistics typically provide service-usage data for billing requirements. Network accounting statistics provide service-quality data for SLA QoS compliance requirements. Subscriber accounting statistics provide subscriber profile usage information for billing and SLA compliance.

The information in an accounting statistics record includes the following:

- collection timestamp
- collection record type
- source device and object identifiers
- statistics counters that contain the raw data

Table 10-4 Accounting statistics record properties

Property	Description
Accounting Policy ID	The NFM-P accounting policy ID, which is optional; inclusion can be enabled from the NFM-P System Preferences form, and is disabled by default
Final Count	Contains a value greater than zero when a subscriber or SAP is deleted after the previous collection. The value represents the ordinal of the deleted object; for example, if the object is the first object deleted since the previous collection, the value is 1; the property has a value of 2 in the record of the second object deleted during the same collection interval, and so on.
Forwarding Class	Forwarding class associated with a queue
Lag Port	Identifies a LAG port
Monitored Object	Detailed information about the object for which the statistics are collected

Table 10-4 Accounting statistics record properties (continued)

Property	Description
Monitored Object Name	Name of the object for which the statistics are collected
Periodic Time	Not applicable to accounting statistics
Port Id	Identifies a port
Queue Id	Identifies a hardware queue
Record Type	Type of collection, for example, Scheduled Full
Sap Id	Physical port identifier portion of a SAP definition
Sdp Id	Identifies an SDP
Site ID	Identifies a site
Site Name	Identifies the name of a site
SlaProfile Id	Identifies the value assigned to an SLA profile
Group ID	Identifies the ISA-AA group
Partition ID	Identifies the partition ID within an ISA-AA group
Subscriber Id	Identifies a subscriber
Subscriber Profile Id	Identifies a subscriber profile
Svc Id	Identifies a service
Time Captured	Time when the NFM-P starts to process the data file collected from the NE
Time Recorded	Time when the statistics were collected on the NE

All accounting statistics share certain common properties, with some additional properties for each statistic type.

Table 10-5 Accounting statistics common properties

Statistic Type	Properties
All	<ul style="list-style-type: none"> • Time Captured • Record Type • Monitored Object • Monitored Object Name • Site Name • Site ID • Periodic Time • Time Recorded

Table 10-5 Accounting statistics common properties (continued)

Statistic Type	Properties
Service	<ul style="list-style-type: none"> • Sap Id • Svc Id • Sdp Id • Queue Id • Final Count
Network	<ul style="list-style-type: none"> • Forwarding Class • Lag Port • Port Id • Queue Id
Subscriber	<ul style="list-style-type: none"> • Subscriber Id • Subscriber Profile Id • Sla Profile Id
AA	<ul style="list-style-type: none"> • Group ID • Partition ID

Table 10-6 Accounting statistics counter definitions

Counter name	Definition
All Octets Offered All Packets Offered	Offered packets enter a queue and are dropped or forwarded. An octet value is a packet value expressed in Bytes.
Forwarding Class	The forwarding class of the traffic; this is a classification rather than a counter
High Octets Dropped High Octets Offered High Pkts Dropped High Pkts Offered High Priority Octets Dropped High Priority Octets Offered High Priority Pkts Dropped High Priority Pkts Offered	High-priority packets for a traffic class are marked as in-profile on the ingress classification or based on the forwarding class of the packet. High-priority forwarding classes include nc, h1, ef, and h2. See the appropriate node's QoS guide for more information about forwarding class scheduler mapping. Offered packets enter a queue and are dropped or forwarded. Dropped packets are not forwarded through a queue. An octet value is a packet value expressed in Bytes.

Table 10-6 Accounting statistics counter definitions (continued)

Counter name	Definition
In Profile Octets Dropped In Profile Octets Forwarded In Profile Packets Forwarded In Profile Pkts Dropped In Profile Pkts Forwarded	In-profile packets are received at a rate that is lower than the queue CIR. Forwarded packets are sent to an egress queue. Dropped packets are not forwarded through a queue. An octet value is a packet value expressed in Bytes.
Long Duration Flows	Subscriber host application sessions that are longer than 180s in duration
Low Octets Dropped Low Octets Offered Low Pkts Dropped Low Pkts Offered Low Priority Octets Dropped Low Priority Octets Offered Low Priority Pkts Dropped Low Priority Pkts Offered	Low-priority packets for a traffic class are marked as out-of-profile on the ingress classification or based on the forwarding class of the packet. Low-priority forwarding classes include l1, af, l2, and be. See the appropriate node's QoS guide for more information about forwarding class scheduler mapping. Offered packets enter a queue and are dropped or forwarded. Dropped packets are not forwarded through a queue. An octet value is a packet value expressed in Bytes.
Medium Duration Flows	Subscriber host application sessions that are 180s or less in duration
Net to Sub Active Flows Net to Sub Admit Flows Net to Sub Admit Octets Net to Sub Admit Pkts Net to Sub Deny Flows Net to Sub Deny Octets Net to Sub Deny Pkts	Net to Sub traffic flows from the network to a subscriber. An active flow is a flow that is opened, closed, opened and closed, or continued during a collection interval. Admit objects are objects that are forwarded. Deny objects are objects that are not forwarded. An octet value is a flow or packet value expressed in Bytes.
Num of Subs	The number of subscribers that have the specified flow type
Out Of Profile Octets Dropped Out Of Profile Octets Forwarded Out of Profile Packets Forwarded Out Of Profile Pkts Dropped Out Of Profile Pkts Forwarded	Out-of -profile packets are received at a rate that is higher than the queue CIR. Forwarded packets are sent to an egress queue. Dropped packets are not forwarded through a queue. An octet value is a packet value expressed in Bytes.
Short Duration Flows	Subscriber host application sessions that are 30s or less in duration

Table 10-6 Accounting statistics counter definitions (continued)

Counter name	Definition
Sub to Net Active Flows Sub to Net Admit Flows Sub to Net Admit Octets Sub to Net Admit Pkts Sub to Net Deny Flows Sub to Net Deny Octets Sub to Net Deny Pkts	Sub to Net traffic flows from a subscriber to the network. An active flow is a flow that is opened, closed, opened and closed, or continued during a collection interval. Admit objects are objects that are forwarded. Deny objects are objects that are not forwarded. An octet value is a flow or packet value expressed in Bytes.
Total Flow Duration	Aggregate duration, in seconds, of the completed flows at the time of collection.
Total Flows Completed	The number of completed flows in each direction
Total Octets Dropped Total Octets Forwarded Total Packets Dropped Total Packets Forwarded	Offered packets enter a queue and are dropped or forwarded. Forwarded packets are sent to an egress queue. Dropped packets are not forwarded through a queue. An octet value is a packet value expressed in Bytes.
Total Term Flows	The number of completed flows in each direction
Uncoloured Octets Offered Uncoloured Packets Offered Uncoloured Pkts Forwarded Uncoloured Pkts Offered	Octets and packets can be modified or colored using color profiles. Color profiling adds the ability to selectively treat packets received on a SAP as in-profile or out-of-profile regardless of the queue forwarding rate. This allows a customer or access device to color a packet out-of-profile with the intention of keeping the in-profile bandwidth for higher-priority packets. Offered packets enter a queue and are dropped or forwarded. Forwarded packets are sent to an egress queue. An octet value is a packet value expressed in Bytes.

Table 10-6 Accounting statistics counter definitions (continued)

Counter name	Definition
Ingress Aggregate Pkts Offered	Aggregate number of high priority, low priority and uncolour packets offered by the ingress Pchip to Qchip.
Ingress Aggregate Octets Offered	Aggregate number of high priority, low priority and uncolour octets offered by the ingress Pchip to Qchip.
Ingress Aggregate Pkts Forwarded	Aggregate number of in and out-of profile packets (rate above CIR) forwarded by the ingress Qchip.
Ingress Aggregate Octets Forwarded	Aggregate number of in and out-of profile octets (rate above CIR) forwarded by the ingress Qchip.
Ingress Aggregate Pkts Dropped	Aggregate number of high and low priority packets discarded by the ingress Qchip.
Ingress Aggregate Octets Dropped	Aggregate number of high and low priority octets discarded by the ingress Qchip.
Egress Aggregate Pkts Forwarded	Aggregate number of in and out-of profile packets (rate above CIR) forwarded by the egress Qchip.
Egress Aggregate Octets Forwarded	Aggregate number of in and out-of profile octets (rate above CIR) forwarded by the egress Qchip.
Egress Aggregate Pkts Dropped	Aggregate number of in and out-of profile packets discarded by the egress Qchip.
Egress Aggregate Octets Dropped	Aggregate number of in and out-of profile octets discarded by the egress Qchip.
Stats Rate Mbps	Rate Statistics in megabits per second.
Stats Rate Mbps Periodic	Rate Statistics in megabits per second since the last period.
Stats Rate Packets	Rate Statistics in packets per second.
Stats Rate Packets Periodic	Rate Statistics in packets per second since the last period.

10.4 Server performance statistics

10.4.1 Overview

Server performance statistics provide NFM-P system information regarding memory usage, alarm counters, and network management activity. The information in a performance statistic record includes the following:

- collection timestamp
- collection record type
- source object identifiers
- statistics counters that contain the raw data

Table 10-7 Server performance statistics properties

Properties	Description
Monitored Object	Detailed information about the object for which the statistics are collected
Monitored Object Name	Name of the object for which the statistics are collected
Periodic Time	Time difference between one data record collection period and the next
Record Type	Type of collection, for example, Scheduled Full or On-demand
Time Captured	Time when the statistics were collected

Table 10-8 Monitored NFM-P servers per statistics policy

Server performance statistics policy	Monitored servers
Alarm Rate	Main server
Application Assurance Statistics Collection	Main server Auxiliary server
Assurance Event Rate	Main server
CPAM BGP AS Events	Main server
JMS Durable Sessions	JMS server
JMS Topic	JMS server
NFM-P Event Log	Main server
Node Resync	Main server
OAM PM event	Main server Auxiliary server
PCMD Collector (obsolete)	Auxiliary server
Publisher CPAM Event	Main server
Publisher MAP Event	Main server
Publisher Object Event	Main server
Publisher Queue Event	Main server
Publisher Realtime Event	Main server
Publisher XML Event	Main server
Server Memory	Main server Auxiliary server JMS server

Table 10-8 Monitored NFM-P servers per statistics policy (continued)

Server performance statistics policy	Monitored servers
SNMP Trap	Main server
Statistics Collection	Main server Auxiliary server
STM Event	Main server
TCA Rate	Main server
XML API Find to File Request	Main server
XML API Request	Main server
XML Find Request	Main server

Table 10-9 Server performance statistics counter information

Server performance statistics class	Counter name	Counter description
Alarm Rate	Alarm Total	Total number of alarms
	Cleared	Count of cleared alarms
	Condition	Count of condition alarms
	Critical	Count of critical alarms
	Indeterminate	Count of indeterminate alarms
	Info	Count of information alarms
	Major	Count of major alarms
	Minor	Count of minor alarms
	Warning	Count of warning alarms

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
Application Assurance Statistics Collection (subset of Statistics Collection)	Maximum Processing Time	Maximum processing time for AA statistics files during the collection interval
	Minimum Processing Time	Minimum processing time for AA statistics files during the collection interval
	Total Processing Time	Total processing time for AA statistics files during the collection interval
	Average Processing Time	Average processing time for AA statistics files during the collection interval
	AA Stats Failure	Number of AA statistics files that have not been processed due to errors related to transfer or parsing. This count is also included in the AA Stats Failure statistic under the generic Statistics Collection record.
	AA Stats Pending	Number of AA statistics data rows that have been read but not processed. This count is also included in the AA Stats Pending statistic under the generic Statistics Collection record.
	AA Stats Processed	Number of AA statistics data rows processed and converted to NFM-P statistics records. This count is also included in the AA Stats Processed statistic under the generic Statistics Collection record.
	AA Stats Total	Total number of AA statistics data rows read from accounting files. This count is also included in the AA Stats Total statistic under the generic Statistics Collection record.
	File AA Stats Pending	Number of AA files transferred from NEs that are waiting to be processed. This count is also included in the File AA Stats Pending statistic under the generic Statistics Collection record.
	File AA Stats Processed	Number of AA statistics files that have been read and processed into statistics records. This count is also included in the File AA Stats Processed statistic under the generic Statistics Collection record.
	File AA Stats Total	Number of AA statistics files that have been transferred from NEs. This count is also included in the File AA Stats Total statistic under the generic Statistics Collection record.

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
JMS Durable Sessions	Filtered Event Vessel Average Size	Average size of filtered event vessels sent during subscriber session
	Filtered Event Vessel Count	Count of filtered event vessels sent during subscriber session
	Filtered Event Vessel Max Size	Size of largest filtered event vessels sent during subscriber session
	Filtered Event Vessel Min Size	Size of smallest event vessel sent during subscriber session
	Current Message Count	Total messages queued for durable XML subscriber session at collection time
JMS Topic	Max Topic Message Count	The maximum message count in a subscription queue for a topic
	Min Topic Message Count	The minimum message count in a subscription queue for a topic
	Subscription Count	The total subscription count for a topic
	Current Topic Message Count	Total messages in all subscription queues for a topic
Node Resync	Scheduled Resync Failure	Count of failed scheduled resynchronizations
	Scheduled Resync Processed	Count of processed scheduled resynchronizations
	Scheduled Resync Received	Count of scheduled resynchronizations received
	Unscheduled Resync Failure	Count of unscheduled resynchronizations
	Unscheduled Resync Processed	Count of processed unscheduled resynchronizations
	Unscheduled Resync Received	Count of unscheduled resynchronizations received

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
XML API Find Request	Average Request Response Time	Average find request execution time
	Maximum Request Response Time	Maximum find request execution time
	Minimum Request Response Time	Minimum find request execution time
	Requests Failed	Count of find requests that failed execution
	Requests Passed	Count of successfully executed find requests
	Requests Pending	Count of find requests that are being executed
	Requests Received	Count of received find requests; each request is executed upon receipt

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
XML API Find To File Request	Asynchronous Requests Failed	Count of asynchronous findToFile requests that failed execution or were rejected due to queue size exceeded
	Asynchronous Requests Passed	Count of successfully queued and executed asynchronous findToFile requests
	Asynchronous Request sPending	Count of asynchronous findToFile requests that are queued or being executed
	Asynchronous Requests Queued	Count of queued asynchronous findToFile requests
	Asynchronous Requests Received	Count of received asynchronous findToFile requests; Asynchronous requests are queued while awaiting execution
	Asynchronous Requests Rejected Maximum Queue Size Exceeded	Count of asynchronous findToFile requests rejected because the maximum queue size is exceeded
	Average Request Response Time	Average findToFile request execution time; queued time is included for asynchronous requests
	Maximum Request Response Time	Maximum findToFile request execution time; queued time is included for asynchronous requests
	Minimum Request Response Time	Minimum findToFile request execution time; queued time is included for asynchronous requests
	Synchronous Requests Failed	Count of synchronous findToFile requests that failed execution
	Synchronous Requests Passed	Count of successfully executed synchronous findToFile requests
	Synchronous Requests Pending	Count of synchronous findToFile requests that are being executed
	Synchronous Requests Received	Count of received synchronous findToFile requests; each synchronous request is executed upon receipt

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
XML API Request	Average Request Response Time	Average request execution time; includes execution time of failed requests, and excludes execution time of asynchronous findToFile requests
	Maximum Request Response Time	Maximum request execution time; includes execution time of failed requests, and excludes execution time of asynchronous findToFile requests
	Minimum Request Response Time	Minimum request execution time; includes execution time of failed requests, and excludes execution time of asynchronous findToFile requests
	Requests Failed	Count of failed XML API requests, excluding failed asynchronous findToFile requests
	Requests Passed	Count of successfully executed XML API requests
	Requests Pending	Count of XML API requests that are being executed, excluding queued and executing asynchronous findToFile requests
	Requests Received	Count of received XML API requests; each request is executed immediately upon receipt
	Requests Rejected Maximum Connections Exceeded	Count of requests rejected because the maximum number of connections is exceeded
	Requests Rejected Maximum User Connections Exceeded	Count of requests rejected because the maximum number of connections for one user is exceeded
Server Memory	Committed Heap Memory	Total amount of memory allocated for use by code
	Committed Non Heap Memory	Initial amount of memory allocated to the heap
	Init Heat Memory	Initial amount of memory allocated to the heap
	Init Non Heap Memory	Initial amount of memory allocated for use by code
	Max Heap Memory	Maximum number of kbytes occupied by the heap
	Max Non Heap Memory	Maximum number of kbytes occupied by the code
	Used Heap Memory	Number of kbytes currently occupied by the heap
	Used Non Heap Memory	Number of kbytes currently occupied by code

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
SNMP Trap	Dropped Backpressure	Count of traps dropped due to back pressure from the server
	Dropped Duplicate	Count of duplicate traps dropped
	Dropped Full Resync	Count of traps dropped during a full resynchronization
	Dropped Not Managed	Count of traps dropped from unmanaged NEs
	Dropped Out of Sequence	Count of out of sequence traps dropped
	Dropped Trap Disabled	Count of disabled traps dropped
	Dropped Trap Unspecified	Count of unspecified traps dropped
	Incoming	Count of incoming traps

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
Statistics Collection	Accounting Stats Failure	Number of accounting statistics files that have not been processed due to errors related to transfer or parsing
	Accounting Stats Pending	Number of accounting statistics data rows that have been read but not processed
	Accounting Stats Processed	Number of accounting statistics data rows processed and converted to NFM-P statistics records
	Accounting Stats Total	Total number of accounting statistics data rows read from accounting files
	File Accounting Stats Pending	Number of accounting files transferred from NEs and waiting to be processed
	File Accounting Stats Processed	Number of accounting statistics files that have been read and processed into NFM-P statistics records
	File Accounting Stats Total	Total number of accounting statistics files that have been transferred from the NEs
	Scheduled Polling Stats Pending	Number of scheduled performance statistics rows read from the NEs and waiting to be processed
	Scheduled Polling Stats Processed	Number of scheduled performance statistics rows processed and saved to the database
	Scheduled Polling Stats Records	Number of scheduled performance statistics records created through statistics processing; in some case, processing one row of statistics from an NE might result in multiple NFM-P statistics records
	Scheduled Polling Stats Total	Total number scheduled performance statistics rows read from NEs
	Scheduled Stats Failure	Number of scheduled performance statistics collection attempts that have failed due to errors related to reading values from the NE
	Unscheduled Polling Stats Pending	Number of on-demand performance statistics rows read from the NEs and waiting to be processed
	Unscheduled Polling Stats Processed	Number of on-demand performance statistics rows processed and converted to NFM-P statistics records

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
Statistics Collection	Unscheduled Polling Stats Records	Number of on-demand performance statistics records created through statistics processing; in some cases, processing one row of statistics from an NE might result in multiple NFM-P statistics records
	Unscheduled Polling Stats Total	Total number of on-demand performance statistics rows read from NEs
	Unscheduled Stats Failure	Number of on-demand performance statistics collection attempts that have failed due to errors related to reading values from the NE
STM Event	Files Received Count	Number of files received
	Result Processed Count	Number of processed results
	Files Received Count Periodic	Number of files received since the last period
	Result Processed Count Periodic	Number of processed results since the last period

Table 10-9 Server performance statistics counter information (continued)

Server performance statistics class	Counter name	Counter description
TCA	Alarm Skipped	Number of TCAs not executed because the maximum number of alarms is reached
	Alarmed	Number of TCAs executed
	Cleared	Number of TCAs cleared
	Custom	Number of custom TCAs created for execution
	Custom Executed	Number of custom TCA rules executed
	Custom Pending	Number of custom TCAs created but not yet executed
	Delta	Number of TCAs raised because of a delta
	Drop	Number of drop TCAs created for execution
	Drop Executed	Number of drop TCA rules executed
	Drop Pending	Number of drop TCAs created but not yet executed
	Error	Number of error TCAs created for execution
	Error Executed	Number of error TCA rules executed
	Error Pending	Number of error TCAs created but not yet executed
	Utilization	Number of utilization TCAs created for execution
	Utilization Executed	Number of utilization TCA rules executed
	Utilization Pending	Number of utilization TCAs created but not yet executed

10.5 Nokia internal statistics

10.5.1 Overview

Internal statistics are server performance statistics that track information related to core NFM-P processes. The information provided by these statistics is best used in conjunction with Nokia technical support. For more information about internal statistics, please contact your Nokia support representative. [Table 10-10, “Internal statistics” \(p. 148\)](#) describes the internal statistics.

10.5.2 NFM-P publisher events

Several internal statistics track the publishing of event messages in the NFM-P, and the status of the publisher queues. Event messages are in XML format and carry information about events in the NFM-P or changes in the database - for example, a user deleting one or more network objects triggers the creation of a Delete Event Vessel event, containing information about the deleted

objects. Event messages are sent to the publisher queues and distributed throughout the NFM-P system. The statistics in this section provide information about event messages and the status of the publisher queue. For more information about using NFM-P XML events, see the [Network Developer Portal](#).

Some internal statistics are simple counts of the number of event messages published during the collection period, and may not directly correlate to the number of times the triggering event occurred. For example, deleting several objects simultaneously publishes a single Delete Event Vessel event containing all the deleted objects.

Table 10-10 Internal statistics

Server performance statistics class	Counter name	Counter description
Publisher Map Event	Edge List Event Count	Count of Edge List Events published. This event is triggered when an edge group is modified, for example when the group acquires new child objects.
	Edge Status Event Count	Count of Edge Status Events published. This event is triggered when the status of an edge group changes.
	Incremental Event Count	Count of Incremental Event Count events published. This event is triggered when a client requests incremental command notification.
	Incremental Highlight Event Count	Count of Incremental Highlight Events published. This event is triggered when a client requests incremental command notification for highlighting.
	Map Service Event Count	Abstract count of all publisher map service events. For more information about publisher map service events, contact your Nokia service representative.
	Vertex Status Event Count	Count of Vertex Status Events published. This event is triggered when the status of a vertex changes.
	Vertex Topology Event Count	Count of Vertex Topology Events published. This event is triggered by changes to a topology map, for example when an edge, vertex or group is created or modified.
	Vertex Tree Event Count	Count of Vertex Tree Events published. This event is triggered by changes to a topology tree, for example when a vertex or group is added to or removed from the tree.
	View Event Count	Count of View Events published. This event is triggered by updates to the topology map view and sending view messages.

Table 10-10 Internal statistics (continued)

Server performance statistics class	Counter name	Counter description
Publisher Object Event	Attribute Value Change Context Event Count	Count of Attribute Value Change Context events published. This event is triggered by changing the attributes of an object, and contains additional attribute information for context.
	Delete Event Vessel Count	Count of Delete Event Vessel events published. This event is triggered when an object is deleted, and contains lists of objects to be deleted.
	Event Vessel Count	Count of event lists sent using the EventVessel event.
	File Available Event Count	Count of File Available Events events published. This event is triggered when the target of an asynchronous findToFile request becomes available.
	Olc Update Completed Event Count	Count of Olc Update Completed events published. This event is triggered by successful updates to the OLC State of an object.
	Text Message Event Count	Count of Text Message Events published. This event is triggered when an NFM-P text message is sent to a user using the text message function.
Publisher Queue	Object Event In Publisher Queue	Number of outstanding events in the object publisher queue
	Reatime Event In Publisher Queue	Number of outstanding events in the real-time publisher queue
	Xml Event In Publisher Queue	Number of outstanding events in the XML publisher queue
Publisher Realtime Event	Realtime Data Communication Failure Count	Count of realtime statistic collection failures caused by a communication error with the node
	Realtime Data Event Count	Count of Realtime Data Events published. This event is triggered when a graph data point is sent.
	Realtime Data Failure Count	Count of realtime statistic collection failures caused by the node failing to retrieve a statistic value.
	Realtime Data Missed Event Count	Count of Realtime Data Missed Events published. This event is triggered when a realtime statistic collection fails due to the NFM-P not receiving the data within the configured collection interval.

Table 10-10 Internal statistics (continued)

Server performance statistics class	Counter name	Counter description
Publisher XML Event	Alarm Status Change Event Count	Count of Alarm Status Change Events. This event is triggered by changes to the status of an alarm or alarm aggregate.
	Attribute Value Change Event Count	Count of Attribute Value Change Events published. This event is triggered by changing the attributes of an object, and contains information about the changes.
	DB Activity Event Count	Count of database switchovers and failover notifications sent using the DBActivityEvent event. This event is triggered when a switchover or failover begins, and again when it ends.
	Db Connection State Change Event Count	Count of Db Connection State Change Events published. This event is triggered by changes to the database connection state, such as the connection going down or being placed on standby.
	Db Error Event Count	This counter is not currently used by any function.
	Db Proxy State Change Event Count	Count Db Proxy State Change Events published. This event is triggered by changes to the status of the NFM-P database proxy.
	Deployer Event Count	Count of Deployer Events published. This event is triggered by attempting to deploy an object.
	Event Vessel Average Size	Average size of event vessels
	Db Error Event Count	This counter is not currently used by any function.
	Db Proxy State Change Event Count	Count Db Proxy State Change Events published. This event is triggered by changes to the status of the NFM-P database proxy.
	Deployer Event Count	Count of Deployer Events published. This event is triggered by attempting to deploy an object.
	Event Vessel Average Size	Average size of event vessels
	Event Vessel Count	Count of event vessels sent
	Event Vessel Max Size	Size, in bytes, of largest event vessel sent
	Event Vessel Min Size	Size, in bytes, of smallest event vessel sent

Table 10-10 Internal statistics (continued)

Server performance statistics class	Counter name	Counter description
Publisher XML Event	Exception Event Count	Count of generic exception events
	Filter Change Event Count	Count of changes to the registration filter for a subscribed client
	Keep Alive Event Count	Count of Keep Alive events, sent every five minutes to the OSS.
	Log File Available Event Count	Count of LogFileAvailableEvent events published. This event is triggered when the target of a LogToFile request becomes available.
	Managed Route Event Count	Count of changes to routes managed by the CPAM
	Object Creation Event Count	Count of objects created in the NFM-P
	Object Deletion Event Count	Count of objects deleted in the NFM-P
	Relationship Change Event Count	Count Relationship Change Events published. This event is triggered by changes to the relationships between objects, such as adding or deleting a relationship.
	Script Execution Event Count	Count of scripts executed
	State Change Event Count	Count of stateChangeEvent events sent, containing changes to the state of an object.
	Stats Event Count	Count of Stats Events published. This event is triggered when a polling interval begins or ends.
	Terminate Client Session Count	Count of Terminate Client Session events published. This event is sent to close a client session.

