NOKIA

NSP
Network Services Platform

Network Functions Manager - Packet (NFM-P)

Release 22.3

Planning Guide

3HE-18148-AAAA-TQZZA

Issue 1

March 2022

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.

Disclaimer

Open Source Software and Red Hat Enterprise Linux Operating System

In case:

- (i) any "Open Source Software and Red Hat Enterprise Linux Operating System ("FOSS & RHEL") is packaged separately or integrated with any Nokia Software and to which third party license obligations apply; or,
- (ii) any FOSS & RHEL is directly licensed by Customer under a separate license or subscription agreement, and such FOSS & RHEL is interacting or interoperating with any Nokia Software or Product:

information will be available either in the FOSS & RHEL itself or on the website from which the download is available indicating the license under which such FOSS was released, and containing required acknowledgements, legends and/or notices.

It is hereby acknowledged and agreed by the Parties that any FOSS & RHEL is distributed on an "as is" basis under the respective FOSS & RHEL license terms. Nokia will not warrant nor will be liable for, and will not defend, indemnify, or hold Customer harmless for any claims arising out of, or in any case related to FOSS & RHEL and their use (or inability to use) by the Parties. This includes, but is not restricted to, any and all claims for direct, indirect, incidental, special, exemplary, punitive or consequential damages in connection with FOSS & RHEL or its components (whether included in the Nokia Software or not) and their use or inability to use. Also, this includes claims for or in connection with the title in, the non-infringement of or interferences and damages caused to Customer or third parties by FOSS & RHEL.

CUSTOMER SHALL HAVE NO RIGHT TO RECEIVE FROM NOKIA ANY CARE (MAINTENANCE & SUPPORT SERVICES) ON FOSS &

RHEL LICENSED BY CUSTOMER UNDER A SEPARATE LICENSE AGREEMENT OR SUBSCRIPTION CONTRACT AND TO WHICH THIRD PARTY LICENSE OBLIGATIONS APPLY WHETHER OR NOT IT INTERACTS WITH ANY NOKIA SOFTWARE OR PRODUCT.

The above shall also apply in case Customer requires - and Nokia accepts under the terms of this Disclaimer to use its reasonable commercial effort to do so - certain installation services on FOSS & RHEL as directly licensed by Customer under a separate license or subscription agreement; and, such FOSS & RHEL are interacting or interoperating with a Nokia Software or Product. For sake of clarity in such a case the following shall also apply:

- Before starting any installation service, Customer must instruct Nokia to start such installation and must confirm in writing to Nokia that its FOSS & RHEL license or subscription contract (for RHEL: Red Hat Enterprise Agreement) with Customer includes the right to use of the specific FOSS & RHEL and the related support for all platforms running the FOSS & RHEL; that such subscription and support contract is in force (not expired) and allows such installation activities.
- Nokia will not warrant nor will be liable for any cost, expense, damage, and will not defend, indemnify, or hold Customer or any third party harmless for any claims arising out of, or in any case related to FOSS & RHEL (and in connection with the installation activities of such FOSS & RHEL) and their use (or inability to use) by the Customer or by any third party, following the installation of such FOSS & RHEL. This includes, but is not restricted to, any and all claims for direct, indirect, incidental, special, exemplary, punitive or consequential damages in connection with FOSS & RHEL or its components and their use or inability to use.
- Nokia will not provide nor will have any liability or obligation to provide any support, maintenance, care service, warranty or indemnity with respect to any (installed) FOSS & RHEL as licensed by the Customer under a separate license agreement or subscription contract.

Any Care service (maintenance and support service) on FOSS & RHEL licensed by Nokia as packaged separately or integrated with any Nokia Software may be made available by Nokia under specific contractual terms to be agreed upon by the Parties.

Contents

About this document

Purpose

This document consolidates the technical information related to the deployment of the NSP NFM-P Release 21. This document does not focus on the functionality offered by NFM-P Release 21, but instead presents the reader with pre-installation information required to plan a successful deployment.

The *NSP NFM-P Planning Guide* does not include a comprehensive list of technologies supported or not supported by NFM-P or the platforms hosting it. The Nokia NSM Product Group should be consulted for clarification when uncertainty exists.

The *NSP NFM-P Planning Guide* details the following aspects of the NFM-P:

- Product deployment overview
- Supported operating systems specifications
- Platform requirements
- Network requirements
- Scaling guidelines
- Workstation configuration
- Firewall information

Intended audience

This document is intended for network engineers, planners and IT staff who are familiar with the functionality of the NFM-P and are planning a product deployment.

Document support

Customer documentation and product support URLs:

- Doc Center
- Technical support

How to comment

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

1 Product deployment overview

1.1 NFM-P architecture

1.1.1 NFM-P platforms

Several types of platforms can be present in an NFM-P deployment:

- NFM-P client
- NFM-P client delegate
- NFM-P server
- NFM-P auxiliary (statistics collector, call trace collector, PCMD collector)
- NSP flow collector controller
- NSP flow collector
- NFM-P auxiliary database
- NFM-P database
- NSP analytics server

NFM-P supports collocation of the NFM-P server and NFM-P database software on a single workstation or VM.

NFM-P also supports a distributed deployment, whereby the NFM-P server and the NFM-P database software components are installed on two different workstations or VMs.

An NFM-P auxiliary can be configured for statistics collection, call trace collection, or PCMD collection but can only be configured to perform one of these functions.

NFM-P supports the distribution of statistics collection through the use of one or multiple NFM-P auxiliary statistics collectors. Statistics collection with an NFM-P auxiliary uses either the NFM-P database or the NFM-P auxiliary database for statistics record retention.

The NFM-P auxiliary database can be deployed as a single instance or in a cluster of a minimum of three servers or VMs. The server BIOS CPU frequency scaling must be disabled on any platform hosting the NFM-P auxiliary database.

The NSP analytics server is deployed on a separate workstation or VM and is used in conjunction with the auxiliary database cluster to generate custom analytics reports.

NFM-P supports redundancy of the NFM-P server, NFM-P database, NFM-P auxiliary collector, and NSP analytics server workstations. The NFM-P auxiliary statistics collector supports 3+1 redundancy.

When the NFM-P auxiliary database is deployed in a cluster of at least three separate instances, it can tolerate a single instance failure with no data loss. All NFM-P auxiliary database nodes in the same cluster must be deployed in the same geographic site, with less than 1ms of latency between the nodes. A second cluster can be deployed to implement geographic redundancy and must contain the same number of nodes as the primary cluster.

An NFM-P auxiliary statistics collector must be installed on an independent workstation or VM and can only be configured in a distributed NFM-P deployment.

An NFM-P auxiliary call trace collector must be installed on an independent workstation or VM to collect the call trace data from CMM and CMG network elements. Up to two active NFM-P auxiliary call trace collectors can be installed to scale the collection of call trace data. Each active NFM-P auxiliary call trace collector can be assigned to a redundant collector where call trace data is synchronized between the redundant collector pairs. The NFM-P auxiliary call trace collectors can be configured in either a distributed or collocated NFM-P deployment.

The NSP flow collector and NSP flow collector controller can be installed on the same server or VM provided that a only a single active NSP flow collector is required. In cases where multiple active NSP flow collectors are required, the NSP flow collector and NSP flow collector controller must be installed on independent workstations or VMs. The NSP flow collector does not use a traditional redundancy model. Instead, the network elements can be configured to send cflowd data to multiple NSP flow collectors. The NSP flow collector controller can be configured as either standalone or redundant. The NSP flow collector and NSP flow collector controller can be configured in either a distributed or collocated NFM-P deployment.

An NFM-P auxiliary PCMD collector must be installed on an independent workstation or VM to collect PCMD data streams from the network. The NFM-P auxiliary PCMD collector can be configured in either a distributed or collocated NFM-P deployment and is supported in a redundant configuration.

More details on redundancy in NFM-P can be found in 1.3 “Redundancy architecture” (p. 17).

NFM-P supports IPv4 and IPv6 connectivity between the NFM-P server/auxiliary and NSP flow collector to the managed network. Connectivity between the NFM-P components fully supports IPv4.

Connectivity between NFM-P components supports IPv6 with certain restrictions where the following is not supported:

- NFM-P deployments that include the management of CMM, and eNodeBs
- NFM-P deployments with NFM-P auxiliary PCMD collector(s)
- EMS integration with NFM-P
- Dual stack between NFM-P components, except clients

A network element can only be managed by one NFM-P standalone or redundant deployment. Having multiple NFM-P deployments managing the same network element is not supported, and will cause unexpected behavior.

The following illustrates a typical deployment of NFM-P in standalone mode when the NFM-P server and NFM-P database functions are collocated.

Figure 1-1 NFM-P standalone deployment - collocated NFM-P server/database configuration

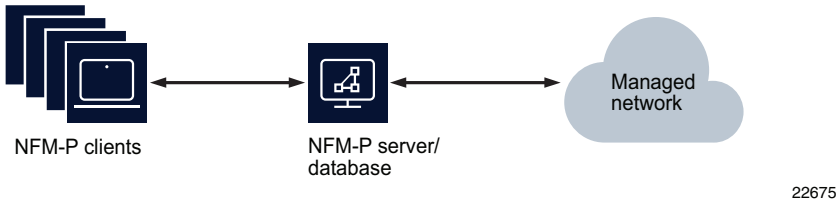

The following illustrates a typical deployment of NFM-P in standalone mode when the NFM-P server and NFM-P database functions are not collocated.

Figure 1-2 NFM-P standalone deployment – distributed NFM-P server and NFM-P database configuration.

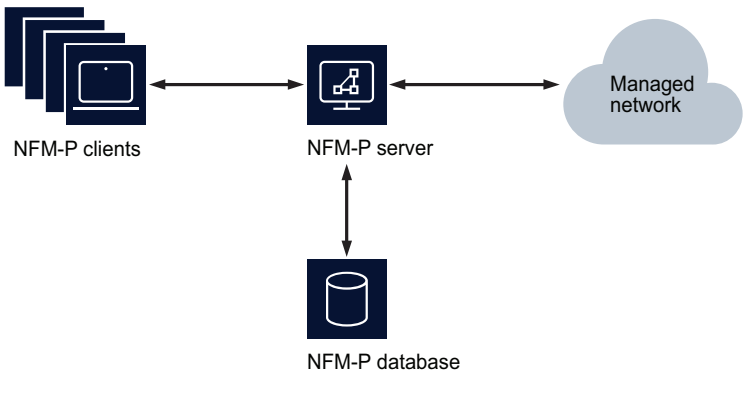

The following illustrates a typical deployment of NFM-P in standalone mode when the NFM-P server and NFM-P database functions are collocated and an NFM-P auxiliary call trace collector is used. The NFM-P auxiliary statistics collector is not supported in this configuration.

Figure 1-3 NFM-P standalone deployment – collocated NFM-P server and NFM-P database configuration and NFM-P auxiliary call trace collector

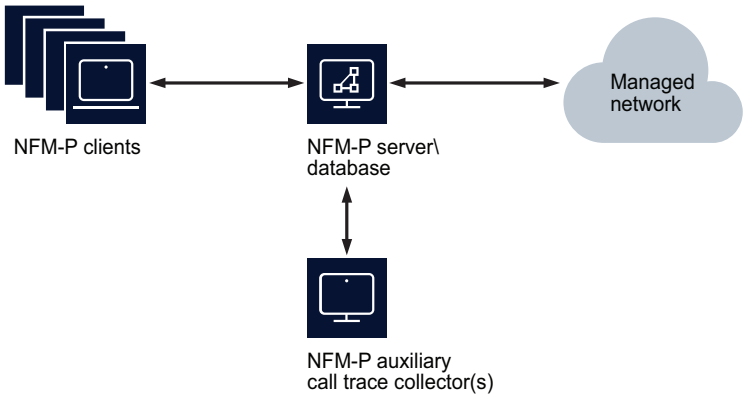

The following illustrates a typical deployment of NFM-P in standalone mode when the NFM-P server and NFM-P database functions are in a distributed configuration and NFM-P auxiliary collectors are used. In this configuration there can be up to three active NFM-P auxiliary statistics collectors or it could be configured redundant, and there can be one or two NFM-P auxiliary call trace collectors collecting call trace data from the network.

Figure 1-4 NFM-P standalone deployment - distributed NFM-P server and NFM-P database configuration and NFM-P auxiliary collectors

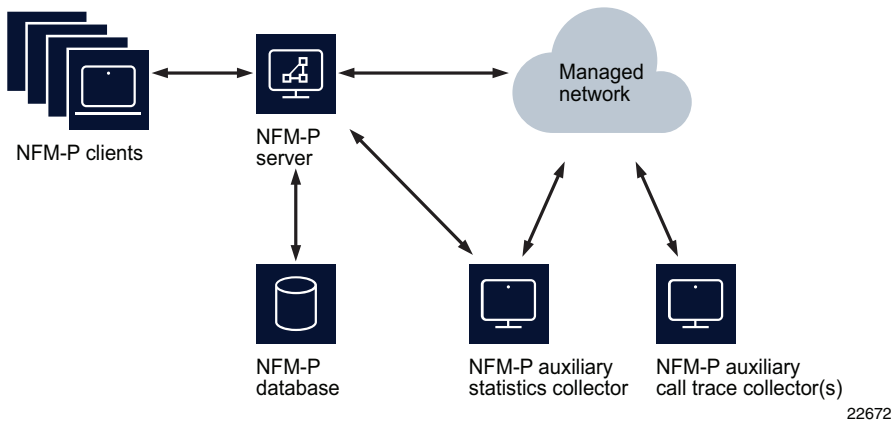

The following illustrates a deployment of NFM-P in standalone mode when the NFM-P server and NFM-P database functions are in a distributed deployment and NFM-P auxiliary collectors are installed with statistics collection using the NFM-P auxiliary database. In this configuration, there can be up to three preferred NFM-P auxiliary statistics collectors. There can be one or two NFM-P auxiliary call trace collectors collecting call trace data from the network. The NFM-P auxiliary database cluster comprises of either a single instance or a cluster of at least three instances.

Figure 1-5 NFM-P standalone deployment - distributed NFM-P server and NFM-P database configuration and NFM-P auxiliary collectors with statistics collection using the NFM-P auxiliary database

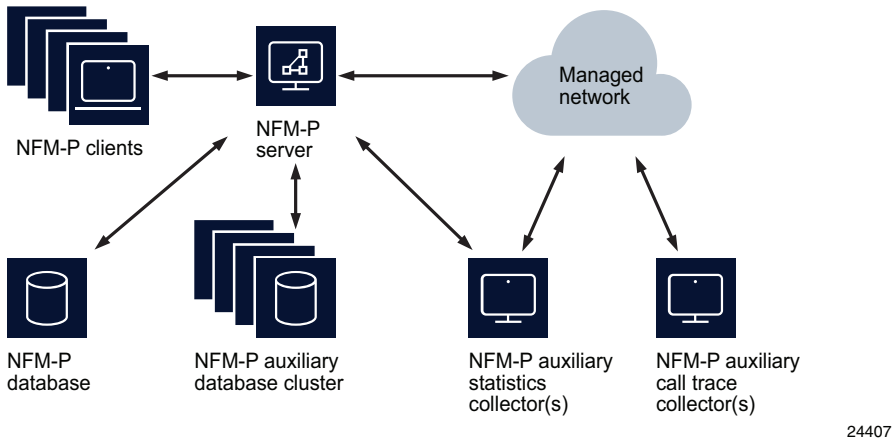

For bare metal installations, the NFM-P server, NFM-P auxiliary collector, NSP flow collector, NSP flow collector controller, NFM-P auxiliary database, NSP analytics server, and NFM-P database are supported on specific Intel x86 based HP and Nokia AirFrame workstations.

The NFM-P client and client delegate software may be installed on workstations running different operating systems from the NFM-P server, NFM-P auxiliary, NFM-P auxiliary database, NSP flow collector, NSP flow collector controller, NSP analytics server, and NFM-P database. The NFM-P client can be installed on RHEL 7 server x86-64, Windows, or Mac OS where the NFM-P client delegate can be installed on RHEL 7 server x86-64, Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019. Refer to Chapter 2, “Operating systems and third party software” for Operating System specifics.

All NFM-P workstations in the NFM-P management complex must maintain consistent and accurate time. It is recommended that NTP be used to accomplish this requirement.

1.1.2 NFM-P auxiliary statistics collector

This type of NFM-P auxiliary collects and processes performance, accounting, application assurance accounting and performance management statistics along with OAM PM test results. This option enables customers to remove the collection and processing load of these record types from the NFM-P server while allowing for increased collection capabilities. An NFM-P auxiliary statistics collector should be used when collection is expected to exceed the capacity of the NFM-P server. Refer to Chapter 3, “Platform requirements” for scalability details of the NFM-P server and dimensioning of the NFM-P auxiliary statistics collector.

The NFM-P auxiliary statistics collector can be configured as Preferred, Reserved, or Remote for a given NFM-P server (active or standby). This allows for a local and geographically redundant NFM-P auxiliary statistics collector configuration. When collecting statistics using the NFM-P auxiliary database or using logToFile only, up to three NFM-P auxiliary statistics collectors can

collect statistics concurrently. Information on the redundancy model of the NFM-P auxiliary statistics collector can be found in 1.3 “Redundancy architecture” (p. 17).

The NFM-P server and the NFM-P auxiliary statistics collector must maintain consistent and accurate time. It is encouraged to use an NTP service to achieve this. An alarm will be raised if the times are not within 30 seconds. Variations in time can cause the system to stop collecting statistics prematurely.

In networks where NFM-P auxiliary statistics collectors are not configured, the NFM-P server handles the statistics collection. In networks where the NFM-P auxiliary statistics collector is configured, the NFM-P server will never collect statistics – regardless of the availability of the NFM-P auxiliary statistics collectors. At least one NFM-P auxiliary statistics collector must be available for statistics collection to occur.

The NFM-P auxiliary statistics collector is only supported with a distributed NFM-P server and NFM-P database.

For collection of performance management statistics from eNodeB network elements, NTP should be used to synchronize the network element and the NFM-P server and NFM-P auxiliary statistics collector to ensure the statistics are successfully retrieved.

1.1.3 NFM-P auxiliary call trace collector

This type of NFM-P auxiliary collects call trace files from CMM and CMG network elements.

Up to two NFM-P auxiliary call trace collectors can be configured to collect call trace data, and each of those collectors can be configured to be redundant. Each NFM-P auxiliary call trace collector is installed on a separate workstation. Each NFM-P auxiliary call trace collector is configured as a preferred for the NFM-P active server and as a reserved for the NFM-P standby server. This allows for a redundant NFM-P auxiliary call trace collector configuration. Only one of the workstations in the NFM-P auxiliary call trace collector redundant pair will collect the call trace information from the network elements at any given time and the call trace information is synchronized between the Preferred and Reserved pair of workstations. Information on the redundancy model of the NFM-P auxiliary call trace collector can be found in 1.3 “Redundancy architecture” (p. 17).

The NFM-P auxiliary call trace collector is supported with both a collocated NFM-P server and NFM-P database or distributed NFM-P server and NFM-P database.

The NFM-P auxiliary call trace collector must maintain consistent and accurate time. It is encouraged to use an NTP service to achieve this. An alarm will be raised if the times are not within 30 seconds.

1.1.4 NSP flow collector

The NSP flow collector collects AA cflowd data or system cflowd (IPFIX/Netflow/CGNAT) data from network elements but cannot collect both.

The NSP flow collector operates in a standalone mode to collect and generate data for the managed network elements that are configured to send flow data to it. Each network element can be configured to send flow data to multiple NSP flow collectors.

The NSP flow collector is supported with both a collocated NFM-P server and NFM-P database or distributed NFM-P server and NFM-P database.

The NSP flow collector must maintain consistent and accurate time. It is encouraged to use an NTP service to achieve this.

1.1.5 NSP flow collector controller

The NSP flow collector controller extracts network data from the NFM-P database to be distributed to the NSP flow collectors.

There is only one active NSP flow collector controller in the network where a second server or VM can be added for redundancy.

The NSP flow collector controller must maintain consistent and accurate time. It is encouraged to use an NTP service to achieve this.

1.1.6 NFM-P auxiliary PCMD collector

The NFM-P auxiliary PCMD collector collects PCMD data from SR based mobile gateways that are configured to stream per call measurement data to it. The NFM-P auxiliary PCMD collector generates CSV files from the PCMD data, that can be sent to a third party application for post processing.

The NFM-P auxiliary PCMD collector is supported with both a collocated NFM-P server and NFM-P database or distributed NFM-P server and NFM-P database.

The NFM-P auxiliary PCMD collector workstation must maintain consistent and accurate time. It is encouraged to use an NTP service to achieve this.

1.1.7 NFM-P client delegate

This option enables customers to launch multiple NFM-P GUI clients from a single Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or RHEL 7 server x86-64 workstation. For RHEL server x86-64 installations, these clients can be displayed using the X11 protocol to other RHEL desktops or native X displays. For Windows Server installations, these GUI clients can be displayed using Windows Remote Desktop. Displaying GUI clients to computers running X-emulation software is not currently supported.

The client delegate platform provides an option to consolidate multiple installations of the NFM-P GUI client on a single workstation. Individual NFM-P clients can be installed on the client delegate. The NFM-P client also supports the ability for multiple users to share a single installation; however, each user must run the client with a unique UNIX id.

Information on dimensioning the NFM-P client delegate platform is given in 3.3 “Minimum platform requirements” (p. 34).

1.2 NFM-P key technologies

1.2.1 Java Virtual Machine

NFM-P applications use Java technology. The installation packages contain a Java Virtual Machine which is installed with the software. This is a dedicated Java Virtual Machine and does not conflict with other Java Virtual Machines which may be installed on the same workstation.

NFM-P uses Java Virtual Machine version 8 from Oracle and OpenJDK 8.

1.2.2 Oracle database

The NFM-P database embeds an installation of Oracle 19c Enterprise Edition, which is installed with the NFM-P database. This database is used to store information about the managed network. The installation of Oracle is customized for use with the NFM-P application and must be dedicated to NFM-P. NFM-P database redundancy uses Oracle DataGuard, and is configured in maximum performance mode.

Nokia will not support any configuration deviations from the Oracle installation as performed by the NFM-P database installation package, as it represents an NFM-P License Agreement Violation. Modifying the Oracle installation can impact system performance, stability and upgrades. Customer support agreements may be violated.

Access to the Oracle database is restricted to the NFM-P application. Direct user access to the database is strictly forbidden.

The Oracle database is embedded with NFM-P and because of this, Oracle requires all licenses to be purchased from Nokia. This applies to customers with Oracle Site licenses as well.

Oracle’s official support position for running Oracle database 19c, embedded within NFM-P, on VMware hosted virtual environments is described in Oracle Support Note 249212.1. Oracle will provide support for those running on VMware virtualized environments. In addition, VMware has a public statement committing to assist with resolving Oracle database issues. Nokia will work with Oracle and VMware to resolve any NFM-P database issues but problem resolution times may be impacted in some cases.

Oracle's official support position for running Oracle Database 19c, embedded within NFM-P, on RHEL KVM hosted virtual environments is that Oracle does not certify any of their products in this environment. Nokia will work with Oracle and Red Hat to resolve any NFM-P database issues but due to the lack of Oracle support, problem resolution times may be impacted in some cases. Customers should be aware of and must accept this risk when choosing to run NFM-P in a RHEL KVM virtualized environment.

1.2.3 Other databases

Embedded within the NFM-P server is a Neo4j database and a PostgreSQL database and embedded within the NFM-P auxiliary database is a Vertica database.

Similar to the Oracle support statement, Nokia will not support any configuration deviations from the installation as performed by the installation package, as it represents an NFM-P License Agreement Violation. Modifying the installation can impact system performance, stability and upgrades. Customer support agreements may be violated..

Access to the PostgreSQL, Neo4j, and Vertica databases is restricted to the NFM-P application. Direct user access to any of the databases is strictly forbidden.

In a redundant configuration, the active NFM-P server hosts the primary PostgreSQL and Neo4j databases. The standby NFM-P server hosts the standby PostgreSQL and Neo4j databases.

1.3 Redundancy architecture

1.3.1 Overview

Redundancy between NFM-P server and database applications is used to ensure visibility of the managed network is maintained when one of the following failure scenarios occur:

- Loss of physical network connectivity between NFM-P server and/or NFM-P database and the managed network
- Hardware failure on workstation hosting the NFM-P server and/or NFM-P database software component

NFM-P supports redundancy of the NFM-P server and NFM-P database components in the following configurations:

- NFM-P server and NFM-P database collocated configuration
- NFM-P server and NFM-P database distributed configuration

The following illustrates an NFM-P redundant installation when the NFM-P server and NFM-P database components are installed in a collocated configuration.

Figure 1-6 NFM-P collocated server/database redundancy deployment

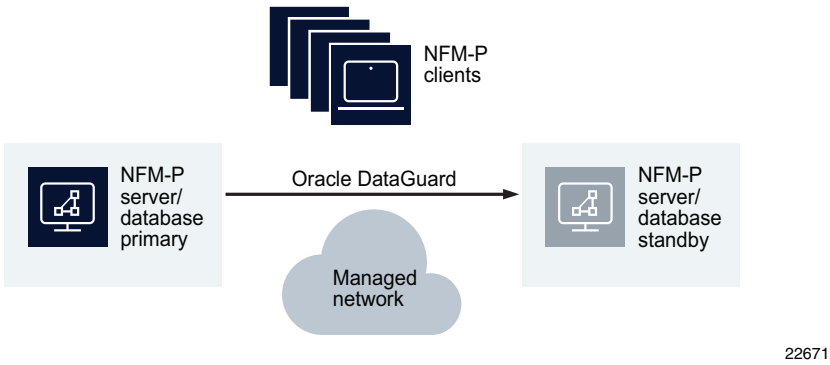

The following illustrates an NFM-P redundant installation when the NFM-P server and NFM-P database components are located on different workstations in a distributed configuration.

Figure 1-7 NFM-P distributed server/database redundancy deployment in a geographically redundant setup.

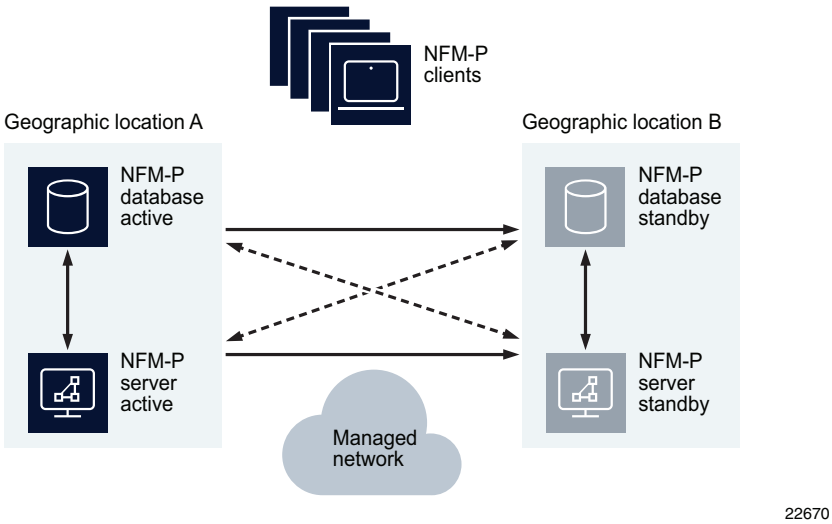

1.3.2 Redundancy and NFM-P auxiliaries and NSP flow collectors

In customer networks, where the statistics collection requirements exceed the scalability capabilities of an NFM-P server, the NFM-P auxiliary statistics collector can be used. As with other high availability components, the NFM-P auxiliary statistics collector can be configured to be redundant. When collecting statistics using the NFM-P database, each NFM-P server can be configured to have one preferred and one reserved NFM-P auxiliary statistics collector. When collecting statistics using the NFM-P auxiliary database or using logToFile only, each NFM-P server can be configured with up to three preferred and one reserved NFM-P auxiliary statistics collector.

When call trace information is being collected from CMM and CMG network elements in customer networks, an NFM-P auxiliary call trace collector must be used. The NFM-P auxiliary call trace collector can be installed in a redundant pair where up to two NFM-P auxiliary call trace collector redundant pairs can be installed.

In customer networks where cflowd data is to be collected from network elements, an NSP flow collector must be used. The NSP flow collector can only be installed in a standalone configuration. To achieve data redundancy, network elements can be configured to forward cflowd data to multiple NSP flow collectors. The NSP flow collector controller can be installed in a standalone or redundant configuration.

For the collection of per-call measurement data from SGW/PGW network elements, an NFM-P auxiliary PCMD collector must be used. The NFM-P auxiliary PCMD collector can be installed in a standalone or redundant configuration. Data collected by the NFM-P auxiliary PCMD collector is not replicated to the redundant collector.

In Figure 1-8, “NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries that crosses geographic boundaries” (p. 20) there are NFM-P auxiliary collectors configured. In the example where redundancy is geographic, there can be up to four NFM-P

auxiliary statistics collectors and up to two NFM-P auxiliary call trace collector workstations configured in each geographic location. The Preferred/Reserved/Remote role of the NFM-P auxiliary statistics collector is dependent upon and configured, based on the NFM-P server that is active. When there are more than one active auxiliary statistics collector, local redundancy (Preferred/Reserved) of the auxiliary statistics collector must be used in conjunction with geographic redundancy, where the same number of auxiliary statistics collectors will be deployed in each geographic site. The NFM-P auxiliary statistics collectors in the opposite geographic location are configured to be Remote. In this scenario, if one of the NFM-P auxiliary statistics collectors for the active NFM-P server were no longer available, the active NFM-P server would use the reserved NFM-P auxiliary statistics collector in the same geographic location to collect statistics. Figure 1-9, “NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries using the NFM-P auxiliary database for statistics collection” (p. 21) shows the same redundant configuration but with statistics collection using a geographically redundant NFM-P auxiliary database. Latency between geographic sites must be less than 200 ms.

Figure 1-8 NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries that crosses geographic boundaries

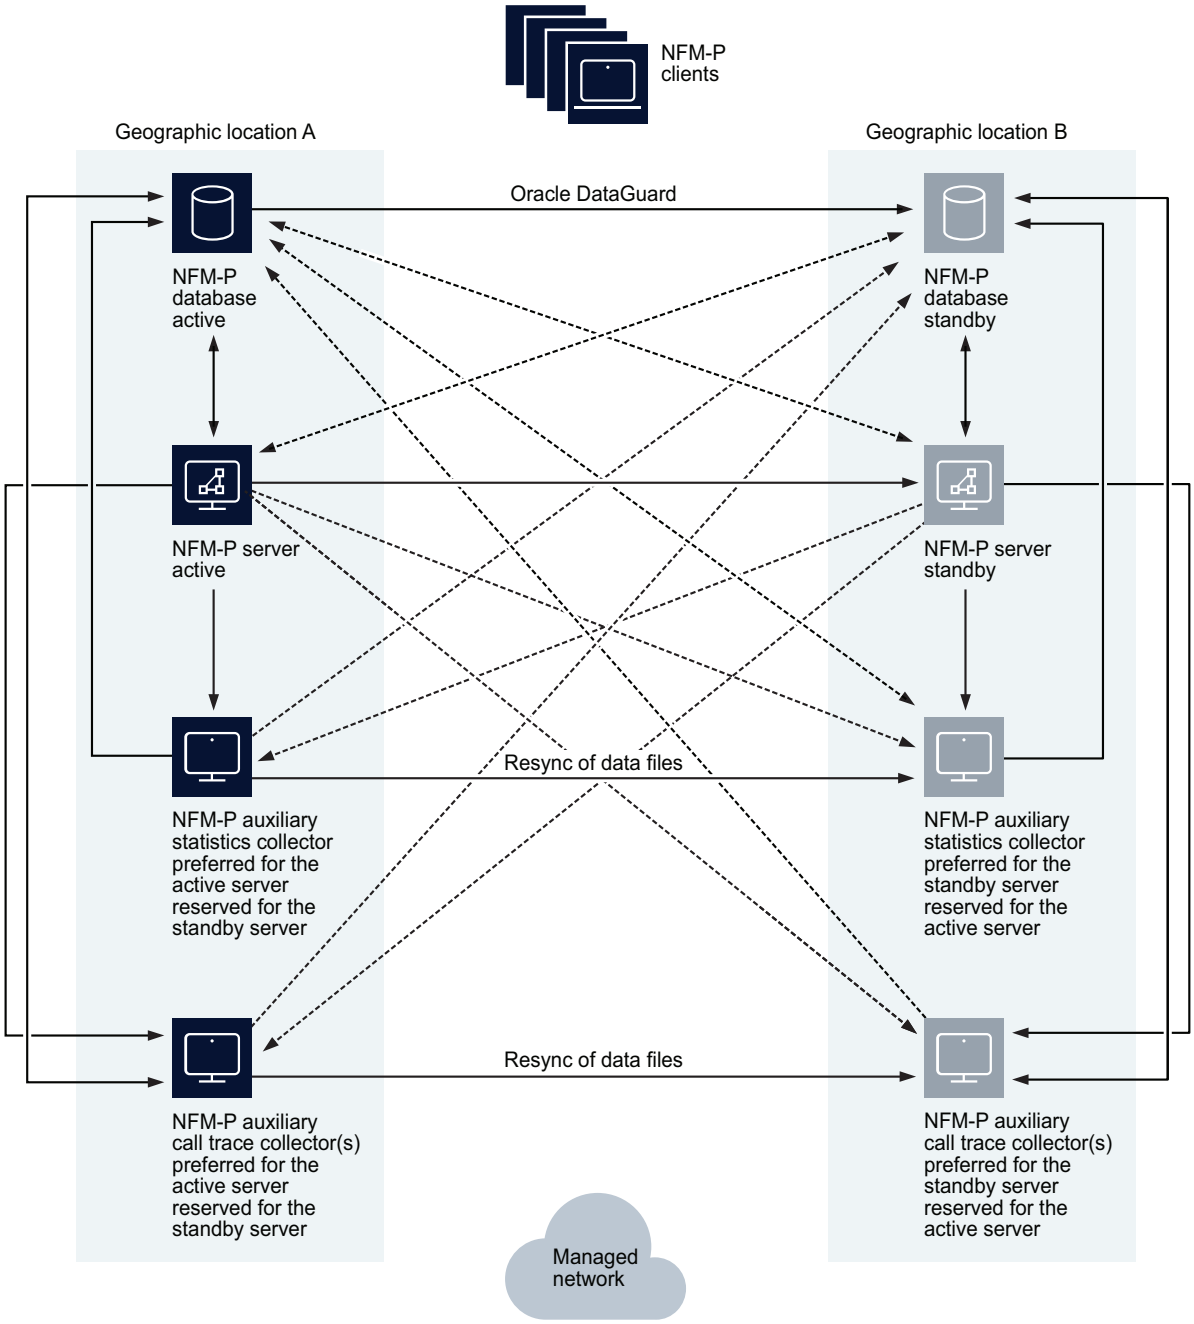

Figure 1-9 NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries using the NFM-P auxiliary database for statistics collection

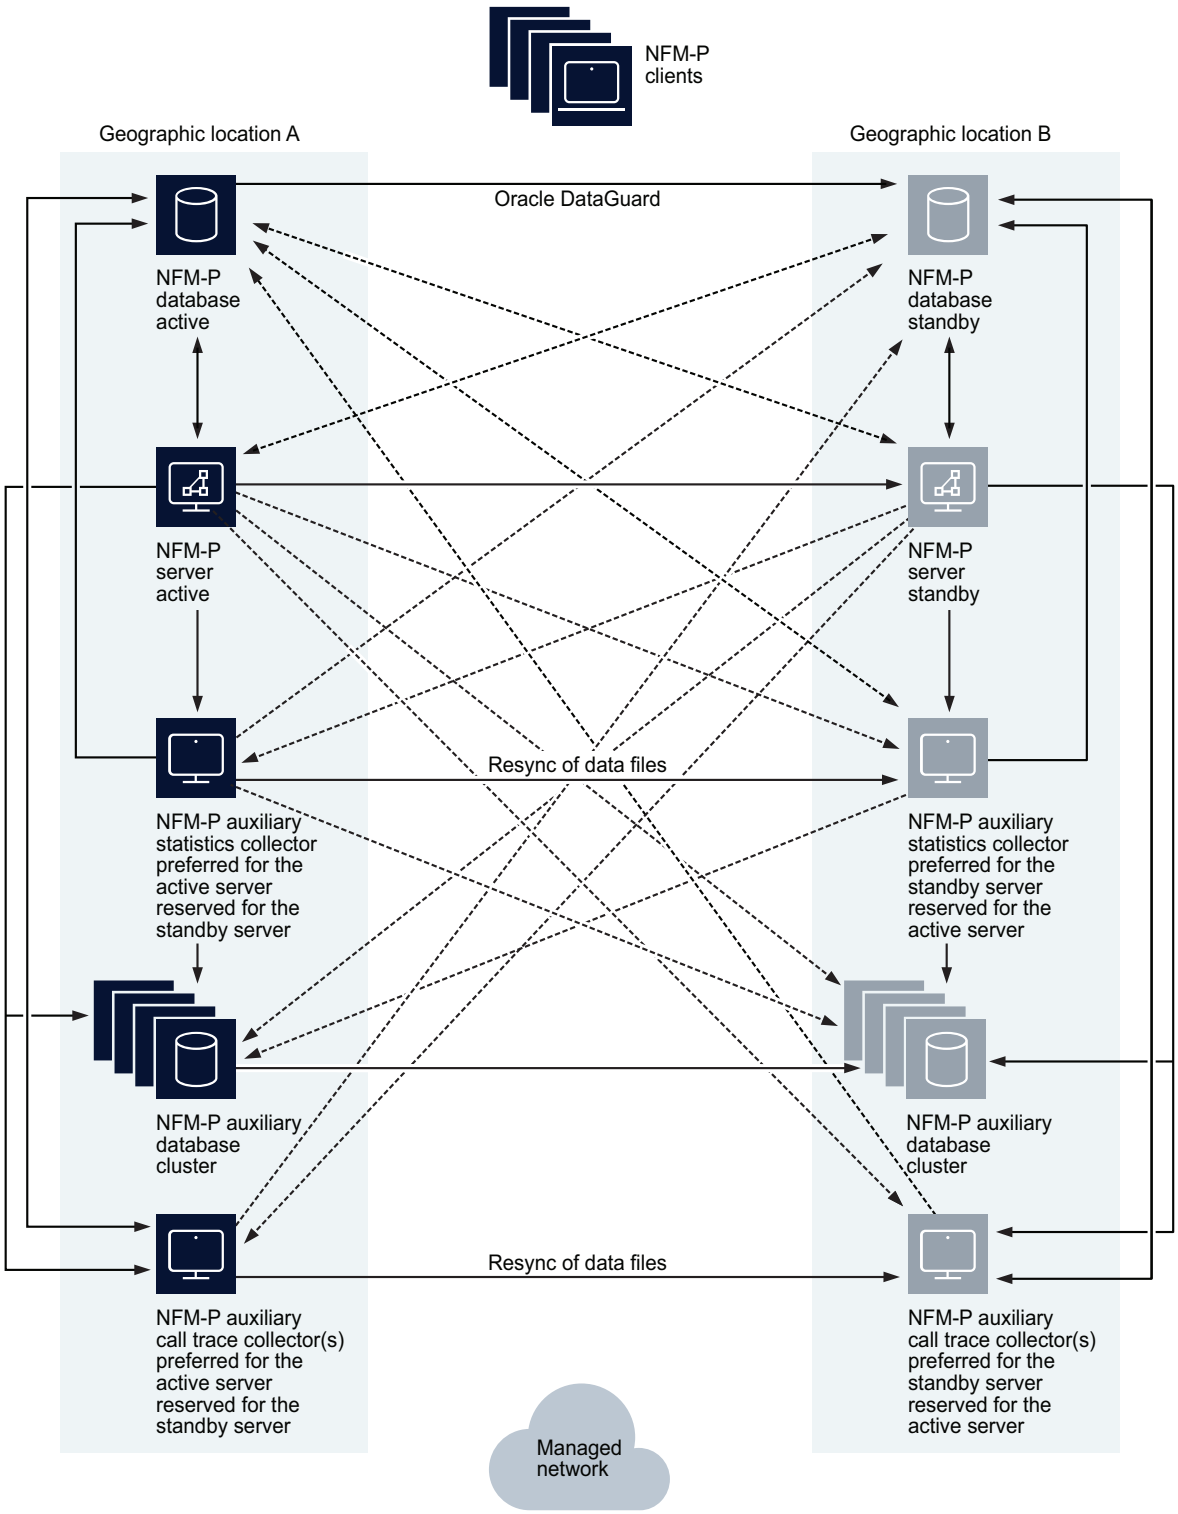

Further information about NFM-P redundancy can be found in the *NSP NFM-P User Guide*.

1.4 Redundancy deployment considerations for NFM-P

1.4.1 Overview

When deploying NFM-P in a redundant configuration, the following items should be considered.

It is a best practice to keep the NFM-P server, NFM-P database, and NFM-P auxiliary collectors in the same geographic site to avoid the impact of network latency. When the NFM-P database or NFM-P server switches sites, the NFM-P auto-align functionality will ensure the NFM-P server, and NFM-P auxiliary collectors are all aligned in the same geographic location. If the auto-align functionality is not enabled, a manual switch of the workstations is desirable.

1.4.2 Redundancy with collocated NFM-P server/database

Requirements:

- The operating systems installed on the primary and standby NFM-P server/database must be of the same versions and at the same patch levels.
- The layout and partitioning of the disks containing the NFM-P software, the Oracle software and the database data must be identical on the active and standby NFM-P server/database.
- The machine which will be initially used as the active NFM-P server/database must be installed or upgraded before the machine that will initially be used as the standby.
- The workstations hosting the NFM-P software should be connected in a way to prevent a single physical failure from isolating the two workstations from each other.
- Workstations running the NFM-P server/database software must be configured to perform name service lookups on the local workstation before reverting to a name service database located on the network such as NIS, NIS+, or DNS. A root user must inspect and modify the /etc/nsswitch.conf file to ensure that files is the first entry specified for each database listed in the file.

1.4.3 Redundancy with distributed NFM-P server and NFM-P database

Requirements:

- The operating systems installed on the primary and standby NFM-P server as well as the primary and standby NFM-P database must be of the same versions and at the same patch levels.
- The layout and partitioning of the disks containing the NFM-P software, the Oracle software and the database data must be identical on the primary and standby NFM-P database.
- The machines which are intended to be used as primary NFM-P server and NFM-P database should be installed on the same LAN as one another with high quality network connectivity.
- The machines which are intended to be used as standby NFM-P server and standby NFM-P database should be installed on the same LAN as one another with high quality network connectivity.
- The pair of workstations to be used as active NFM-P server and NFM-P database should be

connected to the pair of workstations to be used as standby NFM-P server and NFM-P database in a way that will prevent a single physical failure from isolating the two workstation pairs from each other.

- Workstations running the NFM-P server and NFM-P database software must be configured to perform name service database lookups on the local workstation before reverting to a name service located on the network such as NIS, NIS+, or DNS. A root user must inspect and modify the /etc/nsswitch.conf file to ensure that *files* is the first entry specified for each database listed in the file.

1.4.4 Redundancy with distributed NFM-P server and NFM-P database and NFM-P auxiliary collectors

In addition to the rules stated above for distributed NFM-P server and NFM-P database, the following rules apply:

- The operating systems installed on the NFM-P auxiliary collectors must be of the same versions and patch levels as the NFM-P server and NFM-P database workstations.
- If collecting statistics using the NFM-P auxiliary database, the operating systems installed on the NFM-P auxiliary database workstations must be of the same versions and patch levels as the NFM-P server, NFM-P database, and NFM-P auxiliary statistics collector workstations.
- NFM-P auxiliary collectors are intended to be on the same high availability network as the NFM-P server and NFM-P database. NFM-P auxiliary statistics, call trace, and PCMD collectors are intended to be geographically collocated with the active and standby locations of the NFM-P server and NFM-P database. The NSP flow collector typically resides in the managed network, closer to the network elements.
- When using more than one active NFM-P auxiliary statistics collector in a geographic (greater than 1ms latency) configuration, the active and reserved collectors for a give NFM-P server must reside in the same geographic site. The auxiliary statistics collectors in the opposite geographic site would be configured as Remote.
- Workstations running the NFM-P auxiliary collector software must be configured to perform name service database lookups on the local workstation before reverting to a name service database located on the network such as NIS, NIS+, or DNS. A root user must inspect and modify the */etc/nsswitch.conf* file to ensure that *files* is the first entry specified for each database listed in the file.

2 Operating systems and third party software

2.1 Red Hat Enterprise Linux (RHEL)

2.1.1 NFM-P support

NFM-P is supported on Red Hat Enterprise Linux (RHEL) 7, Server Edition x86-64. Previous releases or other variants of Red Hat and other Linux variants are not supported.

NFM-P Release 22.3 supports the following base RHEL versions:

- RHEL server 7 x86-64 - Update 3 (7.3)
- RHEL server 7 x86-64 - Update 4 (7.4)
- RHEL server 7 x86-64 - Update 5 (7.5)
- RHEL server 7 x86-64 - Update 6 (7.6)
- RHEL server 7 x86-64 - Update 7 (7.7)
- RHEL server 7 x86-64 - Update 8 (7.8)
- RHEL server 7 x86-64 - Update 9 (7.9)

See the Host Environment Compatibility Reference for NSP and CLM for compatibility between NFM-P releases and RHEL OS versions.

The Nokia provided RHEL OS qcow2 image is based upon RHEL 7.9 and is available for KVM and OpenStack based deployments only.

The Red Hat Linux support of NFM-P is applicable to specific x86 Intel platforms provided by HP and Nokia only, for bare metal installations, where some systems may require specific updates of the RHEL operating system. See Red Hat's Hardware Certification list on their website. NFM-P does not necessarily support all functionality provided in RHEL 7.

NFM-P supports the use of the RHEL Logical Volume Manager (LVM) on all server types and is limited to the resizing of logical volumes only. To ensure that disk throughput and latency of the resized volume remains consistent, the procedure To test NFM-P disk performance, in the NFM-P Administrator Guide, must be followed.

The RHEL operating system must be installed in 64-bit mode where NFM-P software will be installed.

The NFM-P server, NFM-P auxiliary collector, NSP flow collector, NSP flow collector controller, NFM-P auxiliary database, NSP analytics server, NFM-P client delegate, and NFM-P database RHEL operating system must be installed in English.

Nokia recommends the installation of any OS, driver, or firmware updates that the hardware vendor advises for RHEL.

With the exception of NFM-P documented Operating System parameter changes, all other settings must be left at the RHEL default configuration.

2.1.2 Red Hat support

For customers using the NSP_RHEL7 qcow2 image for NSP guest virtual machines, support for the RHEL instance is available directly from Nokia, not Red Hat. For all other RHEL installations, Red Hat support must be purchased for all platforms running RHEL server with NFM-P. It is strongly recommended to purchase a support package from Red Hat that provides 24x7 support.

The NSP_RHEL7 qcow2 image can only be used as a guest VM hosting an NSP component, and not for the deployment of any other Nokia or third-party product.

2.2 Microsoft Windows

2.2.1 NFM-P support

The Windows operating system is only supported for NFM-P clients and NFM-P client delegate servers. The table below summarizes Microsoft Windows support.

Table 2-1 Windows operating system support summary

Microsoft Windows version	NFM-P client	NFM-P client delegate server
Windows 8 / 8.1 Enterprise	Supported (64-bit)	Not-supported
Windows 10 Professional	Supported	Not-supported
Windows Server 2012 R2	Supported	Supported
Windows Server 2016	Supported	Supported
Windows Server 2019	Supported	Supported

When installing the NFM-P client on Windows, ensure that there is sufficient disk space as identified in the *NSP NFM-P Installation and Upgrade Guide* for the software.

2.2.2 Microsoft support

Support for all installations of the Microsoft Windows operating system must be obtained from Microsoft.

2.3 Apple macOS

2.3.1 NFM-P support

The Mac OS operating system is only supported for NFM-P clients. macOS 11 (Big Sur) has been tested with NFM-P 22.3.

2.3.2 Apple support

Support for all installations of the Apple MacOS operating system must be obtained from Apple.

2.4 Operating system summary

2.4.1 Operating system summary

The following table summarizes the supported configurations for each of the Operating Systems supported by NFM-P.

Table 2-2 NFM-P operating system support summary

NFM-P application	RHEL 7 server x86-64	Microsoft Windows	Mac OS
NFM-P server	7.3 through 7.9	Not supported	Not supported
NFM-P database	7.3 through 7.9	Not-supported	Not supported
Collocated NFM-P server/database	7.3 through 7.9	Not supported	Not supported
NFM-P client	7.3 through 7.9	Supported	Supported
NFM-P auxiliary	7.3 through 7.9	Not supported	Not supported
NFM-P auxiliary database	7.3 through 7.9	Not supported	Not supported
NSP analytics server	7.3 through 7.9	Not supported	Not supported
NSP flow collector	7.3 through 7.9	Not supported	Not supported
NSP flow collector controller	7.3 through 7.9	Not supported	Not supported
NFM-P client delegate	7.3 through 7.9	Supported	Not supported

2.5 Third party software

2.5.1 NFM-P client or client delegate software requirements

NFM-P clients are launched, installed and uninstalled through a web browser using either java webstart or direct installer download. To use the java webstart functionality, each client platform must have a system JRE (Java Runtime Environment) installed along with a supported web browser. The NFM-P java webstart installer/launcher requires Oracle Java version 8, update 192 or later, Oracle Java version 9, or Oracle Java version 10 (18.3) for the system JRE on all Windows, RHEL, and Mac OS platforms. Since Oracle Java version 9 and Oracle Java version 10 are end-of-life, it is strongly recommended to continue using Oracle Java 8 for the system JRE, to ensure access to current security and functional updates. The system JRE needs to be already installed on the client platform. The java webstart installation method will be removed in a future release. The direct installer download does not require the client platform to have a system JRE installed.

The NEtO element manager that is cross launched from the NFM-P client UI requires binding to a specific system port on an NFM-P client and therefore a client delegate can only support a single NEtO instance running amongst all clients connected to a client delegate at any time.

To consolidate NFM-P client UIs to a single server when using the NEtO element manager, a virtualized solution should be used instead, with each NFM-P client residing in a separate VM.

3 Platform requirements

3.1 Hardware platform requirements

3.1.1 Overview

For all bare metal installations, Nokia requires the use of supported HP or Nokia AirFrame Intel based x86 workstations running RHEL.

For optimal disk I/O performance, the read and write caches must be enabled for each disk / volume. Specific HBA controllers may be required for certain platforms to ensure that the read and write caches can be enabled. The server vendor should be consulted to determine the correct HBA controller to allow the creation of the correct number of volumes and enable the read and write caches.

Redundant installations of NFM-P can use different workstations for the active and inactive platforms provided that each of the servers meet the minimum requirements for the intended deployment. In the case where different platforms are used, performance differences are expected depending upon which server is active.

Applications that are not sanctioned by Nokia should not be running on any of the NFM-P server, auxiliary or database workstations. Nokia reserves the right to remove any application from workstations running NFM-P components that are suspected of causing issues.

The hardware platforms do not support running applications that are not specifically identified for that platform. For instance, an NFM-P client is not supported on the hardware platform for a distributed or collocated NFM-P server as there is a significant memory requirement for the NFM-P client that will impact the behavior of the NFM-P server platform.

In exceptional circumstances, a single NFM-P GUI client can be temporarily run from an NFM-P server, but should only be used when remote clients are unavailable.

NFM-P supports the use of the Operating System SNMP agent for monitoring platform availability and system resources. The number of OIDs retrieved must not exceed 100 per minute to ensure NFM-P is not negatively impacted.

3.2 Hardware platform and resource requirements using virtualization

3.2.1 Overview

Virtualization is supported using VMware vSphere ESXi, RHEL KVM, and OpenStack. All other forms of virtualization or virtualization products are not supported.

For installations of the NFM-P server, NFM-P database, NFM-P auxiliary database, NSP flow collector, NSP flow collector controller, and NFM-P auxiliary collector on a Guest Operating System of a virtualized installation, the Guest Operating System must be an NFM-P supported version of RHEL 7 server x86-64. For installations of the NFM-P client and NFM-P client delegate on a Guest Operating System of a virtualized installation, the Guest Operating System can be either an NFM-P supported version of RHEL 7 Server or Windows.

Defined CPU and Memory resources must be reserved and dedicated to the individual Guest OSs and cannot be shared or oversubscribed. As a best practice, VMs should be configured with a single virtual socket with all vCPUs assigned to it, if possible. Additional hardware resources should be reserved for use by the host hypervisor installation to ensure that the resources assigned to the Guest OSs is not impacted. Disk and Network resources should be managed appropriately to ensure that other Guest OSs on the same physical server do not negatively impact the operation of NFM-P.

Virtualized installations of NFM-P are server vendor agnostic but must meet specific hardware criteria and performance targets to be used with NFM-P. Server class hardware must be used, not desktops. Processor support is limited to specific Intel Xeon based x86-64 CPUs with a minimum CPU core speed as outlined in the proceeding tables. The CPU must be from the Haswell microarchitecture, or newer, where the CPU microarchitecture determines the minimum supported CPU speed.

For best performance, storage should be either internal disks (10K or 15K RPM SAS), Fiber Channel attached storage (array or SAN) with dedicated Fiber Channel connections between hosts and Storage Arrays, or 10Gb iSCSI using non-shared infrastructure. All storage must meet the performance metrics provided with NFM-P platform sizing responses. Performance must meet the documented requirements for both throughput and latency.

Nokia support personnel must be provided with the details of the provisioned virtual machine. These details can either be provided through read-only access to the hypervisor or must be available to Nokia support when requested. Failure to provide these details could impact support of NFM-P.

3.2.2 VMware virtualization

NFM-P supports using VMware vSphere ESXi 5.5, 6.0, 6.1, 6.5, 6.7, and 7.0 only, on x86 based servers natively supported by ESXi. VMware's Hardware Compatibility List (HCL) should be consulted to determine specific hardware support. Not all features offered by ESXi are supported when using NFM-P. For example, Memory Compression, or Storage vMotion are not supported. Nokia should be contacted to determine if a specific ESXi feature is supported with an NFM-P installation.

Defined CPU and Memory resources must be reserved for the individual Guest OSs and cannot be shared or oversubscribed. This includes individual vCPUs which must be reserved for the VM.

For new installations, it is recommended to use the latest Virtual Machine Hardware version supported by all of the ESXi hosts in the cluster, from the supported versions. For existing installations, VMware's best practices should be followed regarding Virtual Machine Hardware version changes. Virtual Machine versions 10, 14, and 19 have been tested with NFM-P where the minimum supported Virtual Machine Hardware version is 10 and the latest supported version is 19.

See the following table for additional virtual machine setting requirements.

Table 3-1 VMware virtual machine settings

Resource Type	Parameter	Setting
CPU	Shares	Set to High
	Reservation	Must be set to the number of vCPUs * the CPU frequency. For example, on a 2.4GHz 8 vCPU configuration, the reservation must be set to (8*2400) 19200MHz
	Limit	Unlimited
Memory	Shares	set to High
	Reservation	Reserve all guest memory
	Limit	check box checked for Unlimited
Disk	Shares	set to High
	Limit - IOPs	set to Unlimited
	Type	Thick Provision Eager Zeroed
SCSI controller	Type	VMware Paravirtual
Network Adapter	Type	VMXNET 3

Only one time synchronization mechanism should be used on a guest VM. If using NTP or similar protocol on the guest VM, VMtools time synchronization should be disabled for that guest.

3.2.3 VMware features

The following VMware features have been tested with NFM-P. To ensure NFM-P stability and compatibility, the following recommendations should be noted for each feature:

vMotion

- Always follow VMware best practices
- Testing was performed with dedicated 10Gb connections between all hosts
- Not supported with the NFM-P auxiliary database

High Availability

- Always follow VMware best practices
- Do not use Application Monitoring
- Use Host or VM Monitoring only
- Enable NFM-P database alignment feature to keep active servers in same Data Center

Snapshots

- Always follow VMware best practices
- Do not include memory snapshots
- Always reboot all NFM-P virtual machines after reverting to snapshots
- NFM-P performance can be degraded by as much as 30% when a snapshot exists and therefore NFM-P performance and stability is not guaranteed

- Snapshots should be kept for the least amount of time possible
- Snapshot deletion can take many hours and will pause the VM several times
- NFM-P database failover will occur when VMs are reverted to snapshots, requiring a re-instantiation of the standby database
- Supported on all components except for the NFM-P auxiliary database

Distributed Resource Scheduler (DRS)

- Always follow VMware best practices
- Manual and partially automated DRS is supported
- Fully Automated DRS is not supported
- Supported on all components except for the NFM-P auxiliary database

3.2.4 KVM virtualization

NFM-P supports using RHEL 6.3 through 6.7 KVM using QEMU version 0.12.1.2 and RHEL 7.2 through 7.9 KVM using QEMU version 1.5.3, 2.0.0, 2.3.0, 2.10.0, or 2.12.0 only, on x86 based servers natively supported by KVM. RHEL's Hardware Compatibility List (HCL) should be consulted to determine specific hardware support. Not all features offered by KVM are supported when using NFM-P. For example, Live Migration, Snapshots, or High Availability are not supported. Nokia should be contacted to determine if a specific KVM feature is supported with an NFM-P installation.

Defined CPU and Memory resources must be reserved for the individual Guest OSs and cannot be shared or oversubscribed. This includes individual vCPUs which must be reserved for the VM.

The Disk Controller type must be set to "virtio", the storage format must be configured as "raw", cache mode set to "none", the I/O mode set to "native", and the I/O scheduler set to "deadline". The NIC device model must be "virtio". The hypervisor type must be set to "kvm".

3.2.5 OpenStack

NFM-P tests on open source OpenStack and will support the application running on any OpenStack distribution that is based on the tested versions. Any product issues deemed to be related to the specific OpenStack distribution will need to be pursued by the customer with their OpenStack vendor. Supported OpenStack versions include Newton, Queens, and Train.

To ensure NFM-P stability and compatibility with OpenStack, the following recommendations should be noted:

Hypervisor

- KVM is the only hypervisor supported within an OpenStack environment. See the preceding section for supported versions.

CPU and Memory resources

- Defined CPU and Memory resources must be reserved and dedicated to the individual Guest OSs and cannot be shared or oversubscribed. The OpenStack Nova configuration for cpu_

allocation_ratio and ram_allocation_ratio must both be set to 1.0 on either the control node or each individual compute node where a VM hosting NFM-P could reside.

Hyperthreading

- Hyper-threaded CPU usage must be consistent across all compute nodes. If there are CPUs that do not support hyper-threading, hyper-threading must be disabled on all compute nodes, at the hardware level, where NFM-P components could be deployed.

CPU Pinning

- CPU pinning is supported but not recommended as it restricts the use of OpenStack migration

Availability zones / affinity / placement:

- Nokia does not provide recommendations on configuring OpenStack for VM placement.

Migration

- Only Regular migration is supported. Live migration is not supported.

Networking

- Basic Neutron functionality using Open vSwitch with the ML2 plugin can be used in an NFM-P deployment. OpenStack floating IP address functionality can be used on specific interfaces used by NFM-P that support the use of NAT. This would require a Neutron router using the neutron L3 agent.

Storage

- All storage must meet the performance metrics provided with NFM-P Platform Responses. Performance must meet the documented requirements for both throughput and latency.

VM Storage

- VM storage must be persistent block (Cinder) storage and not ephemeral. For each VM to be deployed, a bootable Cinder volume must be created. The size of the volume is indicated in the NFM-P Platform sizing response.

Flavors

- Flavors should be created for each “Station Type” indicated in the NFM-P platform sizing response.

Firewalls

- Firewalls can be enabled using OpenStack Security Groups or on the VMs using firewalld. If firewalld is used, an OpenStack Security Group that allows all incoming and outgoing traffic should be used.

3.3 Minimum platform requirements

3.3.1 Minimum hardware platform requirements

The following tables specify the minimum hardware platform requirements necessary to successfully operate the NFM-P application in a bare metal configuration.

The minimum platform requirements also represent the smallest configurations suitable for lab evaluations and demonstrations of the NFM-P product.

3.3.2 Bare metal hardware configurations

Table 3-2 NFM-P bare metal minimum collocated platform

For networks not exceeding: • 675 equivalent MDAs • 1,000 GNEs • 5 simultaneous NFM-P clients (GUI or XML-API) • 3,000 elemental STM tests every 15 minutes • 50,000 performance or 100,000 accounting statistics records every 15 minutes • 50,000 TCAs	
NFM-P application	x86 architecture
NFM-P server and database (Collocated)	6* x86 CPU Cores (12 Hyper-threads), minimum 2.0GHz [1] 64 GB RAM recommended (54 GB RAM minimum) 4*10K RPM SAS disk drives of at least 300 GB in size is required for performance and storage capacity

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz.

Table 3-3 NFM-P bare metal minimum distributed platform

For networks not exceeding: • 1,875 equivalent MDAs and 5 simultaneous NFM-P clients (GUI or XML-API) OR • 1,275 equivalent MDAs and 25 simultaneous NFM-P clients (GUI or XML-API) • Maximum of 5,000 GNEs • 6,000 elemental STM tests every 15 minutes • 150,000 performance or 200,000 accounting statistics records every 15 minutes • 150,000 TCAs	
NFM-P application	x86 architecture
NFM-P server	6* x86 CPU Cores (12 Hyper-threads), minimum 2.0GHz [1] 64 GB RAM recommended (55 GB RAM minimum). 2*10K RPM SAS disk drives of at least 300 GB each in size
NFM-P database	4* x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 32 GB RAM minimum. 4*10K RPM SAS disk drives of at least 300 GB in size is required for performance and storage capacity

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz.

3.3.3 Minimum hardware resource requirements

The following tables list the minimum hardware resource requirements for deployments of NFM-P using VMware vSphere ESXi or RHEL KVM.

The minimum platform requirements also represent the smallest configurations suitable for lab evaluations and demonstrations of the NFM-P product.

3.3.4 Virtual machine minimum collocated resource requirements

Table 3-4 NFM-P virtual machine minimum collocated configuration

For networks not exceeding: • 675 equivalent MDAs • 1,000 GNEs • 5 simultaneous NFM-P clients (GUI or XML-API) • 3,000 elemental STM tests every 15 minutes • 50,000 performance or 100,000 accounting statistics records every 15 minutes • 50,000 TCAs	
NFM-P application	VM Guest hardware resource requirements
NFM-P server and database (collocated)	12 vCPUs, minimum 2.0GHz [1] 64 GB RAM recommended (54 GB RAM minimum) 800GB disk space I/O requirements found in 3.5 “Storage” (p. 46)

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz.

   The minimum resource requirements above are also applicable in situations where the NFM-P application is installed in a redundant configuration.

3.3.5 Virtual machine minimum distributed resource requirements

Table 3-5 NFM-P virtual machine minimum distributed configuration

For networks not exceeding: • 1,875 equivalent MDAs and 5 simultaneous NFM-P clients (GUI or XML-API) OR • 1,275 equivalent MDAs and 25 simultaneous NFM-P clients (GUI or XML-API) • Maximum of 5,000 GNEs • 6,000 elemental STM tests every 15 minutes • 150,000 performance or 200,000 accounting statistics records every 15 minutes • 150,000 TCAs	
NFM-P application	VM Guest hardware resource requirements
NFM-P server	12 vCPUs, minimum 2.0GHz [1] 55 GB RAM minimum (64 GB recommended) 500 GB disk space I/O throughput and latency as provided in NFM-P sizing response

Table 3-5 NFM-P virtual machine minimum distributed configuration (continued)

For networks not exceeding: • 1,875 equivalent MDAs and 5 simultaneous NFM-P clients (GUI or XML-API) OR • 1,275 equivalent MDAs and 25 simultaneous NFM-P clients (GUI or XML-API) • Maximum of 5,000 GNEs • 6,000 elemental STM tests every 15 minutes • 150,000 performance or 200,000 accounting statistics records every 15 minutes • 150,000 TCAs	
NFM-P application	VM Guest hardware resource requirements
NFM-P database	8 vCPUs, minimum 2.0GHz [1] 32 GB RAM minimum 1000 GB disk space I/O throughput and latency as provided in NFM-P sizing response

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz.

All of the minimum hardware platforms above are also applicable in situations where the NFM-P application is installed in a redundant configuration.

3.3.6 Scaling limits for collocated configurations

Collocated configurations have been capped at the maximums described in the following table. Higher numbers may be achievable, but Nokia will only support the stated maximums. Note that all stated maximums may not be achievable simultaneously.

Table 3-6 Scaling limits for collocated configurations

Scaling parameter	Maximum
Number of MDAs	1,875
Number of Simultaneous NFM-P clients (GUI or XML-API)	25
Number of SAPs	600,000
Number of OAM tests per 10 minute interval	1,000
Performance statistics per 15 minute interval	50,000
Accounting statistics per 15 minute interval	200,000
TCAs	50,000

3.3.7 Minimum platform requirements for NFM-P auxiliary collectors

Table 3-7 NFM-P auxiliary platforms - Bare Metal

Architecture	Supported NFM-P auxiliary type	Configuration
Bare Metal x86	statistics collector	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 12 GB RAM minimum. 16GB RAM is recommended. 4*10K RPM SAS disk drives of at least 300 GB each in size
Bare Metal x86	call trace collector	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 32 GB RAM minimum - CMM call trace collection 42 GB RAM minimum - CMG or CMM+CMG call trace collection 4*10K RPM SAS disk drives of at least 300 GB each in size - CMM call trace collection 7*10K RPM SAS disk drives of at least 300 GB each in size - CMM or CMM+CMG call trace collection
Bare Metal x86	PCMD collector	12 * x86 CPU Cores (24 Hyper-threads), minimum 2.6GHz 64 GB RAM minimum. 2*10K RPM SAS disk drives of at least 300GB each in size (RAID 1) + 6*0K RPM SAS disk drives of at least 300GB each in size (RAID 0) Minimum of two 1Gb network interfaces. One dedicated to PCMD data collection.

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

Table 3-8 NFM-P auxiliary platforms - VM

Architecture	Supported NFM-P auxiliary type	Configuration
VMware/KVM	statistics collector	8 vCPUs, minimum 2.0GHz [1] 12 GB RAM minimum. 16GB RAM is recommended. 600 GB disk space I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	call trace collector	8 vCPUs, minimum 2.0GHz [1] 24 GB RAM minimum - CMM call trace collection 42 GB RAM minimum - CMG or CMM+CMG call trace collection 800 GB disk space - CMM call trace collection 1,900 GB disk space - CMG or CMM+CMG call trace collection I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	PCMD collector	24 vCPUs, minimum 2.6GHz 64 GB RAM minimum. 2,000 GB disk space Dedicated network interface for PCMD data collection. I/O throughput and latency as provided in NFM-P sizing response

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

3.3.8 Platform requirements for NSP flow collector and NSP flow collector controller

Table 3-9 NSP flow collector and NSP flow collector controller platforms for labs

Architecture	Type	Configuration
Bare Metal x86	NSP flow collector	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 16 GB RAM minimum. 1*10K RPM SAS disk drives of at least 300 GB each in size
Bare Metal x86	NSP flow collector controller	2 * x86 CPU Cores (4 Hyper-threads), minimum 2.0GHz [1] 4 GB RAM minimum. 1*10K RPM SAS disk drives of at least 300 GB each in size
Bare Metal x86	NSP flow collector and NSP flow collector controller (collocated)	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 16 GB RAM minimum. 1*10K RPM SAS disk drives of at least 300 GB each in siz
VMware/KVM	NSP flow collector	8 vCPUs, minimum 2.0GHz [1] 16 GB RAM minimum. 300 GB disk space I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	NSP flow collector controller	4 vCPUs, minimum 2.0GHz [1] 4 GB RAM minimum. 300 GB disk space I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	NSP flow collector and NSP flow collector controller (collocated)	8 vCPUs, minimum 2.0GHz [1] 16 GB RAM minimum. 300 GB disk space I/O throughput and latency as provided in NFM-P sizing response

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

Table 3-10 NSP flow collector and NSP flow collector controller platforms for production deployments

Architecture	Type	Configuration
Bare Metal x86	NSP flow collector	12 * x86 CPU Cores (24 Hyper-threads), minimum 2.0GHz [1] 64 GB RAM minimum. 2*10K RPM SAS disk drives of at least 300 GB each in size

Table 3-10 NSP flow collector and NSP flow collector controller platforms for production deployments (continued)

Architecture	Type	Configuration
Bare Metal x86	NSP flow collector controller	2 * x86 CPU Cores (4 Hyper-threads), minimum 2.0GHz [1] 4 GB RAM minimum. 1*10K RPM SAS disk drives of at least 300 GB each in size
Bare Metal x86	NSP flow collector and NSP flow collector controller (collocated)	12 * x86 CPU Cores (24 Hyper-threads), minimum 2.0GHz [1] 64 GB RAM minimum. 2*10K RPM SAS disk drives of at least 300 GB each in size
VMware/KVM	NSP flow collector	24 vCPUs, minimum 2.0GHz [1] 64 GB RAM minimum. 500 GB disk space I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	NSP flow collector controller	4 vCPUs, minimum 2.0GHz [1] 8 GB RAM minimum. 300 GB disk space I/O throughput and latency as provided in NFM-P sizing response
VMware/KVM	NSP flow collector and NSP flow collector controller (collocated)	24 vCPUs, minimum 2.0GHz [1] 64 GB RAM minimum. 500 GB disk space I/O throughput and latency as provided in NFM-P sizing response

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

3.3.9 Minimum platform requirements for NFM-P auxiliary database

Table 3-11 NFM-P auxiliary database platform - single node cluster

Architecture	Configuration
Bare Metal x86	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.4GHz 64 GB RAM minimum. 2 SAS 10K RPM disk drives of at least 300GB each in size (RAID 1) + 4 SAS 10K RPM disk drives of at least 1.2TB each in size (RAID 1+0) + 4 SAS 10K RPM disks of at least 1.2TB each in size (RAID 5)
VMware/KVM	8 vCPUs, minimum 2.4GHz 64 GB RAM minimum. 500+ GB disk space I/O throughput as measured with the vioperf utility

Table 3-12 NFM-P auxiliary database platform - three+ node cluster

Architecture	Configuration (for each node of the auxiliary database cluster [1])
Bare Metal x86	12 * x86 CPU Cores (24 Hyper-threads), minimum 2.6GHz 128 GB RAM minimum. 2 SAS 10K RPM disk drives of at least 300GB each in size (RAID 1) + 12 SAS 10K RPM disk drives of at least 600GB each in size (RAID 1+0) + 5 SAS 10K RPM disks of at least 1.2TB each in size (RAID 5) Minimum of two 1Gb network interfaces. One dedicated to inter-cluster communication.
VMware/KVM	24 vCPUs, minimum 2.6GHz 128 GB RAM minimum. Dedicated network interface for inter-cluster communication only. 500+ GB disk space I/O throughput as measured with the vioperf utility simultaneously on all nodes in the cluster

Notes:

1. Minimum of three nodes required in the cluster.

3.3.10 Minimum platform requirements for NSP analytics server

Table 3-13 NSP analytics server platform

Architecture	Configuration
Bare Metal x86	4 * x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz [1] 8 GB RAM minimum (24 GB recommended) 1 SAS 10K RPM disk drive of at least 300GB in size
VMware/KVM	8 vCPUs, minimum 2.0GHz [1] 8 GB RAM minimum (24 GB recommended) 300 GB disk space I/O throughput and latency as provided in the NFM-P sizing response

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

Ad-hoc report design is a computationally intensive process. If it is expected that the ad-hoc feature will be used on a regular basis, customers are strongly encouraged to meet the recommended specifications.

3.3.11 Platform requirements for NFM-P client delegate workstations

NFM-P allows multiple clients to be installed on a single HP x86 or Nokia AirFrame workstation running RHEL 7 server x86-64, or specific versions of Windows. This option enables customers to launch multiple NFM-P clients from a single workstation. These GUI clients can be displayed using a Citrix client/Server, or additionally in the case of RHEL, the X11 protocol to other desktops, or native X displays.

The client delegate platform provides an option to consolidate multiple installations of the NFM-P client on a single workstation or the option of installing one instance of the NFM-P client run by many users (with unique Operating System accounts). Regardless of the method of the client installation, the platform requirements per client are the same.

The amount of memory listed includes the requirement for the NFM-P java UI and web browser. Additional memory for each NFM-P client will be required for management of the network elements described in 5.8 “Network Element specific requirements” (p. 66) .

Management of certain network elements may include the cross-launch of additional software that may not be compatible with certain operating systems. The information in Table 3-16, “Element manager operating system support summary” (p. 44) lists these element managers and their operating system support. The documentation of these element managers should be consulted to determine current operating system support.

The NFM-P client delegate configuration is only supported on specific HP or Nokia AirFrame x86 workstations running RHEL server x86-64, or specific versions of Windows. Additionally, the NFM-P client delegate installation is supported on a Guest Operating System of a VMware vSphere ESXi or RHEL KVM installation. The Guest OS must be one of those supported for GUI clients found in 2.1 “Red Hat Enterprise Linux (RHEL)” (p. 25). Table 3-14, “Minimum NFM-P client delegate resource requirements” (p. 41) describes resource requirements for this type of workstation.

Table 3-14 Minimum NFM-P client delegate resource requirements

Architecture	Configuration
Bare Metal x86	4* x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz 28 GB RAM minimum, 36 GB for networks with greater than 15,000 NEs, 76 GB for networks managing AirScale eNodeBs 1*10K RPM SAS disk drive, 146GB in size
VMware/KVM	8 vCPUs, minimum 2.0GHz 28 GB RAM minimum, 36 GB for networks with greater than 15,000 NEs, 76 GB for networks managing AirScale eNodeBs 70 GB disk space

The configurations in the preceding table will support up to 15 GUI clients. Additional GUI clients can be hosted on the same platform provided that the appropriate additional resources found in Table 3-15, “Additional client NFM-P client delegate resource requirements” (p. 42) are added to the platform.

Table 3-15 Additional client NFM-P client delegate resource requirements

Architecture	Additional resources per client
Bare Metal x86	1/4 * x86 CPU Core (1/2 Hyper-thread), minimum 2.0GHz 1.75 GB RAM, 2.25 GB for networks with greater than 15,000 NEs, 5 GB for networks managing AirScale eNodeBs 1 GB Disk Space

Table 3-15 Additional client NFM-P client delegate resource requirements (continued)

Architecture	Additional resources per client
VMware/KVM	1/2 vCPU, minimum 2.0GHz 1.75 GB RAM, 2.25 GB for networks with greater than 15.000 NEs, 5 GB for networks managing AirScale eNodeBs 1 GB Disk Space

For situations where more than 60 simultaneous GUI sessions are required, Nokia recommends deploying multiple NFM-P client delegates.

Displaying GUI clients to computers running X-emulation software is not currently supported. In cases where the GUI client is to be displayed to a PC computer running Windows, Nokia supports installing the GUI client directly on the PC.

NFM-P supports using Citrix for remote display of NFM-P clients. Supporting Citrix on the delegate platform will require extra system resources that will need to be added to those that are required by the NFM-P delegate. Refer to Citrix documentation to determine the additional Citrix resource requirements.

The following Citrix software has been tested with the Windows client delegate:

- Windows Server 2012 R2 — Citrix Server - XenApp Version 7.18
- Windows Server 2016 — Citrix Server - XenApp Version 7.18
- Windows Server 2019 — Citrix Server - XenApp Version 7.18
- Windows Server 2019 — Citrix Client - Receiver Version 4.1.2.0.18020
- Windows 10 — Citrix Client - Receiver Version 4.1.2.0.18020

Due to an incompatibility between 64-bit versions of the Firefox web browser and Citrix Server XenApp, the default web browser must be set to either Google Chrome, if using a version of XenApp older than 7.14.

The client delegate can be published in XenApp by installing the delegate server on your delivery controller and then publishing the <delegate install directory>\nms\bin\nmsclient.bat file as a manually published application.

3.3.12 NFM-P client platform requirements

Nokia recommends a minimum of 1.75 GB of dedicated RAM – regardless of the operating system, for the NFM-P client which includes the java UI and web browser memory requirements. In cases where other applications are running on the same platform as the NFM-P client, it is important to ensure 1.75 GB RAM is available to meet the NFM-P client requirements.

Additional memory for each NFM-P client will be required for management of the network elements described in 5.8 “Network Element specific requirements” (p. 66) .

Management of certain network elements may include the cross-launch of additional software that may not be compatible with certain operating systems. The information in the following table lists these element managers and their operating system support. The documentation of these element managers should be consulted to determine current operating system support.

All platforms used to display NFM-P applications must have a WebGL compatible video card and the corresponding drivers installed.

Table 3-16 Element manager operating system support summary

Element manager	Node type	RHEL 7 server support	Microsoft Windows support
NEtO	9500 MPR / Wavence SM	Not-supported	Supported

The following table provides the minimum requirement for the hardware that will host NFM-P GUI client software. Additional memory and disk resources will be required by the Operating System.

Table 3-17 NFM-P client hardware platform requirements

NFM-P client hardware platform requirements	
RHEL platforms	Microsoft Windows
1 CPU (2 Hyper-threads)@ 2.0 GHz or higher 1.75 GB RAM dedicated to NFM-P client, 2 GB for networks with greater than 15,000 NEs, 5 GB for networks managing AirScale eNodeBs 1 GB available disk space 1280*1024 Display resolution for Java GUI (recommended) 1280*720@72ppi Display resolution for NFM-P applications (minimum) 1920*1080@72ppi Display resolution for NFM-P applications (recommended) Example platform: DL380 Gen10	1 CPU (2 Hyper-threads)@ 2 GHz or higher 1.75 GB RAM dedicated to NFM-P client, 2 GB for networks with greater than 15,000 NEs, 5 GB for networks managing AirScale eNodeBs 1 GB available disk space 1280*1024 Display resolution for Java GUI (recommended) 1280*720@72ppi Display resolution for NFM-P applications (minimum) 1920*1080@72ppi Display resolution for NFM-P applications (recommended)

An NFM-P client installation is supported on a Guest Operating System of a VMware vSphere ESXi or RHEL KVM installation. The Guest OS must be one of those supported for GUI clients found in 2.1 “Red Hat Enterprise Linux (RHEL)” (p. 25).

The following table provides the dedicated NFM-P resource requirements for each Guest OS running under VMware vSphere ESXi or RHEL KVM that will be used to host the NFM-P client GUI. This does not include the specific operating system resource requirements which are in addition to the hardware resources listed below. CPU and Memory resources must be reserved and dedicated to the individual Guest OSs and cannot be shared or oversubscribed. Disk and network resources should be managed appropriately to ensure that other Guest OSs on the same physical server do not negatively impact the operation of the NFM-P GUI client.

Table 3-18 Virtualized NFM-P client resource requirements

Virtual machine resource requirements	
RHEL Guest OS resources	Microsoft Windows Guest OS resources
2 vCPUs @ 2GHz or higher 1.75 GB dedicated RAM, 2 GB for networks with greater than 15,000 NEs 1 GB available disk space 1280*1024 Display resolution for Java GUI (recommended) 1280*720@72ppi Display resolution for NFM-P applications (minimum) 1920*1080@72ppi Display resolution for NFM-P applications (recommended)	2 vCPUs @ 2 GHz or higher 1.75 GB dedicated RAM, 2 GB for networks with greater than 15,000 NEs 1 GB available disk space 1280*1024 Display resolution for Java GUI (recommended) 1280*720@72ppi Display resolution for NFM-P applications (minimum) 1920*1080@72ppi Display resolution for NFM-P applications (recommended)

3.4 Platform requirements for larger networks

3.4.1 Determining platform requirements for larger networks

NFM-P may require larger workstations in order to successfully manage networks that exceed any of the dimensions supported by the minimum hardware platforms. In order to determine workstation requirements to successfully manage larger networks, the following information is required:

- Expected number and types of Network Elements to be managed
- Expected number of MDAs in the network to be managed
- Expected number of services and SAPs in the network to be managed
- Expected number of Dynamic LSPs to be deployed in the network
- Maximum expected number of NFM-P clients (GUI) simultaneously monitoring the network
- Expected number of XML-API applications that will connect as clients to the XML API interface
- Expected number of subscribers, specifically for triple-play network deployments
- Expected statistics collection and retention
- Expected number of STM tests
- Expected number of managed GNEs
- Whether NFM-P redundancy is to be utilized
- Whether NEBS compliance is required
- Whether CPAM is required
- Whether RAID 1 is required

The information above must then be sent to an Nokia representative who can provide the required hardware specifications.

Ensure that any projected growth in the network is taken into account when specifying the expected network dimensioning attributes. For existing NFM-P systems, the user may determine the number of MDAs deployed in the network using the help button on the GUI. It is also possible to determine the number of statistics being handled by the system by looking at the “Statistics Collection” information window. Select the “Tools”, then “Statistics”, then “Server Performance Statistics” menu.

List the “Statistics Collection” objects. From this list window, check the “Scheduled Polling Stats Processed Periodic” and the “Accounting Stats Processed Periodic” columns for the performance and accounting statistics that your system is currently processing within the time interval defined by the collection policy (15 minutes by default).

3.5 Storage

3.5.1 Storage overview

This section provides information on configuring workstations that will host NFM-P software.

Specific partition sizes and configuration procedures are available in the *NSP NFM-P Installation and Upgrade Guide*.

When using the RHEL server OS, ext4 is the required file system for all application specific mount points. No other file systems are supported with NFM-P. OS specific mount points can be either xfs or ext4 as the file system. Windows based clients must use a local file system for client files. Network based files systems, including Samba are not supported.

While Nokia identifies areas of the disk that are not specifically required for NFM-P and are partitionable for customer use, workstation resources are expected to be dedicated for NFM-P. As such, these “Remainder” portions of the disks should only be used for static storage purposes. Consideration should also be made to the expected growth of the network. If the “Remainder” is not to be used, then it should not be created.

For all network sizes, Nokia requires the use of at least four disks on workstations running the NFM-P database. This disk configuration allows for better performance by distributing the database across multiple disks. Customized disk configurations may be required for larger network deployments or where large scale statistics collection is required. Request a formal platform sizing for further details. NAS disk configurations are not supported.

Disk configurations for workstations running the NFM-P database with less than four physical disks greatly limits the NFM-P system performance, managed-network size, and data storage capacity, and is therefore only supported for lab trials.

Refer to 6.3 “Scaling guidelines for statistics collection” (p. 71) for statistics collection recommendations.

In NFM-P upgrade scenarios, previous disk configurations may still be valid.

3.5.2 Using RAID technologies

In bare metal deployments, Nokia requires the use of RAID 0 (striping), unless otherwise specified, provided by a hardware based RAID controller. Software based RAID 0 is not supported. Nokia will provide disk layout and configuration details for customers requiring a Storage Array or layouts not specified in the *NSP NFM-P Installation and Upgrade Guide*. The increased disk I/O performance offered by RAID 0 is required for all NFM-P deployments. The *NSP NFM-P Installation and Upgrade Guide* provides details of these configurations. A RAID 0 stripe size of 512 Kbytes is required for optimal NFM-P disk performance. If a platform does not support a stripe size of 512 Kbytes, choose the next largest stripe size, for example, 256 Kbytes. Specifying a smaller or larger stripe size may result in degraded performance that compromises NFM-P network management.

Nokia supports the use of RAID 1 (Mirroring). Deployments requiring increased resiliency are encouraged to use NFM-P platform redundancy. If RAID 1 is required, a platform providing hardware RAID 1 and that has sufficient number of disk to meet the increased disk requirements must be selected.

To reduce the chance of data loss or application down time, Nokia recommends the use of RAID 1, in a RAID 1+0 configuration.

For specific applications, Nokia supports the use of RAID 5 to increase storage resiliency and maximize available space. The NFM-P auxiliary database backup partition is supported with RAID 5.

Note: Nokia is not responsible for installation, administration or recovery of RAID on an NFM-P platform.

3.5.3 Using SAN storage

Nokia supports the use of SAN storage. SAN connectivity must consist of 4Gb or faster optical connections or 10Gb iSCSI connections. It is recommended that these connections are dedicated connections between the hosts and storage arrays. The SAN must be available to NFM-P without interruption in a low latency environment.

NFM-P platform sizing responses will provide the required performance targets when using NFM-P with a SAN. Note that certain mount points may not be required due to deployment options. Refer to the *NSP NFM-P Installation and Upgrade Guide* for required mount points based upon the type of NFM-P workstations deployed.

Note: Nokia is not responsible for installation, administration or recovery of SANs on an NFM-P platform.

3.5.4 Virtualization I/O requirements

When using NFM-P on a guest operating system of a hosted virtualized installation, specific storage requirements must be met. For optimal performance, storage should be either internal disks (10K or 15K RPM SAS), Fiber Channel attached storage (array or SAN) with dedicated fiber channel connections between hosts and Storage Arrays, or 10Gb iSCSI using non-shared infrastructure. All storage must meet the performance metrics provided with NFM-P platform sizing responses. Storage I/O shares must be set to “High” and IOPs set to “Unlimited” for best performance and low latency.

Refer to Table 3-19, “Minimum collocated configuration throughput and latency” (p. 48) for the minimum required throughput and latency for a collocated NFM-P configuration. Higher scale networks and distributed configurations may require alternate throughput and latency targets that will be provided with the NFM-P platform sizing response that is required for every NFM-P deployment.

NFM-P includes a benchmarking utility to be used for determining the throughput and latency of the storage device to be used with the virtual server hosting NFM-P. The utility is installed with an NFM-P server in the /opt/nsp/nfmp/server/nms/bin/unsupported/IOTest directory and is called NSP_IOTest.pl. If NFM-P has not yet been installed, the utility can be obtained from Nokia or from the NFM-P software package.

Executing the utility with the -h flag will present the user with a help menu, explaining different options and presenting execution examples. Each mount point must be tested and must meet the throughput and latency requirements for the specific deployment. These throughput and latency requirements must be obtained from Nokia as they are specific to each deployment. The throughput and latency targets must be met, irrespective of any other activity on the underlying storage device and the targets must be achievable concurrently. For this reason, it is important to understand the underlying storage configuration to ensure that the output of the benchmarking utility is interpreted correctly. For example, each of the listed targets may be achievable using a single 10K RPM SAS disk but concurrently, the listed targets would not be achievable using the same single 10K RPM SAS disk. The performance of NFM-P would be degraded using this configuration.

Table 3-19 Minimum collocated configuration throughput and latency

Mount point	Read (MB/s)	Write (MB/s)	Latency (ms)
/opt/nsp	37	15	< 1.0
/opt/nsp/os	15	15	< 1.0
/opt/nsp/nfmp/server/xml_output	37	15	< 1.0
/opt/nsp/nfmp/dbbackup	14	21	< 1.0
/opt/nsp/nfmp/db/tablespace	158	8	< 1.0
/opt/nsp/nfmp/server/nms/log	1	1	< 1.0
/opt/nsp/nfmp/db/archivelog	14	38	< 1.0
/opt/nsp/nfmp/nebackup	6	6	< 1.0

The *NSP NFM-P Installation and Upgrade Guide* should be consulted for recommended partition sizes.

4 Maintaining current state of network elements

4.1 Mechanism to maintain current state of network elements

4.1.1 Overview

NFM-P uses several mechanisms to maintain and display the current state of the network elements it manages. These mechanisms can include:

- IP connectivity (ping) verification
- SNMP connectivity verification
- SNMP traps
- SNMP trap sequence verification
- Scheduled SNMP MIB polling

These mechanisms are built into the Nokia 7950, 7750, 7450, 7710, 7210, and 7705 Network Elements and the NFM-P network element interaction layers.

4.1.2 IP connectivity (ping) verification

NFM-P can be configured to ping all network elements at a configurable interval to monitor IP connectivity. If the network element is unreachable, an alarm will be raised against the network element. Details of the alarm are the following:

- Severity: Critical
- Type: communicationsAlarm
- Name: StandbyCPMManagementConnectionDown, OutOfBandManagementConnectionDown or InBandManagementConnectionDown
- Cause: managementConnectionDown.

Ping verification is disabled by default. IP connectivity checks using ping must be scheduled through the default policy.

4.1.3 SNMP connectivity verification

NFM-P performs an SNMP communication check every 4 minutes. If NFM-P can not communicate via SNMP with a network element, it will raise a communications alarm against that network element. NFM-P will also color the network element red on the map to indicate the communication problem. NFM-P will clear the alarm and color the network element as green once NFM-P detects SNMP connectivity to the network is re-established. Details of the alarm are the following:

- Severity: Major
- Type: communicationsAlarm
- Name: SnmpReachabilityProblem
- Cause: SnmpReachabilityTestFailed

This behavior occurs by default and is not configurable.

4.1.4 SNMP traps

NFM-P listens to SNMP traps to receive changes from the network elements. NFM-P configures the trap log ID on each network element when it is first discovered. The network element then uses that trap log ID to send all configuration changes and updates to NFM-P. The NFM-P will react to the traps it receives and make appropriate changes to the database, alarms and related object as required.

4.1.5 SNMP trap sequence verification

NFM-P retrieves the last trap sequence number sent from all network elements at a configurable interval. This interval is configurable on a per resource group basis. Resource groups allow the user to configure the communications behavior of a group of network elements. By default, the core resource group includes all network elements, and verifies the trap sequence number every 4 minutes. NFM-P compares that sequence number with the sequence number of the last trap it received from that network element. If they do not match, NFM-P will request only the missing traps from the network element. If at any point NFM-P realizes that it is missing more than 200 traps from a network element, or if the network element no longer has the missed trap, NFM-P will request a full resynchronization on that network element rather than just request the missing traps.

This behavior occurs by default and is not configurable.

4.1.6 Scheduled SNMP MIB polling

NFM-P can poll all data SNMP MIBs from the network elements at a configurable interval. The Polling Policy is disabled by default. This behavior is configurable via the Polling tab of the Network Elements properties form.

4.2 Network outages

4.2.1 Network outages and recovery

When a Nokia 7x50 based network element loses visibility of the NFM-P, it is unable to send traps to the network manager, and the traps are queued on the network element. 4.1.5 “SNMP trap sequence verification” (p. 50) describes NFM-P behavior with regards to trap handling. When a network outage occurs, the network element configuration in NFM-P will be made consistent with the network element, but any event notifications, such as SNMP traps, that occurred during the network outage will not have been processed. This will cause intermediate state change alarms to not be reflected in NFM-P during the network outage.

5 Networking

5.1 Network requirements

5.1.1 Overview

The network interconnecting the NFM-P systems, network elements, and XML-API systems is of significant importance to the effective management of the network. The following sections describe the requirements for the network links between NFM-P workstations and the connection to the network being managed. Nokia recommends that sufficient bandwidth be made available to the NFM-P workstations within the Data Communication Network.

For SNMP management of Nokia network elements, all network segments that carry NFM-P management traffic must allow the successful transmission of 9216 byte SNMP packets. The *NSP NFM-P Troubleshooting Guide* contains more information on packet fragmentation issues.

Be sure to include the tables with the bandwidth required for statistics collection in the total bandwidth required between the NFM-P components, as they are in separate tables.

The tables do not specify the underlying infrastructure required to support these bandwidth requirements.

See Chapter 8, “Multiple network interface NFM-P deployments” for information on configuring the NFM-P components with multiple interfaces.

5.2 Network elements

5.2.1 Network element connectivity support

NFM-P supports both IPv4 and IPv6 connectivity to network elements. The following network elements may be managed by NFM-P using IPv6:

- 7950
- 7750
- 7450
- 7710
- 7705
- 7705
- 7250
- 7210
- vCPAA
- eNodeB: AirScale and FlexiZone
- 5G Classic
- 5G Cloud CU and 5G DU

- 9500 MPR / Wavence SM
- OmniSwitch 6350, 6465, 6560, 6865
- CMM

Call trace data can only be retrieved from CMM and CMG network elements with IPv4 connectivity.

NFM-P supports the use of multiple interfaces for network element management communication. If a network element uses both an in-band and out-of-band address for management, these interfaces must reside on the same server interface.

5.3 Bandwidth requirements

5.3.1 Bandwidth requirements for collocated NFM-P installations

The following table lists the bandwidth requirements for the connections between the components of an NFM-P collocated installation. It is a good practice to measure the bandwidth utilization between the various components to determine a suitable bandwidth. There are a number of factors that could require an increase above our bandwidth utilization recommendations, including: GUI activity, XML-API activity, network events, number of network elements being managed.

Table 5-1 NFM-P collocated server/database bandwidth requirements

Available bandwidth required from primary NFM-P server/database workstation	Recommended bandwidth: excluding statistics bandwidth requirements
NFM-P client (GUI)	1 Mbps
XML API client (The bandwidth will depend on the XML-API application)	1 Mbps
Between primary and standby NFM-P server/database workstation NOTE: When network element database backup synchronization is enabled, the bandwidth requirement between the NFM-P servers will vary significantly depending on the size of the network element backup file sizes.	5-10 Mbps (sustained) 16-26 Mbps (during re-instantiation or database backup synchronization)

5.3.2 Bandwidth requirements for distributed NFM-P installations

The following tables list the requirements for the connections between the components of an NFM-P distributed installation. It is a good practice to measure the bandwidth utilization between the various components to determine a suitable bandwidth. There are a number of factors that could require an increase above our bandwidth utilization recommendations – including: GUI activity, XML-API activity, network events, number of network elements being managed.

Table 5-2 NFM-P distributed server/database bandwidth requirements

Available bandwidth requirements for NFM-P	Recommended bandwidth: excluding statistics and call trace bandwidth requirements
NFM-P server to an NFM-P database NOTE: This depends on GUI changes and lists, # of changes occurring in the network, and network objects managed.	5 to 10 Mbps (3 Mbps minimum)
NFM-P server to an NFM-P client	1 Mbps
NFM-P server to an XML API client (The bandwidth will depend on the XML-API application)	1 Mbps
Between a primary and a standby NFM-P server NOTE: When network element database backup synchronization is enabled, the bandwidth requirement between the NFM-P servers will vary significantly depending on the size of the network element backup file sizes.	1 Mbps
NFM-P server to an NFM-P auxiliary statistics collector	1 Mbps
Between primary and standby NFM-P databases NOTE: The higher bandwidth is required to handle re-instantiation and is also required immediately after a database backup when database backup synchronization is enabled.	6 Mbps (sustained) 15-25 Mbps (during re-instantiation or database backup synchronization) 3 Mbps (minimum)

Table 5-3 Additional bandwidth requirements for file accounting STM results collection

Bandwidth requirements for installations collecting file accounting STM results using the logToFile method only	Increased bandwidth per 50,000 file accounting STM records
NFM-P server to an XML API client if using registerLogToFile NOTE: a higher bandwidth may be desirable	3.5 Mbps
NFM-P server to NFM-P database workstation	1.5 Mbps
Between the NFM-P database workstations – required for sufficient bandwidth for database re-instantiations NOTE: The higher bandwidth is required to handle re-instantiation during STM collection	2 Mbps (sustained) 12 Mbps (during re-instantiation or database backup synchronization)

5.3.3 Additional bandwidth requirements for statistics collection

The size of the network and the number of statistics that are collected will impact the recommended bandwidth. The following tables should be used to determine how much additional bandwidth will be required between the NFM-P workstations when statistics collection is added to the system. The collecting server, in the tables below, would be either the NFM-P auxiliary statistics collector or, in the absence of the NFM-P auxiliary statistics collector, the NFM-P server. The additional bandwidth requirements are per 200,000 collected records per interval. The bandwidths of connections not listed do not change dramatically with the addition of statistics.

The registerLogToFile method of retrieving statistics can be compressed or uncompressed. Using the compressed option will require additional CPU requirements on the workstation that is collecting the statistics (either NFM-P server or NFM-P auxiliary statistics collector). In this case, the bandwidth required will be reduced.

Table 5-4 Additional bandwidth requirements for accounting statistics collection

Record Storage Location	Bandwidth between collecting server and NFM-P database	Bandwidth between collecting server and NFM-P auxiliary database	Bandwidth between NFM-P databases	Bandwidth between NFM-P auxiliary database clusters	Bandwidth between NFM-P databases during re-instantiation	Bandwidth between collecting server and XML-API client
NFM-P database	2.2 Mbps	N/A	3.2 Mbps	N/A	18 Mbps	N/A
NFM-P auxiliary database	N/A	2.2 Mbps	N/A	0.8 Mbps per NFM-P auxiliary database node	N/A	N/A
logToFile (NFM-P server or NFM-P auxiliary statistics collector)	N/A	N/A	N/A	N/A	N/A	3.5 Mbps
findToFile (NFM-P server)	N/A	N/A	N/A	N/A	N/A	3.5 Mbps

Table 5-5 Additional bandwidth requirements for application assurance accounting statistics collection

Record Storage Location	Bandwidth between collecting server and NFM-P database	Bandwidth between collecting server and NFM-P auxiliary database	Bandwidth between NFM-P databases	Bandwidth between NFM-P auxiliary database clusters	Bandwidth between NFM-P databases during re-instantiation	Bandwidth between collecting server and XML-API client
NFM-P database	3.1 Mbps	N/A	4.2 Mbps	N/A	20 Mbps	N/A
NFM-P auxiliary database	N/A	3.1 Mbps	N/A	0.8 Mbps per NFM-P auxiliary database node	N/A	N/A
logToFile (NFM-P server or NFM-P auxiliary statistics collector)	N/A	N/A	N/A	N/A	N/A	4.6 Mbps

Table 5-6 Additional bandwidth requirements for performance statistics collection

Record Storage Location	Bandwidth between collecting server and NFM-P database	Bandwidth between collecting server and NFM-P auxiliary database	Bandwidth between NFM-P databases	Bandwidth between NFM-P auxiliary database clusters	Bandwidth between NFM-P databases during re-instantiation	Bandwidth between collecting server and XML-API client
NFM-P database	5.4 Mbps	N/A	14.4 Mbps	N/A	72 Mbps	N/A
NFM-P auxiliary database	5.4 Mbps	5.4 Mbps	14.4 Mbps	0.8 Mbps per NFM-P auxiliary database node	72 Mbps	N/A
logToFile (NFM-P server or NFM-P auxiliary statistics collector)	N/A	N/A	N/A	N/A	N/A	3.5 Mbps
findToFile (NFM-P server)	N/A	N/A	N/A	N/A	N/A	3.5 Mbps

Table 5-7 Additional bandwidth requirements for LTE performance management statistics collection

Bandwidth requirements for installations collecting LTE performance management statistics	Increased bandwidth per 200,000 LTE performance management statistics records
Between a primary and a standby NFM-P server: If the NFM-P server is collecting the statistics:(NFM-P auxiliary statistics collector is NOT installed)	1.0 Mbps
Between a preferred and a reserved NFM-P auxiliary statistics collector if the NFM-P auxiliary statistics collector is collecting the statistics	1.0 Mbps

When an NFM-P auxiliary statistics collector is installed to collect statistics using the NFM-P database, the bandwidth requirements between two geographic locations will need to reflect the state where an NFM-P auxiliary statistics collector in geographic location A may send information to the active NFM-P server in geographic location B which will - in turn – send information back to the NFM-P database in geographic location A. For this reason, the bandwidth between geographic location A and B must be the sum of the bandwidth requirements between the NFM-P auxiliary statistics collector to NFM-P server and NFM-P server to NFM-P database. It is also a best practice to ensure that the NFM-P auxiliary statistics collector, NFM-P server, and NFM-P database are all collocated in the same geographic site.

5.3.4 NSP analytics server

When an NSP analytics server is deployed with NFM-P, the following bandwidth requirements should be noted. The connections between the NSP analytics server and the NFM-P auxiliary database(s) and NFM-P database(s) require minimal bandwidth.

Table 5-8 Additional bandwidth requirements for the NSP analytics server

Bandwidth requirements for installations with NSP analytics	Recommended Bandwidth
Between the NSP analytics server and the end user (web browser)	2.5 Mbps minimum requirement, 10 Mbps optimal
Between the NSP analytics server and the server hosting nspOS	3 Mbps
Between the NSP analytics server and external FTP host (if used for sending results of scheduled reports)	3 Mbps

5.3.5 NFM-P auxiliary call trace collectors

When an NFM-P auxiliary call trace collector is installed, there are a number of bandwidth requirements listed below. Any bandwidths not listed are not impacted significantly by call trace data collection.

To handle the redundant pairs appropriately, the bandwidth requirements between two geographic locations will need to reflect the state where an NFM-P auxiliary call trace collector in geographic location A may need to provide information to the API client in geographic location B. The synchronization of call trace and debug trace files will be impacted by the number client application ftp sessions retrieving call trace and debug trace files. To minimize this impact, it is recommended to limit the number of ftp sessions.

Table 5-9 Additional bandwidth requirements for call trace collection

Bandwidth requirements for installations with call trace collection	Bandwidth usage characterization
NFM-P server to an XML API client	Low bandwidth XML-API requests and responses
XML API client to NFM-P auxiliary call trace collector workstation NOTE: a higher bandwidth may be desirable	Higher bandwidth to retrieve via FTP the call trace files from the NFM-P auxiliary
NFM-P auxiliary call trace collector Preferred workstation to it's Reserved redundant pair. NOTE: a higher bandwidth may be desirable	Higher bandwidth to ensure timely synchronization of call trace files

5.3.6 NFM-P auxiliary database

When an NFM-P auxiliary database is part of an NFM-P deployment, there are a number of bandwidth requirements listed below. Any bandwidths not listed are not impacted significantly by the use of the NFM-P auxiliary database for statistics collection.

When the auxiliary database cluster is deployed with a minimum of three nodes, the NFM-P auxiliary database server requires a minimum of two network interfaces; one for communication to the NFM-P management complex and one for internal data communication between each of the NFM-P auxiliary database servers in the cluster. The interface for internal data communication needs to be dedicated with a minimum interface speed of 1Gbps and part of a private network.

Table 5-10 Additional bandwidth requirements for NFM-P auxiliary database

Bandwidth requirements for installations with NFM-P auxiliary database	Bandwidth usage characterization
NFM-P auxiliary statistics collector to NFM-P auxiliary database cluster	Higher bandwidth to write statistics data into the NFM-P auxiliary database cluster
NSP analytics server to NFM-P auxiliary database cluster NOTE: a higher bandwidth may be desirable	Higher bandwidth to generate reports based upon raw and aggregated data
NFM-P auxiliary database node to NFM-P auxiliary database node (intra-cluster)	High — must use a dedicated, minimum 1Gbps interface
NFM-P auxiliary database to NFM-P auxiliary database (redundant cluster)	High — up to 500Mb

5.4 Contributors to Bandwith Requirements

5.4.1 NFM-P GUI clients

The bandwidth specifications provided above for NFM-P GUI clients are based on the fact that information about changes in the network is forwarded to the NFM-P GUI clients. The NFM-P client updates information visible to the user based on recent changes in the network.

A few examples of network changes which will be reported to NFM-P include status changes of physical equipment, status changes of Layer 2 or Layer 3 interfaces, configuration of network elements, provisioning of new equipment or services, status changes in services or any attributes thereof, configuration changes of routing protocols and several others.

In situations where the frequency of changes sent to the NFM-P GUI is significant and exceeds the bandwidth specification, the performance of the NFM-P client will degrade, and there is a possibility that the connection to the server will be dropped. An NFM-P GUI restart will be required to reconnect to the server to receive change notifications.

5.4.2 NFM-P GUI clients on X displays

NFM-P GUI clients can be displayed remotely on terminals using the X11 protocol for graphical displays. In these cases, it is important to ensure the bandwidth availability between the workstation running the NFM-P client and the host displaying the NFM-P client be at least 1024 Kbps. Also, it is important to ensure the round-trip network latency between these two hosts is quite low (20-30ms). To achieve acceptable performance on bandwidth limited links, X-compression should be used by using the ssh -XC command. If not using compression, it is recommended that the minimum bandwidth be higher than 1024 Kbps. Situations where the available bandwidth is lower or the network latency is higher will result in poor usability of the NFM-P GUI client. A bandwidth of 1024 Kbps will impact GUI start time and will not meet the published time of less than 2 minutes.

Extra bandwidth may be required to support the network elements described in 5.8 “Network Element specific requirements” (p. 66)

Note that NFM-P GUI client startup may be impacted when using minimum bandwidth links.

5.4.3 XML API clients

There are two main factors affecting the bandwidth requirements between the NFM-P server and an XML API client:

- Design and behavior of the application using the XML-API interface
- Rate of changes in the network

Applications which listen to network changes via the JMS interface provided by NFM-P XML API or applications which retrieve large pieces of information via the API, such as statistics information or network inventory information, will require access to dedicated bandwidth from the machine hosting the application to the NFM-P server according to the tables above. Applications which do not require real time event and alarm notification may operate with acceptable performance when the bandwidth between the machine hosting the application and the NFM-P server is less than the quantity specified in the tables above.

It is a best practice to minimize event and alarm notifications using a JMS filter to reduce bandwidth requirements and the possible effects of network latency.

In an environment where network changes are infrequent, it is possible to successfully operate an application using the API when the bandwidth between the machine hosting this application and the NFM-P server is less than the quantity specified in the tables above, possibly as little as 128 kbps. However, in situations where the frequency of network changes increases, the performance or responsiveness of the application will degrade.

5.4.4 NFM-P auxiliary statistics collector

The main factors impacting communication to and from the NFM-P auxiliary statistics collector are:

- Number of performance statistics being collected. The NFM-P server needs to tell the NFM-P auxiliary statistics collector which statistics to collect every interval.
- Number of statistics collected from the network elements.
- Number of statistics written to the NFM-P database.

The more performance statistics are collected, the more significant the bandwidth utilization between the NFM-P server and the NFM-P auxiliary statistics collector. Similarly, this will require more significant bandwidth utilization between the NFM-P auxiliary statistics collector and the NFM-P database workstations. The bandwidth requirements are not dependent on network activity.

5.4.5 NFM-P call trace collector

The main factors impacting communication to and from the NFM-P auxiliary call trace collector are:

- Number of NEs where call traces are enabled.
- Size of files being retrieved by the NFM-P XML-API client requesting the call trace.

The more call traces that are enabled, the higher the bandwidth requirement from the CMM and CMG network elements to the NFM-P auxiliary call trace collector. Enable and Disable messages are sent to the NFM-P auxiliary call trace collector from the NFM-P server. NFM-P XML-API clients can ask the NFM-P server for the list of NFM-P call trace collector workstations, and ftp connect directly to the NFM-P auxiliary call trace collector to retrieve the call trace log files.

5.4.6 NSP flow collector and flow collector controller

The main factors impacting communication to and from the NSP flow collector are:

- Size of the NFM-P managed network for the network extraction
- Size of generated IPDR files
- Number of network elements sending cflowd records

The main factors impacting communication to and from the NSP flow collector controller are:

- Size of the NFM-P managed network for the network extraction
- Number of NSP flow collectors connected to the NSP flow collector controller

Table 5-11 Additional bandwidth requirements for the NSP flow collector

Bandwidth requirements for NSP flow collector	Bandwidth usage characterization
NSP flow collector controller to an NSP flow collector This is for Network Snapshot Transfer (FTP/SFTP) By default this operation should only occur weekly if the NFM-P server and NSP flow collector controller remain in sync. The amount of bandwidth required is dependent on network size.	Bandwidth requirement will depend upon network size, which determines the network extraction file size, and the desired time complete the file transfer from the NSP flow collector controller to the NSP flow collector
Managed Network to NSP flow collector In the case of Redundant NSP flow collectors, the amount of dedicated bandwidth is required for each NSP flow collector.	40 Mbps per 20,000 flows per second
NSP flow collector to IPDR file storage server Approximate amount of Stats per a 1 MB IPDR Stats File: 2,560 TCP PERF statistics (all counters) or, 3,174 RTP statistics (all counters) or, 9,318 Comprehensive statistics (all counters) or 9,830 Volume statistics (all counters) In the case of Redundant NSP flow collectors, the amount of dedicated bandwidth calculated on the right is for each NSP flow collector to the workstation where IPDR files are being transferred.	Use the information on the left to calculate the amount of data generated for the expected statistics. Use this to calculate the time to transfer at a given bandwidth. The total time must be less than 50% of collection interval. For example – if 1GB of IPDR files are expected per interval, and the collection interval is 5min, a 45 Mbps connection will take 3min,2sec to transfer. This is more than 50% and a larger network connection is required.

Table 5-12 Additional bandwidth requirements for the NSP flow collector controller

Bandwidth requirements for NSP flow collector controller	Bandwidth usage characterization
NFM-P server to an NSP flow collector controller This is for Network Snapshot Transfer (FTP/SFTP) By default this operation should only occur weekly if the NFM-P server and NSP flow collector remain in sync. The amount of bandwidth required is dependent on network size.	Bandwidth requirement will depend upon network size, which determines the network extraction file size, and the desired time complete the file transfer from the NFM-P server to the NSP flow collector controller.

5.4.7 NFM-P PCMD auxiliary collector

The main factors impacting communication to and from the NFM-P auxiliary PCMD auxiliary collector are:

- Number of bearers.
- Number of and size of events per bearer.

- Size of files being retrieved by the NFM-P XML-API client requesting the PCMD files.

On average, each bearer will generate 100 events per hour where each event is approximately 250 bytes in size.

5.5 Network bandwidth

5.5.1 Bandwidth requirements

In order to effectively manage the network, NFM-P must have access to sufficient bandwidth between the NFM-P server(s), NFM-P auxiliary(s) and the network elements.

This bandwidth will be used to carry the management traffic between NFM-P and the network element. The following table describes the bandwidth requirements for a particular network element.

Table 5-13 NFM-P server to network bandwidth requirements

Network element example	Bandwidth requirement from NFM-P server(s) to the network element
7950 XRS	2-4 Mbps
7750 SR-12E (fully loaded)	2 Mbps
7750 SR-12 (fully loaded)	2 Mbps
7750 SR-2s	2 Mbps
7750 SR-a4	1 Mbps
7750 SR-c12 (fully loaded)	600 kbps
7450 ESS-7 (fully loaded)	1 Mbps
7450 ESS-1	200 kbps
7705 SAR (fully loaded)	200 kbps – 400 kbps
7250 IXR-6 / 7250 IXR-R4 / 7250 IXR-R6 / 7250 IXR-x	800 kbps – 1,000 kbps
7250 IXR-e	300 kbps
7210 SAS-E, 7210 SAS-M, 7210 SAS-K	200-300 kbps
7210 SAS-D, 7210 SAS-X, 7210 SAS-T, 7210 SAS-R, 7210 SAS-Mxp, 7210 SAS-Sx	500-600 kbps
7701 CPAA / vCPAA	250 kbps
9500 MPR / Wavence SM	200 kbps
OmniSwitch 6250, 6350 6400, 6450, 6465, 6560, 6850, 6855, 6865, 9000 Series	300 kbps
OmniSwitch 6860, 6860E, 6860N, 6900, 10K	400 kbps
CMM	200 kbps
1830 VWM OSU	400 kbps

5.5.2 Details on the bandwidth requirements

The recommended bandwidth described above is a conservative figure that is meant to ensure that the performance of NFM-P and its ability to manage successfully each network element will not be affected by unusual network conditions.

Specifically, the bandwidth recommendation ensures that NFM-P can fully discover (or resynchronize) all of the objects contained in the network element, within a reasonable amount of time, varying heavily based upon the specific network element type and configuration.

The following are the main operations that result in significant amounts of information being exchanged between NFM-P and the network elements. These factors are therefore the principal contributors to the bandwidth requirements.

- Network Element Discovery: Upon first discovery of the network element, a significant amount of data is exchanged between NFM-P and the network element.
- SNMP traps: SNMP traps do not result directly in significant data being sent from the network element to the NFM-P. Several of the SNMP traps however do not contain all of the information required for NFM-P to completely represent the new status of the network element. As a result, NFM-P will subsequently perform a poll of a certain number of the SNMP MIBs to obtain the required information from the network element. Consequently, SNMP traps do result in a certain quantity of data and therefore cause bandwidth utilization. The exact quantity of bandwidth utilized will vary based on the number and the type of trap that is sent from the network element. In the worst case however, this bandwidth utilization will be less than that utilized during a network element discovery.
- SNMP polling: It is possible to configure NFM-P to poll the SNMP MIBs on the network elements at various intervals. By default, NFM-P will perform a complete poll of the SNMP MIBs every 24 hours on non-SR–OS based network elements. During the polling cycle, the amount of data transferred between NFM-P and the network element is equivalent to the amount of data transferred during the network element discovery.
- Statistics collection: It is possible to configure NFM-P to poll the SNMP MIBs on the network elements that contain performance statistics information. During the polling cycle, the amount of data transferred between NFM-P and the network element is less than the amount of data transferred during the network element discovery. With the configuration of an NFM-P auxiliary statistics collector, the communication from and to the network elements will be distributed between the NFM-P server and an NFM-P auxiliary statistics collector.
- Network element backup: It is possible to configure NFM-P to request a backup of the network element at specified interval. During the NE backup cycle, the amount of data transferred between NFM-P and the network element is less than half of the amount of data transferred during the network element discovery.
- Provisioning of services and deployment of configuration changes: When network elements are configured or when services are provisioned via the NFM-P GUI or via application using the API, a small quantity of network bandwidth is utilized. The amount of data transferred is significantly less than during the network element discovery.
- Initiation and collection of STM tests and their results: When STM tests are initiated, the NFM-P server sends individual requests per elemental test to the network elements. Once the test is complete, the network elements report back using a trap. The NFM-P server then requests the information from the network element, and stores it in the database. This can result in a significant increase in network traffic to the network elements.

- Software Downloads: The infrequent downloading of network element software loads is not included in the bandwidth levels stated in Table 5-13, “NFM-P server to network bandwidth requirements” (p. 60). Bandwidth requirements will depend upon the size of the network element software load and the desired amount of time to successfully transfer the file to the NE.

For some network elements, management of the NE includes methods other than standard MIB/SNMP management – for example web-based tools. These network elements may require additional bandwidth above the bandwidth levels stated in Table 5-13, “NFM-P server to network bandwidth requirements” (p. 60).

5.5.3 Possible consequences of insufficient bandwidth

In situations where there is less than the recommended bandwidth between the NFM-P and the network element, the following are possible consequences:

- The length of time required to perform a network element discovery will increase
- The length of time required to perform a SNMP poll of the network element will increase
- The length of time required to retrieve statistics from the network element will increase
- The proportion of SNMP traps that will not reach NFM-P because of congestion will increase. This is significant since NFM-P will detect it has missed traps from the network element and will result in NFM-P performing additional SNMP polling to retrieve the missing information. This will result in additional data being transferred, which will increase the bandwidth requirements, possibly exacerbating the situation.

5.5.4 Determining total bandwidth requirements for NFM-P managed networks

The amount of bandwidth required for each of the network elements should be obtained from Table 5-13, “NFM-P server to network bandwidth requirements” (p. 60).

The total amount of bandwidth that is required for NFM-P to manage the complete network will vary based on the topology of the infrastructure that is used to carry the management traffic. From NFM-P’s perspective, there must be sufficient bandwidth (as per Table 5-13, “NFM-P server to network bandwidth requirements” (p. 60)) between itself and each of the network elements that is under management.

In cases where the management traffic is carried over physical point-to-point links between the NFM-P server and NFM-P auxiliary network and each of the network elements, sufficient bandwidth must be reserved on the physical links. The NFM-P server complex can simultaneously communicate to several NEs for the following functions:

- NE discovery, NE resync, resyncing for trap processing
- NE backups, NE software downloading, and sending configurations to NEs
- Collecting performance statistics
- Collecting accounting statistics
- Initiating STM tests on NEs
- Retrieve STM Test Results - also via (s)FTP
- NE reachability checks and NE trap gap checks

Rarely are all of the above performed simultaneously so it is recommended to assume for link aggregation points that NFM-P can communicate with a minimum of 20-30 NEs simultaneously – this can increase to 60-70 NEs on a 16 CPU core NFM-P server workstation. For Networks of over 1,000 NEs or where an NFM-P auxiliary statistics collector is being used, that number should be increased by 20-30 NEs. Higher bandwidth maybe required under special cases where above average data is attempted to be transferred between NFM-P and the network elements. For example, large statistics files, NE backups, or software images.

5.6 Network latency

5.6.1 Network latency considerations

Network latency can potentially impact the performance of NFM-P. The following are known impacts of latency between the various NFM-P components:

- NFM-P server to NFM-P clients (GUI/XML-API): event notification rates of network changes
- NFM-P auxiliary statistics collector to the network elements: ftp connection for statistics collection and SNMP stats collection
- NFM-P server to the network elements: resync times, provisioning, ftp connections for statistics and network element backups, trap handling, and SNMP stats collection (See 6.3 “Scaling guidelines for statistics collection” (p. 71) for more information on latency impact on SNMP stats collection)
- NFM-P server and NFM-P auxiliary collectors to NFM-P database: NFM-P performance is sensitive to latency in this area. The round trip latency between the active NFM-P components (server, database, auxiliary) must be no longer than 1 ms., otherwise overall NFM-P performance will be significantly impacted. The NFM-P auxiliary database can tolerate up to 200 ms of latency between it and the rest of the NFM-P management complex.

Since SNMP communication to a single Network Element is synchronous, the impact of latency is directly related to the number of SNMP gets and responses. Operations to a Network Element with a round trip latency of 50 ms will have the network transmission time increase by ten times compared to a Network Element with a round trip latency of only 5 ms. For example, is a specific operation required NFM-P to send 1,000 SNMP gets to a single network element, NFM-P will spend a total of 5 seconds sending and receiving packets when the round trip latency to the network element if 5 ms. The time that NFM-P spends sending and receiving the same packets would increase to 50 seconds if the round trip latency were increased to 50 ms.

Network Element re-sync can be especially sensitive to latency as the number of packets exchanged can number in the hundreds of thousands. For example, if a re-sync consists of the exchange of 100,000 packets (50,000 gets and 50,000 replies), 50 ms of round trip latency would add almost 42 minutes to the overall re-sync time and 100 ms of round trip latency would add almost 84 minutes to the overall re-sync time.

NFM-P can use a proprietary mechanism to discover and resync specific node types and versions, that can dramatically reduce resync and discovery times to network elements with high network latency. TCP Streaming is supported on the following Network Element types with a release of 11.0R5 or later:

- 7950 XRS

- 7750 SR
- 7450 ESS
- 7250 IXR
- 7710 SPR

5.6.2 Geographical redundancy of NFM-P components

It is ideal to ensure that all NFM-P workstations and the NFM-P XML-API clients are collocated within a geographical site on a high availability network to avoid the impact of network latency.

In cases where geographic redundancy is configured, all active NFM-P workstations (NFM-P server, NFM-P auxiliaries, and NFM-P database) should be located within a geographical site on a high availability network to avoid the impact of network latency between components, which must remain at less than 1ms. When an NFM-P component (server, auxiliary, or database) switchover or failover occurs, manual intervention may be required to align the workstations on the same geographical site to minimize the performance impact of network latency. This task can be automated by enabling the database alignment feature within NFM-P.

NFM-P has been tested with up to 250ms of geographic latency. Specifically for the NFM-P database, Oracle doesn't provide any guidance on latency, other than adjusting TCP socket buffer sizes. If the NFM-P deployment includes the NFM-P auxiliary database, the latency between the active NFM-P auxiliary statistics collectors and the NFM-P auxiliary database must be less than 200ms, effectively reducing the tested geographic redundancy limit from 250ms to 200ms.

5.6.3 Optimizing throughput between NFM-P components

In high-speed, high-latency networks the TCP socket buffer size controls the maximum network throughput that can be achieved. If the TCP socket buffer is too small it will limit the network throughput, despite the fact that the available bandwidth might support much higher transfer rates.

Adjusting the TCP socket buffer size to achieve optimal network throughput may be necessary if the network bandwidth is more than 10Mbps and round-trip latency is higher than 25ms.

The optimal TCP socket buffer size is the bandwidth delay product (BDP). The bandwidth delay product is a combination of the network bandwidth and the latency, or round-trip time (RTT); basically, it is the maximum amount of data that can be in transit on the network at any given time.

For example, given a 20Mbps network with a RTT of 40ms the optimal TCP socket buffer size would be computed as follows:

```
BDP = 20 Mbps * 40ms = 20,000,000 bps * .04s = 800,000 bits / 8 = 100,000 bytes socket
buffer size = BDP  = 100,000 bytes
```

The RHEL documentation should be consulted to determine how to modify the TCP socket buffer size and ensure that the change is persistent.

It is important to note that increasing the TCP socket buffer size directly affects the amount of system memory consumed by each socket. When tuning the TCP socket buffer size at the operating system level, it is imperative to ensure the current amount of system memory can support the expected number of network connections with the new buffer size.

5.6.4 Additional NFM-P database throughput optimizations

In addition to the optimizations above, the NFM-P database workstation requires changes to the sqlnet.ora and listener.ora files that are contained in the oracle/network/admin directory. The lines with the SEND_BUF_SIZE and RECV_BUF_SIZE should be uncommented (delete the “#” character), and set to 3 times the BDP value calculated above. The database should be shutdown when this change is made.

5.7 Network reliability

5.7.1 Network reliability considerations

This section describes network reliability considerations.

5.7.2 Reliability between NFM-P components

The NFM-P requires reliable network communications between all the NFM-P components:

- NFM-P servers
- NFM-P databases
- NFM-P auxiliaries
- NSP flow collector
- NSP flow collector controller
- NFM-P auxiliary databases
- NSP analytics server
- NFM-P clients and NFM-P client delegate server
- NFM-P XML API clients

The performance and operation of NFM-P can be significantly impacted if there is any measurable packet loss between the NFM-P components. Significant packet loss can cause NFM-P reliability issues.

Nokia supports the deployment of NFM-P using the RHEL IP Bonding feature. The support for IP Bonding is intended only to provide network interface redundancy configured in active-backup mode for IP Bonding on the OS instance hosting the application software. All other modes of IP Bonding are not supported. RHEL documentation should be consulted on how to configure IP Bonding.

5.7.3 NFM-P server to NE network reliability

The NFM-P server requires reliable network connectivity between the NFM-P server/Auxiliary to the managed network elements. The mediation layer in NFM-P is designed to recover from lost packets between the NFM-P server and the network elements; however, these mechanisms come with a cost to performance. Any measurable packet loss will degrade performance of NFM-P's ability to manage the Network Elements. The loss of packets between NFM-P and NE will have an impact on (but not limited to):

- Any SNMP operations to the network elements:

- SNMP Trap processing performance
- Provisioning performance
- Provisioning failures
- Performance statistics collection (possibly to the point where statistics collection will be incomplete)
- STM test operation (initiating test and collecting results retrieval)
- NE discovery and resync performance
- NE discovery and resync failures
- scheduled polling for reachability checks
- Accounting statistics retrieval (possibly to the point where statistics collection will be incomplete)
- CLI session operation
- NE backup retrieval and software download performance

The following example highlights the significant impact of lost packets. It only considers the SNMP communication times with one network element. With the default mediation policy configured with an SNMP retry time-out of 10 seconds, and an average round trip latency of 50 ms between NFM-P server and the network element, NFM-P will spend a total of 25 seconds sending and receiving 1000 packets (500 SNMP gets and 500 SNMP responses). With a 0.1% packet loss (1 packet out of the 1,000) the NFM-P server will wait for the retry time-out (10 seconds) to expire before retransmitting. This will cause the time to complete the 500 SNMP gets to increase by 10 seconds – for a total of 35 seconds of communication time, or an increase of 40% over the time with no packet loss. With 0.5% packet loss, the 500 SNMP gets would increase by 50 seconds – for a total of 75 seconds to complete or an increase of 200%.

5.8 Network Element specific requirements

5.8.1 GNE, Nokia OmniSwitch, and 9500 / Wavence considerations

NFM-P clients support the web-based WebView functionality on OmniSwitch family of switches which requires direct network connectivity to the Network Element from the NFM-P client.

NFM-P clients support web-based clients on Generic Network Elements (GNEs) but require direct network connectivity between the NFM-P client and GNE.

9500 MPR / Wavence SM support includes the use of NEtO for specific management functions for these network element types. NEtO is a separate application that is installed along with the NFM-P client and launched through the NFM-P client UI. Please consult the node documentation for current memory requirements that are in addition to the NFM-P client memory requirements. The 9500 MPR / Wavence SM also uses a web interface for management.

6 Scaling

6.1 Scalability guidelines

6.1.1 Scalability limits

Table 6-1, "NFM-P Release 22.3 scalability limits" (p. 67) represents the scalability limits for Release 22.3. Note that:

- These limits require particular hardware specifications and a specific deployment architecture.
- Scale limits for all network elements assume a maximum sustained trap rate of 100 traps/second for the entire network. NFM-P's trap processing rate depends on many factors including trap type, NE type, NE configuration, NE and network latency, network reliability as well as the size and speed of the servers hosting the NFM-P application. NFM-P scalability testing runs at a sustained trap rate exceeding 100 per second for the largest deployment and server configurations.

Chapter 3, "Platform requirements" contains information on identifying the correct platform for a particular network configuration. To achieve these scale limits, a distributed NFM-P configuration is required, and may also require an NFM-P auxiliary statistics collector and a storage array for the NFM-P database workstation.

Consult Nokia personnel to ensure you have the correct platform and configuration for your network size.

Table 6-1 NFM-P Release 22.3 scalability limits

Attribute of managed network	Scaling limit
Maximum number of managed MDAs	60,000
Maximum number of Network Elements	50,000
Maximum number of GNEs [1]	50,000
Maximum number of managed services	4,000,000
Maximum number of optical transport services	20,000
Maximum number of 1830 VWM RMUs	60,000
Maximum number of SAPs	12,000,000
Maximum number of simultaneous NFM-P GUI sessions	250
Maximum number of simultaneous web UI client sessions	4–250 Table 6-3, "NFM-P apps maximum number of concurrent sessions" (p. 69)
Maximum number of simultaneous active XML API HTTP applications	30
Maximum number of simultaneous active XML API JMS applications	20

Table 6-1 NFM-P Release 22.3 scalability limits (continued)

Attribute of managed network	Scaling limit
Maximum number of outstanding alarms	50,000
Maximum number of outstanding alarms - Distributed Configuration	250,000
Maximum number of Historical Alarms	9,600,000
Maximum number of TCAs	250,000
Maximum number of monitored services in the Service Supervision application	1,000,000
Maximum number of concurrent NSP analytics users	10

Notes:

1. The number of interfaces on a GNE and the traps that may arise from them is the key factor determining the number of GNE devices that can be managed. As GNE devices are expected to be access devices the sizing is based on an average of 10 interfaces of interest on each device (10 x 50,000 = 500,000 interfaces). Processing of traps from interface types that are not of interest can be turned off in NFM-P. Under high trap load, NFM-P may drop traps.

   NFM-P uses the number of MDAs as the fundamental unit of network dimensioning. To determine the current or eventual size of a network, the number of deployed or expected MDAs, as opposed to the capacity of each router, must be calculated.

Table 6-2 Network element maximums and equivalency

Network element type	Maximum number of network elements supported	MDA equivalency
7750, 7450, 7710	50,000	1 MDA = 1 MDA [1] [2]
7705	50,000	50,000
7250 IXR-6 / 7250 IXR-10 / 7250 IXR-R4 / 7250 IXR-R6	50,000	1 MDA = 1MDA
7250 IXR-s / 7250 IXR-e	25,000	50,000
7210	50,000	50,000
OMNISwitch 6250, 6400, 6450, 6850, 6855 (each shelf in the stackable chassis)	50,000	50,000
OMNISwitch 6350, 6465, 6560, 6865 (each shelf in the stackable chassis)	5,000	5,000
OMNISwitch 6860, 6860E, 6860N	5,000	5,000
OMNISwitch 6900	800	800
OMNISwitch 9600, 9700, 9700E, 9800, 9800E (each NI)	1,000	1,000
OMNISwitch 10K (each NI)	400	400
CMM	320	[3]

Table 6-2 Network element maximums and equivalency (continued)

Network element type	Maximum number of network elements supported	MDA equivalency
CMU	320	[3]
9500 MPR / Wavence SM	15,000	15,000
1830 VWM OSU	2,000	[4]
VSC	1	N/A

Notes:

1. The IMM card has an MDA equivalency of 2 MDAs per card.
2. The CMA card has an MDA equivalency of 1 MDA per card.
3. The CMM has an MDA equivalency of 1 MDA per blade / VM.
4. The 1830 VWM OSU Card Slot has an MDA equivalency of 1/4 MDA per card to a maximum MDA equivalency of 30,000

Table 6-3 NFM-P apps maximum number of concurrent sessions

NFM-P application	Maximum number of concurrent sessions
Analytics	10
Fault Management	250
Help Center	250
Network Supervision	50
Service Supervision	250
Subscriber Management	5
Wireless Supervision	50
Wireless NE Views	50

6.1.2 NFM-P performance targets

Table 6-4, “NFM-P Release 21 performance targets” (p. 70) represents the performance targets for the NFM-P. Factors that may result in fluctuations of these targets include:

- NFM-P server and NFM-P database system resources
- network activity
- user/XML-API activity
- database activity
- network size
- latency

Table 6-4 NFM-P Release 21 performance targets

Performance item description	Target
NFM-P client GUI performance	
Time to launch an NFM-P client GUI	1 - 2 minutes
Time to launch an NFM-P client GUI configuration form	~5 seconds
Time to save an NFM-P client GUI configuration form	~2 seconds
NFM-P server performance	
Time to restart the NFM-P server	15 - 30 minutes (subject to network dimensions)
NFM-P database Backup (without statistics)	Up to 60 minutes (subject to network size)
NFM-P database Restore	~45 minutes
NFM-P server activity switch	10 - 30 minutes (subject to network dimensions)
NFM-P DB switchover (by invoking through the GUI)	<10 minutes
NFM-P DB failover	30 minutes when managing maximum number of devices
Recovery of standby NFM-P database after failover	<75 minutes
Upgrade Performance	
NFM-P client Upgrade	~10 minutes
NFM-P complex upgrade (server, database, auxiliaries) [1]	<6 hours
NFM-P upgrade maximum visibility outage with NFM-P redundant system [2]	15 - 30 minutes

Notes:

1. The target includes the installation of the software on the existing servers and NFM-P database conversion. Operating System installation/upgrades, patching, pre/post-upgrade testing and file transfers are excluded from the target.
2. Provided proper planning and parallel execution procedures were followed.

6.2 Scaling guidelines for NFM-P XML API clients

6.2.1 XML-API client limits

There can be a maximum of 20 NFM-P XML API JMS clients. Greater than 10 NFM-P XML API JMS clients requires an NFM-P server with a minimum of 16 CPU Cores.

The number of NFM-P XML API HTTP clients supported by an NFM-P server workstation is 2 times the number of CPU cores with at least 10 and at most 30 clients supported.

The maximum number of concurrent findToFile operations supported is five. The maximum number of concurrent control script executions is five.

6.2.2 XML API JMS client messaging rates

Network latency between the NFM-P server and an NFM-P XML API client will reduce the JMS message rate. For durable JMS clients, the *Duplicate OK* method will allow for a higher message rate than the *Auto Acknowledge* method. Refer to the *NSP NFM-P XML API Developer Guide* for more information.

NFM-P is also able to deliver hundreds of messages per second to a non-durable NFM-P XML API client.

Table 6-5 JMS durable messaging rates

JMS messaging	Round-trip latency from the XML API client to the NFM-P server		
	0ms	20ms	40ms
Durable connection with Auto-acknowledge (messages/s)	1000	692	349
Durable connection with Duplicates-OK (messages/s)	1000	693	347

6.3 Scaling guidelines for statistics collection

6.3.1 Statistics collection

NFM-P provides the ability to collect statistics information from the network elements. This section provides guidelines that can be used to determine the extent to which statistics collection can be retrieved from the network.

6.3.2 Statistics collection definitions

Performance statistics: These statistics are associated with various network objects such as ports, interfaces, channels and network elements (routers). These statistics are retrieved by NFM-P using SNMP polling according to the MIB policies that are configured by the user.

Accounting statistics: These statistics are associated with Services, Subscribers, and Network Interfaces and contain data that can be used for accounting, billing and SLA management purposes. These statistics are collected on the 7x50 and retrieved by NFM-P via a file that is transferred via ftp/sftp.

Application Assurance Accounting statistics: These statistics are associated with Subscribers, SAPs, and spoke SDP bindings and contain data related to traffic flows that can be used for QoS and traffic management, and application aware reporting. These statistics are collected on the 7x50 ISA cards and retrieved by NFM-P via a file that is transferred via ftp/sftp.

Statistics Item: An individual statistics counter, such as RxOctets or TxFrames.

Statistics Record: A collection of statistics items which is retrieved from the router and stored in the NFM-P database as an atomic operations. In the various statistics forms on the NFM-P GUI client, a statistics record appears to the user as a single row which contains the collection or retrieval timestamp and a set of individual statistics items. In the case of performance statistics, a statistics record corresponds to a row in the MIB table.

6.3.3 Determining the number of statistics records that will be collected

Statistics can be collected and processed by the NFM-P server or by the NFM-P auxiliary statistics collector for dedicated statistics handling. The NFM-P auxiliary statistics collector provides a dedicated workstation for statistics collection. The following sections should be used to determine the maximum performance and accounting statistics for different hardware setups.

6.3.4 Performance statistics

Refer to the *NSP NFM-P Statistics Management Guide* to find the steps required to configure NFM-P to retrieve and process performance statistics. Note that two steps are required to enable the collection of performance statistics from the network. First, a policy is defined which specifies a set of polling periods for various MIBs. Second, the policy is applied to a number of network elements.

In general, enabling the statistics collection of a MIB will result in one statistics record being collected, at the specified polling period, for each network object to which the MIB applies.

For example, consider a policy is created with only the rtr.L2AccessDhcpRelayCfgStats MIB enabled for collection at 15-minute intervals. That policy is assigned to only two network elements which each contain 500 L2 Access Interfaces. As a result of this action, NFM-P will collect 1,000 statistics records from the network every 15 minutes.

The quantity of resources which are allocated to the retrieval and processing of performance statistics does not depend significantly on the number of CPU Cores available to the NFM-P server or auxiliary statistics collector software. The tables below show the maximum number of performance statistics that can be retrieved and processed by the NFM-P server and the NFM-P auxiliary statistics collector every 15 minutes.

Table 6-6 Maximum number of performance statistics records processed by an NFM-P server

Number of CPU cores on NFM-P server workstations	Maximum number of performance statistics records per 15-minute interval	
	Collocated configuration	Distributed configuration
6 or greater	50,000	150,000

Table 6-7 Maximum number of performance statistics records processed by an NFM-P statistics auxiliary

Number of active auxiliary statistics collectors	Maximum number of performance statistics records per 15-minute interval				
	Statistics collection with NFM-P database		Statistics collection with single auxiliary database	Statistics collection with three+ auxiliary database cluster	logTofile only
	8 CPU Cores, 32GB RAM	12 CPU Cores, 32GB RAM	8 CPU Cores, 32GB RAM	12 CPU Cores, 32GB RAM	12 CPU Cores, 32GB RAM
1	500,000	2,000,000	500,000	2,000,000	2,000,000
2	500,000	2,000,000	500,000	4,000,000	4,000,000
3	500,000	2,000,000	500,000	4,000,000	4,000,000

In situations where NFM-P is asked to collect more performance statistics than it can process in the specified polling period, the *PollerDeadlineMissed* alarms will start appearing. These alarms

indicate to the user that the polling mechanisms within NFM-P cannot retrieve the requested information within the specified polling period. Should this situation arise, the polling period for statistics should be increased or the number of objects that are applied to Statistics Poller Policies should be reduced.

6.3.5 Performance statistics collection and network latency

NFM-P collection of performance statistics from a single network element may be limited due to the round trip delay caused by network and network element latency. NFM-P collects performance statistics records using SNMP. One record is collected at a time to limit the load on the network element. Therefore, round trip latency will directly impact the maximum number of performance statistics records collected. As an example, if the round trip latency is 100ms, and we target a completion time of 66% of the collection interval (to allow for processing variances and other system impacts), the maximum number of performance statistics records that can be collected from one network element in a 15 minute interval would be 6000 records (66% of 900 seconds divided by 100 ms latency).

6.3.6 Accounting statistics

Refer to the *NSP NFM-P Statistics Management Guide* to find the steps required to configure NFM-P to retrieve and process accounting statistics.

The quantity of resources which are allocated to the retrieval and processing of accounting statistics within the NFM-P server or auxiliary statistics collector are set at the installation time and depend on the number of CPU Core available to the NFM-P server or auxiliary statistics collector software. The number of CPU Cores available to the server depends on the number of CPU Cores on the workstation and whether the NFM-P database software is collocated with the NFM-P server software on the same workstation.

An accounting statistic record is the statistic for one queue for one SAP. For example, if 2 ingress and 2 egress queues are configured per SAP, the “Combined Ingress/Egress” statistic represents 4 NFM-P accounting statistic records.

It is recommended that the Accounting Policy Interval and the File Policy Interval be aligned to the same period. Misalignment of the policy periods can cause NFM-P resource contention for both performance and accounting statistics processing.

The following tables provide the maximum number of accounting statistics records that can be retrieved and processed by the NFM-P server or NFM-P auxiliary statistics collector in various situations.

To reach the peak accounting statistics collection from the NFM-P auxiliary statistics collector workstation, the NFM-P database workstation requires a customized configuration that can be obtained from Nokia personnel.

Table 6-8 Maximum number of accounting statistics records processed by an NFM-P server workstation

Number of CPU cores on NFM-P server workstations	Maximum number of accounting statistics records per 15-minute interval	
	Collocated configuration	Distributed configuration
6	100,000	200,000
8 or greater	200,000	400,000

Table 6-9 Maximum number of accounting statistics records processed by an NFM-P statistics auxiliary

Number of active auxiliary statistics collectors	Maximum number of accounting statistics records per 15-minute interval				
	Statistics collection with NFM-P database		Statistics collection with single auxiliary database	Statistics collection with three+ auxiliary database cluster	logToFile only
	8 CPU cores, 32GB RAM	12 CPU cores, 32GB RAM	8 CPU cores, 32GB RAM	12 CPU cores, 32GB RAM	12 CPU cores, 32GB RAM
1	10,000,000	10,000,000	5,000,000	20,000,000	20,000,000
2	10,000,000	10,000,000	5,000,000	40,000,000	40,000,000
3	10,000,000	10,000,000	5,000,000	60,000,000	60,000,000

In situations where NFM-P is asked to collect more accounting statistics records than it can process in the specified retrieval period, the extra statistics will not be retrieved from the network.

There are two methods to export accounting and performance statistics from NFM-P; registerLogToFile, and findToFile. The registerLogToFile method is the preferred method and is required for situations where more than 400,000 accounting statistics records are retrieved in 15 minutes or 500,000 performance statistics are retrieved in 15 minutes.

6.3.7 Application assurance accounting statistics

Refer to the *NSP NFM-P Statistics Management Guide* to find the steps required to configure NFM-P to retrieve and process application assurance accounting statistics.

The quantity of resources which are allocated to the retrieval and processing of application assurance accounting statistics within the NFM-P server are set at the installation time and depend on the number of CPUs available to the NFM-P server software. The number of CPUs available to the NFM-P server depends on the number of CPUs on the workstation and whether the NFM-P database software is collocated with the NFM-P server software on the same workstation.

Scaling of application assurance collection is related to the number of objects configured for collection as opposed to the number of records collected per interval.

The following tables provide the maximum number of application assurance objects that can be configured for collection by the NFM-P server or NFM-P auxiliary statistics collector in various situations.

Table 6-10 Maximum number of application assurance accounting objects configured for collection by an NFM-P server workstation

Number of CPU cores on NFM-P server workstations	Maximum number of application assurance accounting objects configured for collection per 15-minute interval	
	Collocated configuration	Distributed configuration
6	50,000	100,000
8 or greater	100,000	200,000

Table 6-11 Maximum number of application assurance accounting objects configured for collection by an NFM-P statistics auxiliary

Number of active auxiliary statistics collectors	Maximum number of application assurance accounting objects configured for collection per 15-minute interval			
	Statistics collection with NFM-P database		Statistics collection with single auxiliary database	Statistics collection with three+ auxiliary database cluster
	8 CPU Cores, 32GB RAM	12 CPU cores, 32GB RAM	8 CPU cores, 32GB RAM	12 CPU cores, 32GB RAM
1	5,000,000	7,500,000	1,000,000	5,000,000
2	5,000,000	15,000,000	1,000,000	10,000,000
3	5,000,000	15,000,000	1,000,000	20,000,000

In situations where NFM-P is asked to collect more application assurance accounting records than it can process in the specified retrieval period, the extra statistics will not be retrieved from the network.

6.3.8 Exporting performance and accounting statistics records

There are two methods to export accounting and performance statistics from NFM-P; registerLogToFile, and findToFile. The registerLogToFile method is the preferred method and is required for situations where more than 400,000 accounting statistics records are retrieved in 15 minutes or 500,000 performance statistics are retrieved in 15 minutes. This recommendation also minimizes collection latency and reduces system load.

6.3.9 NFM-P database hardware platform requirements

To collect large numbers of statistics using the NFM-P database, there are RAM and storage I/O requirements for the NFM-P database workstation. The table below highlights these requirements.

Table 6-12 NFM-P database workstation hardware requirements for a distributed configuration

Maximum number of simultaneous statistics records per 15-minute interval			NFM-P auxiliary statistics collector(s)	Requires the following NFM-P database workstation setup
Accounting statistics records	Application assurance accounting objects configured for collection	Performance statistics records		
400,000	0	0	No	4 CPU cores, minimum 2.0GHz [1] 4 disks (RAID 0) 32 GB RAM
0	200,000	0		
0	0	150,000		
800,000	0	0	Yes	4 CPU cores, minimum 2.0GHz [1] 4 disks (RAID 0) 48 GB RAM
0	400,000	0		
0	0	200,000		
10,000,000	0	500,000	Yes	8 CPU cores, minimum 2.0GHz [1] 6 disks (RAID 0) 64 GB RAM
0	5,000,000	500,000		
10,000,000	5,000,000	2,000,000	Yes	12 CPU cores, minimum 2.0GHz [1] 6 disks (RAID 0) 64 GB RAM
0	15,000,000	0		

Notes:

1. 2.0GHz only supported on Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz

6.3.10 Simultaneous collection of performance, application assurance accounting and accounting statistics records

NFM-P can collect performance, application assurance, and accounting statistics records simultaneously. However, it is important to consider that enabling the collection of one type of statistics will reduce the capability of NFM-P to collect and process the other type of statistics. It is therefore not possible to achieve the maximum stated limits for performance, application assurance, and accounting statistics records simultaneously, in certain configurations. Table 6-12, “NFM-P database workstation hardware requirements for a distributed configuration” (p. 76) shows an example of simultaneous collection.

6.3.11 Determining the number of performance and accounting statistics records being collected by NFM-P

To ensure the number of performance and accounting statistics records that NFM-P is asked to collect and process every 15 minutes remains below the stated scalability guidelines, it is important to carefully assess the impact of creating and assigning statistics policies. Review the number of objects that are assigned to statistics policies and ensure the polling and retrieval periods are set such that the numbers will remain below the stated guidelines.

Using NFM-P server performance statistics, NFM-P can assist in determining how many polled and accounting statistics are being collected.

NFM-P performance can be adversely affected by increasing the number of historical statistics entries recorded by the NFM-P. NFM-P system impacts include increased time listing log records from the GUI and XML API clients, increased database space, and increased database backups times.

6.3.12 Statistics record retention

The table below shows the different retention rates that are achievable depending upon the collection rate and statistic type.

Table 6-13 Maximum statistics interval retention - NFM-P database

Statistics type	Total number of statistics records to be stored in the database	Maximum number of retention intervals
Performance	<40M	672
	>40M	96
Accounting	<40M	672
	>40M	16

Table 6-14 Maximum statistics interval retention - auxiliary database

Statistics type	Maximum number of retention intervals
Performance	35,040
Accounting	35,040

When using the logToFile method only, for collection, the maximum retention of data on the file system is 600 minutes (10 hours).

6.4 Scaling guidelines for scheduled tests (STM)

6.4.1 Scheduled tests (STM)

NFM-P provides the ability to generate, manage and schedule STM tests within the network. This section provides guidelines that can be used to determine the extent to which STM tests can be scheduled and launched within a network.

There are a number of factors which will influence NFM-P's ability to concurrently manage and schedule a large number of tests. NFM-P keeps track of how many tests are running concurrently. This is to limit the initiation of the tests, and the processing of the results without interfering with the system's other functions.

To understand the STM guidelines, the following terminology is required:

Elemental Test: An OAM test to be sent to a router such as an LSP ping

Elemental Test Result: An OAM test result received from a network element

Accounting file Test: An OAM test that is initiated in the default manner, however, the test results are retrieved from the network element via FTP on a periodic basis.

Test Policy: A definition or configuration that tells NFM-P the specifics about how to generate a test. A test policy can contain multiple test definitions. The policies are used by test suites.

Test Suite: A collection of elemental tests that can be assigned to a specific schedule. There are three defined sections in which tests can be placed within a test suite: First run, Generated and Last run. The tests are executed in order by these sections. It is possible to configure the execution order of tests within the First Run and Last Run sections to be parallel or sequential. The tests in the Generated position are run by the system as concurrently as possible. If the Generated section contains tests from several different test definitions, then all the tests belonging to one definition will be executed before the tests of the next definition begin. Within a definition, the system will attempt to execute the tests as concurrently as possible. This is important to note, as a test suite containing a large number of tests in the Generated section (or in the First Run/Last Run sections set to parallel) may tax the system. Part of the increased stress placed on the system by concurrent tests is a result of the need for the system to use greater amounts of resources in order to initiate, wait for and process many tests concurrently. As well, tests that result in a large amount data to be returned from the routers will place increased demands on the NFM-P.

Schedule: A start time that can have a test suite or test suites assigned to it to produce scheduled tasks. When the schedule's start time is reached, the suite or suites assigned to it will commence. The schedule may be set to continuously repeat after a configurable period of time.

Scheduled Task: An instance of a test suite assigned to a schedule

Non -NE Schedulable STM Tests: NFM-P provides the ability to execute and process results for non NE schedulable tests. Non NE schedulable tests are elemental tests which are not persistently defined on network elements; rather, these tests are defined/configured from NFM-P per test execution. Elemental test results from non-NE schedulable tests are always regular (SNMP mediated) and share the same scale limits/considerations as regular scheduled STM tests.

Table 6-15 Maximum number of STM elemental test results

NFM-P platform	Maximum regular STM elemental test results (SNMP mediated schedulable/ non-NE schedulable) in a 15–minute period	Maximum accounting file STM elemental test results in a 15–minute period with results stored in the NFM-P database or NFM-P database and using logToFile	Maximum accounting file STM elemental test results in a 15–minute period using logToFile only
Distributed NFM-P configuration with minimum 8 CPU Core NFM-P server	15,000	1,500,000 [1]	1,500,000 [1]
Distributed NFM-P configuration NOTE: It may be possible to achieve higher numbers depending on the NFM-P server activity and hardware platform	6,000	22,500	60,000

Table 6-15 Maximum number of STM elemental test results (continued)

NFM-P platform	Maximum regular STM elemental test results (SNMP mediated schedulable/ non-NE schedulable) in a 15–minute period	Maximum accounting file STM elemental test results in a 15–minute period with results stored in the NFM-P database or NFM-P database and using logToFile	Maximum accounting file STM elemental test results in a 15–minute period using logToFile only
Minimum Supported Collocated NFM-P configuration NOTE: It may be possible to achieve higher numbers depending on the NFM-P server activity and hardware platform	3,000	1,500	15,000

Notes:

1. may require a dedicated disk or striped disks for the xml_output partition

6.4.2 Guidelines for maximizing STM test execution

By default, NFM-P will only allow test suites with a combined weight of 80,000 to execute concurrently. The test suite weights are identified in the NFM-P GUI's Test Suites List window. Running too many tests that start at the same time will cause the system to exceed the previously mentioned limit, and the test will be skipped. Ensuring the successful execution of as many STM tests as possible requires planning the schedules, the contents, and the configuration of the test suites. The following guidelines will assist in maximizing the number of tests that can be executed on your system:

- When configuring tests or test policies, do not configure more packets (probes) than necessary, as they increase the weight of the test suite.
- Test suites with a smaller weight will typically complete more quickly, and allow other test suites to execute concurrently. The weight of the test suite is determined by the number of tests in the test suite, and the number of probes that are executed by each test. See Table 6-16, "OAM test weight" (p. 80) for test weight per test type.
- Assign the time-out of the test suite in such a way that if one of the test results has not been received it can be considered missed or failed without stopping other test suites from executing.
- Rather than scheduling a test suite to execute all tests on one network element, tests should be executed on multiple network elements to allow for concurrent handling of the tests on the network elements. This will allow the test suite results to be received from the network element and processed by NFM-P more quickly freeing up available system weight more quickly.
- Rather than scheduling a test suite to run sequentially, consider duplicating the test suite and running the test suites on alternating schedules. This allows each test suite time to complete or time-out before the same test suite is executed again. Remember that this may cause double the system weight to be consumed until the alternate test suite has completed.
- Create test suites that contain less than 200 elemental tests. This way you can initiate the tests at different times by assigning the test suites to different schedules thereby having greater control over how many tests are initiated or in progress at any given time.

- Prioritize which tests you wish to perform by manually executing the test suite to determine how long it will take in your network. Use that duration with some added buffer time to help determine how much time to leave between schedules or repetitions of a schedule and how to configure the test suite time-out.
- A test suite time-out needs to be configured to take effect before the same test suite is scheduled to run again, or it will not execute if it does not complete before the time-out.
- NFM-P database backups can impact the performance of STM tests.

Table 6-16 OAM test weight

Test type	Weight
Regular Elemental STM Test	10 per Test Packet
Accounting File Elemental STM Test	1

6.4.3 Accounting file STM test configuration

Accounting file collection of STM test results requires 7750 and 7450 network elements that are version 7.0 R4 and above. To take advantage of accounting file STM test execution, the test policy must be configured to be NE schedulable with “Accounting file” selected. This will produce STM tests that will be executed on the network element, while the test results are collected by the NFM-P server by way of an accounting file in a similar way to accounting statistics. Accounting file STM test results are collected by the NFM-P server only.

NFM-P supports the use of logToFile for file accounting STM results. When using this method only for results, the number of tests that can be executed per 15 minute interval is increased. Refer to Table 6-15, “Maximum number of STM elemental test results” (p. 78) for specific scaling limits. The logToFile method for file accounting STM results supports a maximum of two JMS clients.

6.4.4 Examples of STM test configuration

The following examples describe the configuration of STM tests on different network configurations.

6.4.5 Example 1

Assume there is a network with 400 LSPs and that the objective is to perform LSP pings on each LSP as frequently as possible. The following steps are to be followed:

1. Create 4 test suites each containing 100 elemental LSP ping tests
2. One at a time, execute each test suite and record the time each one took to complete. Assume that the longest time for executing one of the test suites is 5 minutes.
3. Create a schedule that is ongoing and has a frequency of 15 minutes. This doubles the time taken for the longest test suite and ensures that the test will complete before it is executed again. Assign this schedule to the 4 test suites.
4. Monitor the test suite results to ensure that they are completing. If the tests are not completing (for example getting marked as “skipped”), then increase the frequency time value of the schedule.
5. In the above case, there are 200 elemental tests configured to be executed each 10 minutes.

6.4.6 Example 2

Assume there are eight test suites (T1, T2, T3, T4, T5, T6, T7 and T8), each containing 50 elemental tests. Assume the test suites individually take 5 minutes to run. Also, assume the objective is to schedule them so that the guideline of having less than 200 concurrently running elemental tests is respected.

The recommended approach for scheduling these tests suites is as follows:

- Test suites T1, T2, T3, T4 can be scheduled on the hour and repeat every 10 minutes
- Test suites T5, T6, T7, T8 can be scheduled on the hour + 5 minutes and repeated every 10 minutes

This will ensure no more than 200 elemental tests are scheduled to run concurrently.

6.4.7 Factors impacting the number of elemental tests that can be executed in a given time frame

The following factors can impact the number of elemental tests that can be executed during a given time frame:

- The type of tests being executed. Each type of elemental test takes varying quantities of time to complete (e.g. a simple LSP ping of an LSP that spans only two routers may take less than 2 seconds; an MTU ping could take many minutes).
- The amount of data that is generated/updated by the test within the network elements. NFM-P will have to obtain this information and store it in the NFM-P database. The quantity of data depends on the type of tests being performed and the configuration of the objects on which the tests are performed.
- The number of test suites scheduled at or around the same time
- The number of tests in a test suite
- The number of routers over which the tests are being executed. Generally, a large number of tests on a single router can be expected to take longer than the same number of tests distributed over many routers.
- An NFM-P database backup may temporarily reduce the system’s ability to write test results into the database.
- The workstation used to perform the tests will dictate how many physical resources NFM-P can dedicate to executing elemental tests. On the minimum supported workstation (collocated NFM-P server and NFM-P database on a single server), the number of concurrent tests must be limited to 3,000.

6.4.8 Possible consequences of exceeding the capacity of the system to perform tests

NFM-P will exhibit the following symptoms if the number of scheduled tests exceeds the system’s capacity.

6.4.9 Skipped tests

If a test suite is still in progress at the time that its schedule triggers again, then that scheduled task will be marked as skipped and that test suite will not be attempted again until the next scheduled time.

6.4.10 Failed tests (time-out)

Tests may time-out and get marked as failed. If any of the tests take more than 15 minutes it may get purged from an internal current test list. For example, a test may be successfully sent to a router and the system does not receive any results for 15 minutes. The system marks the test as failed and purges its' expectation of receiving a result. However, later, the system could still receive the results from the router and update its result for the test to success.

6.4.11 Disk space requirements for STM test results

STM test results are stored in the tablespace DB partition. The STM database partitions start with a total size of 300MB of disk space. When the maximum number of test results is configured at 20,000,000 (maximum), the disk space requirement for the STM tests may increase by up to 80GB. A larger tablespace partition should be considered.

The maximum number of test results stored in the database reflects the sum of the aggregate results, test results, and probe results.

Running 10 tests with 1 probe each versus 1 test with 10 probes consumes the same amount of disk space.

When using logToFile for accounting file STM test results, the maximum time-to-live on the disk is 24 hours. At the maximum collection rate of 1,500,000 test results per 15 minutes, the storage requirements on the NFM-P server in the xml_output directory is 600GB per JMS client. The storage requirements are doubled if using the maximum number of JMS clients for file accounting STM results. The disk storage requirements can be decreased by using the compress option for logToFile but will result in increased CPU utilization on the NFM-P server.

6.5 cflowd statistics collection

6.5.1 Scaling guidelines for cflowd statistics collection

The table below shows the scaling limits for an NSP flow collector in its ability to process cflowd flow records from the network and produce IPDR formatted files. The guidelines are divided into the NSP flow collector collecting in a Residential/Mobile deployment and the NSP flow collector collecting in a Business deployment.

For Residential and Mobile deployments, the only statistics types that should be in use are Volume, Comprehensive, Unknown and corresponding Special Study types. TCP and RTP are not supported in Mobile, and although the statistics types are available for Residential deployments, the use case is not, and there are no reports for this data. Additionally, even for Residential, the only reports available are for Comprehensive. Comprehensive statistics have a fixed 60min aggregation interval. Due to the amount of data generated in a Mobile deployment, Volume statistics require an aggregation interval of 60 minutes. As an alternative, Volume Special Study statistics on specific subscribers can be used. The only key factor of difference is whether or not additional counters are enabled for Comprehensive statistics.

Table 6-17 cflowd statistics scaling limits for residential and mobile deployments

NSP flow collector processing rate in flows per second	Counter selection [1]	Maximum number of unique objects in memory [2]	Packet loss per hour [3]
100,000 FPS	Default two counters	100M Objects	<= 2%
	All counters	60M Objects	<= 1%

Notes:

1. Default: two counters. Volume: total bytes/total packets. Comp-volume: total bytes StoC/CtoS sum unknown. Only one counter exists. Vol SS: should be minuscule. All counters: Comp-volume has a total of ten counters that can be enabled.
2. Number of aggregated output requests that are sent to the server every 60 minutes. Assumes transfer has sufficient bandwidth to complete in a timely manner.
3. Packet loss may increase if communication between the NSP flow collector and target file server is interrupted.

For business deployments, in addition to the statistics types with a small number of records; Comprehensive, Volume, Unknown, and Volume Special Study, there are also statistics types with a larger number of records; TCP Performance, and RTP (Voice/Audio/Video). The main distinction is whether or not the TCP/RTP statistics types use the default enabled counters, or if all counters have been enabled. Enabling all of the TCP/RTP counters increases the amount of memory used by the NSP flow collector. Aside from the incoming FPS (Flows Per Second) that the NSP flow collector can process, the other main factor putting pressure on the NSP flow collector is the memory used by the number of unique objects/records (or unique routes, i.e. the # of output records the NSP flow collector produces in the IPDR files) in NSP flow collector memory at any one time. And finally the interval size – the smaller the aggregation interval, the greater percentage of the next interval time will overlap with the transfer time of the previous interval – during this time the NSP flow collector must store objects in memory from two different intervals. Comprehensive statistics types are fixed at 60 minute intervals.

A unique object/route for TCP/Volume records in the business context is:

SAP, App/AppGroup, Interval ID, Src Group ID, Source Interface ID, Dest Group ID, Dest Interface ID

A Volume record will also have a direction field. Volume records coming from the router to the NSP flow collector will result in two output records in the IPDR files (one for each direction). For TCP, two incoming records from the NSP flow collector (one for each direction) will be combined by the NSP flow collector into a single output TCP record in the IPDR files.

A unique object/route for COMPREHENSIVE record in the business context is:

SAP, App/AppGroup, Interval ID, Src Group ID, Source Interface ID, Dest Group ID, Dest Interface ID

and either a hostname field, or three device identification fields.

A unique object/route for RTP is defined as:

Every single flow into the NSP flow collector is a unique route and an equal number of flow records are produced in the IPDR file. The expected number of RTP records sent from 7750 SR Routers is expected to be a small percentage of the total flows (i.e. <5% total flows TCP/VOL/RTP)

Table 6-18 cflowd statistics scaling limits for business deployments

NSP flow collector processing rate in flows per second	Statistic types used and counters used [1]	Maximum number of unique objects in memory [2]	Packet loss per hour [3]
100,000 FPS	Comprehensive/Volume/ Unknown/Vol S.S Only All Counters	60M objects	<= 1%
	TCP/TCP S.S Only: Default Counter	25M objects	<= 1%
	TCP/TCP S.S Only: All Counters	15M objects	<= 1%
	RTP Only: Default Counters	10M objects	<= 1%
	RTP Only: All Counters	3M objects	<= 1%
	Combined Comprehensive/Volume/ Unknown/TCP/RTP (including Special Study)	20M Comp/Volume/ Unknown + 5M TCP (All Cnt) + 0.5 RTP (All Cnt)	<= 1%

Notes:

1. Comprehensive/Volume/ Unknown/Volume SS: All Counters RTP/TCP/TCP S.S Counter Selection Default Counters: Leaving default enabled counters on All Counters: Enabling all available counters for given stat type. There are 40-60 total counters available for TCP and RTP types.
2. Number of aggregated output requisitions that are sent to the server every 60 seconds. Assumes transfer has sufficient bandwidth to complete in a timely manner.
3. Packet loss may increase if communication between the NSP flow collector and target file server is interrupted

6.6 PCMD collection

6.6.1 Scaling guidelines for PCMD collection

LTE management networks support the collection of per call measurement data (PCMD) from SGW, PGW, and ePDG network elements when using the NFM-P auxiliary PCMD collector. A single NFM-P auxiliary PCMD collector can process up to 18 million records per minute where multiple auxiliary PCMD collectors can be deployed to increase overall scaling limits. The scaling limits for a single auxiliary PCMD collector represents approximately 10M Bearers.

6.7 CPAM and vCPAA

6.7.1 Scaling guidelines for CPAM and vCPAA

The following section outlines the tested limits for the CPAM component of NFM-P.

Note: All CPAAs monitoring a contiguous network must belong to the same administrative domain.

The scalability of 7701 CPAA Hardware revision 2 and vCPAA is described in the following table:

Table 6-19 7701 CPAA Hardware revision 2 and vCPAA scalability

Item description	NFM-P 22.3
Maximum Number of Routers Supported with both IGP and BGP turned-on in the same 7701 CPAA (larger node count must use two separate 7701 CPAAs for IGP and BGP) with 2,200,000 BGP combined routes	500
Maximum Number of IGP Routers per 7701 CPAA if BGP is deployed on separate 7701 CPAA (routers can all be in one or multiple areas and count includes Nokia P/PE & 3rd party routers)	700
Maximum Number of OSPF Domains/Areas per 7701 CPAA	One Administrative Domain per 7701 CPAA Up to 50 areas per 7701 CPAA
Maximum Number of ISIS regions/area IDs per 7701 CPAA	One Administrative Domain per 7701 CPAA Up to 32 L1s + L2 per 7701 CPAA
Maximum Number of iBGP Peers mesh/RR	170
BGP/ MP-BGP Route Count	4,000,000 combined
Maximum Number of IPv4/VPN IPv4 Monitored Prefixes (combined)	2,000 configured 2,000 monitored
Maximum number of Sub-ASes	1
BGP stats	48,000/5 minutes 144,000/15 minutes
IP (LDP LSP, GRE, IP) Path Monitor – Configured	20,000

The scalability of CPAM is described in the following table:

Table 6-20 CPAM scalability

Item description	NFM-P 22.3
Maximum number of routers per admin domain	4,000
Maximum Active 7701 CPAAs	40
Maximum Number of IGP Routers	12,000
Maximum Number of OSPF Domains/Areas	150
Maximum Number of ISIS regions/area IDs	50
Maximum Number of BGP/MP-BGP Routes	4,000,000 combined
Maximum number of Sub-ASes	40
BGP stats	48,000/5 minutes 144,000/15 minutes
Combined IP and LSP path monitors configured	The numbers below can be combined (60,000)
RSVP LSP Path Monitored – Configured	60,000

Table 6-20 CPAM scalability (continued)

Item description	NFM-P 22.3
IP (LDP LSP, GRE, IP) Path Monitor – Configured	20,000
Recommended LSP Path monitor statistics polling interval	15 min
Maximum NFM-P database space allocated for IGP Checkpoints	10 GB
Maximum number of supported paths (IP / LSP) that can be imported into Simulated Impact Analysis simultaneously	20,000

6.8 NFM-T integration

6.8.1 Scaling guidelines for NFM-T integration

When NFM-P integration with NFM-T is required, a common nspOS component is must be used. Provided that the total number of NEs managed by both NFM-P and NFM-T combined, is less than 50, the nspOS component embedded within NFM-P can be used. If the number of managed NEs exceeds this value, a standalone nspOS component is required. See the NSP Planning Guide for more information.

7 Security

7.1 Securing NFM-P

7.1.1 Overview

Nokia recognizes the importance of deploying important software such as the NFM-P in secure environments and, as such, supports the use of security techniques to enhance the security of the NFM-P.

NFM-P communications is secured using TLS 1.2 by default, SNMPv3 and HTTPS. See the *NSP NFM-P Installation and Upgrade Guide* for configuration information.

NFM-P implements a number of safeguards to ensure protection of private data. Additional information can be found in the Security section of the NSP Deployment and Installation Guide.

Nokia recommends the following steps to achieving NFM-P workstation security:

- Install a clean operating system environment with the minimum required packages documented in the *NSP NFM-P Installation and Upgrade Guide*
- Install the latest Recommended Patch Cluster from Red Hat (not supported on NFM-P provided qcow2 images)
- Harden the RHEL operating system installation based upon the CIS Benchmarks best practices. Reference the NSP System Architecture Guide for the Recommendations and Compliance statements. The supported CIS Benchmark best practices are already implemented on the NFM-P provided qcow2 images.
- If installing RHEL, disable the mDNS Service.
- Implement firewall rules for NFM-P to control access to ports on NFM-P platforms as described in 7.1.5 “Deploying NFM-P with firewalls” (p. 89) . NFM-P systems have no ingress or egress requirements to access the public internet and should be isolated with properly configured firewalls.
- If installing RHEL, enable the RHEL firewall filter rules lists. See 7.4 “Firewall and NAT rules” (p. 104) for more details
- Installation of NFM-P with a secure configuration described in 7.1.3 “Installing the NFM-P components” (p. 88)
- Network Element connection configuration as described in 7.1.4 “NFM-P network element communication” (p. 89)
- Configure NFM-P to run at runlevel 3 as opposed to the default runlevel 5
- Update the supported TLS versions and ciphers to remove older versions, if not required
- Consider using a Certificate Authority signed certificate instead of self-signed certificates
- Use TLS certificates signed with stronger hashing algorithms
- Enable SELinux in permissive/enforcing mode for the components that support it

- Enable Federal Information Processing Standards (FIPS) security. Note that using FIPS and SNMPv3 algorithms larger than SHA1/AES128 will add CPU load and NE response times may be diminished.

7.1.2 Operating system installation for NFM-P workstations

Nokia supports customers applying RHEL, or Windows patches provided by Red Hat, or Microsoft which will include security fixes as well as functional fixes. If a patch is found to be incompatible with NFM-P, the patch may need to be removed until a solution to the incompatibility is provided by Red Hat, Microsoft, or Nokia. Consult the *NSP NFM-P Release 21.x Release Notice* documents for any last minute changes in regards to the recommended RHEL maintenance update and patch levels. Operating system patches of NSP provided qcow2 images must be obtained from the NSP product group.

NFM-P is supported on RHEL installed with the list of required RHEL Packages documented in the *NSP NFM-P Installation and Upgrade Guide*. SELinux is supported in permissive and enforcing mode on the NFM-P server, NFM-P database, and NFM-P statistics auxiliary only. All other NSP components support SELinux in permissive mode only.

Additional efforts to secure the system could impact NFM-P's operation or future upgrades of the product. Customers should perform some level of basic testing to validate additional platform hardening does not impact NFM-P's operation. The NFM-P Product Group makes no commitment to make NFM-P compatible with a customer's hardening requirements.

7.1.3 Installing the NFM-P components

Nokia recommends the following steps when installing the NFM-P components:

- Configure the NFM-P server IP validation during the NFM-P database installation to ensure that only the specified IP address can communicate with the NFM-P database. This is documented in the *NSP NFM-P Installation and Upgrade Guide*.
- Maintain secure communication between the NFM-P server and NFM-P clients (XML-API and UI) as documented in the *NSP NFM-P Installation and Upgrade Guide*. This is enabled by default.

Nokia also recommends the configuration (as documented in the *NSP NFM-P User Guide*) of the following options to secure communication with the NFM-P client UI and the NFM-P client XML API interfaces:

- Password history count
- Password expiry periods
- Client time-outs
- Security statements
- Scope of command and span of control
- Client IP validation

7.1.4 NFM-P network element communication

The following configurations are documented in the *NSP NFM-P User Guide*, and help secure communication between the network elements and NFM-P server installations:

- SNMPv3
- SSH for remote access to the network elements
- SCP/SFTP for secure file transfer
- NETCONF

7.1.5 Deploying NFM-P with firewalls

A firewall can be deployed to protect the NFM-P server from the managed network and to protect the server from the network hosting the NFM-P clients. The diagrams below illustrate this and show the communications services that are required through the firewalls. Installations of NFM-P can make use of the RHEL built in firewall using firewalld. Standalone Firewall products must not be collocated on servers hosting NFM-P components. Only the built-in RHEL firewall used to enable filter rules lists can be collocated with NFM-P components. See 7.4 “Firewall and NAT rules” (p. 104) for more details.

Some NFM-P operations require idle TCP ports to remain open for longer periods of time. Therefore, customers using a firewall that closes idle TCP connections should adjust Operating System TCP keepalives to a value that ensures that the firewall will not close sockets in use by NFM-P.

For some of the network elements described in 5.8 “Network Element specific requirements” (p. 66) there is a requirement for the NFM-P GUI client to communicate directly with the network element using specialized configuration tools.

Figure 7-1 Firewalls and NFM-P standalone deployments

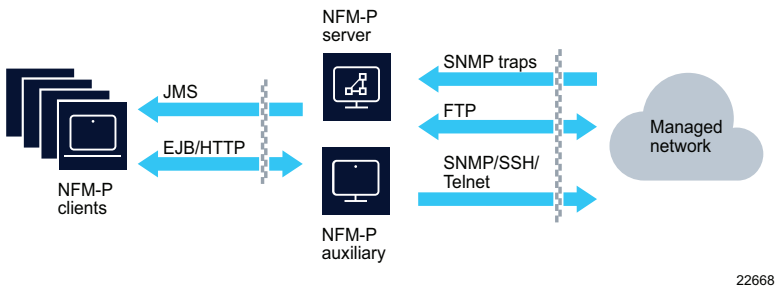

Figure 7-2 Firewalls and NFM-P redundant deployments

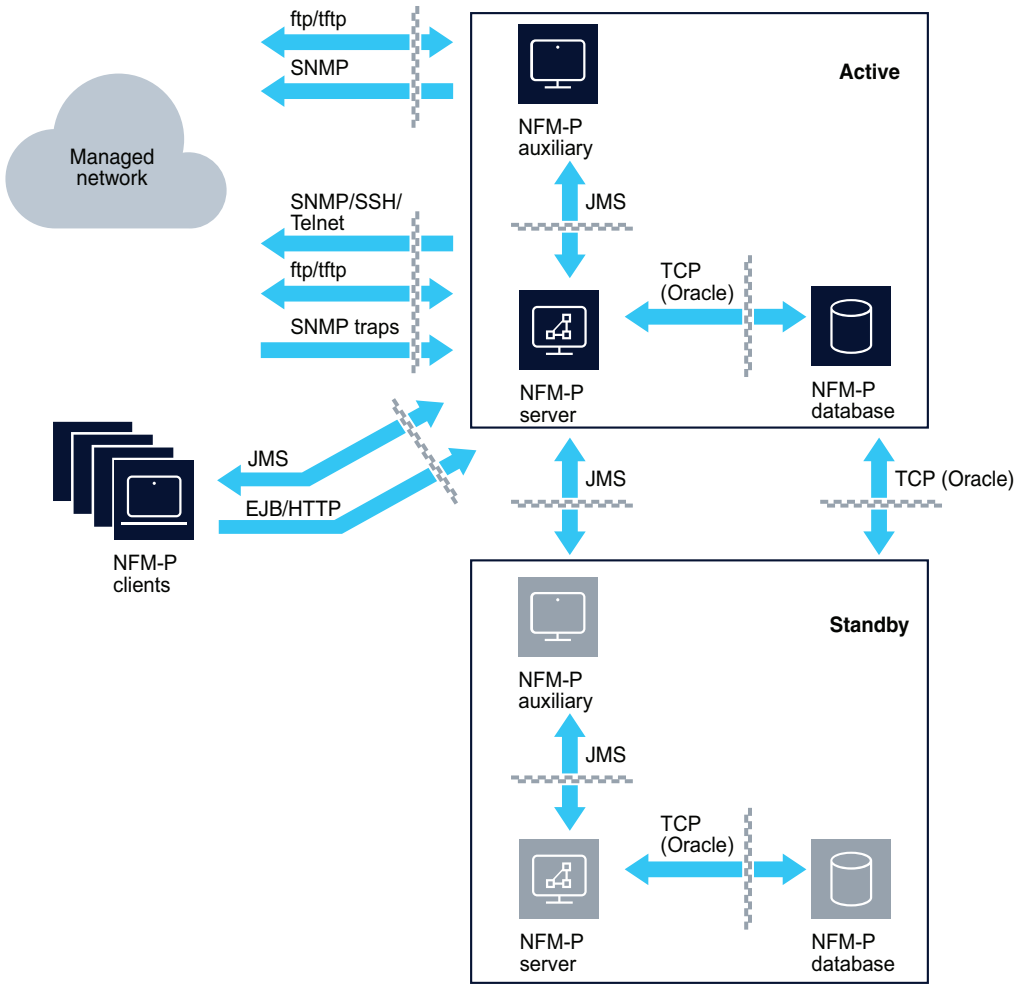

7.2 Port information

7.2.1 Port changes

This version of the NFM-P planning guide contains the following changes, since the 21.11 release:

- Added firewall rules for ports 2390, 5007, 6007 between NFM-P servers

7.2.2 Default ports

The following table describes the listening ports on the various NFM-P applications.

Table 7-1 NFM-P firewall requirements

Default port	Type	Encryption	Description
NFM-P server and NFM-P auxiliary (statistics, call trace, and PCMD)			
N/A	ICMP	N/A	ICMP Ping The active NFM-P server will periodically ping the NFM-P delegate server to ensure reachability.
21 Ports from 1023 - 65536	TCP	None. See SCP and SFTP as secure alternatives	FTP (Passive) This port is used to enable ftp communication from a XML API client to either the NFM-P server or auxiliary. Ftp is used by the XML API client to retrieve logToFile statistics or findToFile results. (See 7.3 “FTP” (p. 103))
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH/SCP/SFTP This port is used for remote access, rsync between NFM-P servers, rsync between the NFM-P databases, and scp/sftp to NFM-P XML API clients.
69	UDP	None. See SFTP for a secure alternative	
80	TCP	None. See port 443 for secure communications.	HTTP This port re-directs to port 443.
162	UDP	Static Encryption When SNMPv3 is configured. Cipher and strength is NE dependant.	SNMP traps By default, this port on the NFM-P server receives SNMP traps from the network elements. This item is specified during the installation of the server and can be changed. (Not required by the NFM-P auxiliary)
443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS interface for the Web Applications through the Launchpad. Also provides a WebDav Server for snapshots and workorders
758	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	nlogin Secure port used for connection to and from the 1830 SMS HSM server
1095	TCP	None.	Internal system communications protocol (JBoss messaging) These ports are used by commands on the NFM-P auxiliary workstation to adjust the NFM-P auxiliary behavior. (Example: adjusting log levels, shutting down the auxiliary server, etc)

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
1097	TCP	None.	Internal system communications protocol (JMS naming/messaging service) Used by the NFM-P client (GUI and XML API) and NFM-P server and NFM-P auxiliary applications to register for JMS notifications and messages. This is used to ensure that the client, server, and auxiliary are aware of system events (i.e.: database changes or alarm notifications, etc)
1099	TCP	None.	Internal system communications protocol (JBoss Naming Service -JNDI) This port is required to ensure the NFM-P GUI, XML API clients, auxiliaries and standby NFM-P server properly initialize with the active NFM-P server. When initially logging into the NFM-P server, NFM-P GUI and XML API clients use this port to find the various services that are available. This port is also used by the NFM-P GUI and XML API clients to register with the NFM-P server to receive notification of network changes.
1664	TCP	None.	LTE NWI3 Corba Service Port This port is required to communicate with OMS for Flexi MR BTS management.
2181	TCP	None. See port 2281 for secure communications.	Java ZooKeeper client connections
2281	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Java ZooKeeper client connections
2390	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	nspdctl
3528	TCP	None.	JBoss jacorb LTE NWI3 Corba Service Port This port is required to communicate with OMS for Flexi MR BTS management.
3529	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss jacorb-ssl LTE NWI3 Corba Service Port This port is required to communicate with OMS for Flexi MR BTS management.
4447	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss messaging port for JMS

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
5007	TCP	None.	Neo4j cluster control
6007	TCP	None.	Neo4j cluster data
6362	TCP	None.	Used by the Web Server This is a local port to the host.
6363	TCP	None.	Neo4j database backup port This is a local port to the host.
6432	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	postgresql communications port
6633	TCP	None.	OpenFlow Used to exchange openflow protocol messages with 7x50 NEs.
7473	TCP	Dynamic Encryption (if TLS is configured) Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Neo4j https web server
7474	TCP	None.	Neo4j web server This is a local port to the host. NFM-P server only
7687	TCP	None.	Neo4j bolt connector
7879	TCP	Dynamic Encryption (if TLS is configured) Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	RPC Layer Used for FM correlation engine to NFM-P server communications. Used for CPROTO communication with the NSP flow collector
7889	TCP	None.	telemetry monitor connection for kpi-engine This is a local port to the host..
8080	TCP	None. See port 8443 for secure communications	HTTP This port provides an HTTP interface for XML API clients to access the NFM-P server.
8085	TCP	None. See port 8444 for secure communications.	HTTP This port provides an HTTP interface for NFM-P client. The NFM-P client uses this port to verify the existence of the server.
8086	TCP	None. See port 8445 for secure communications.	HTTP This port provides an HTTP interface to the WebDav Server for WTA. This port is only required on the CallTrace auxiliary.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
8087	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS Servlet connector used for communication between tomcat and NFM-P server to handle requests with a normal processing time.
8088	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS Webapp services such as correlation.
8089	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS Servlet connector used for communication between tomcat and NFM-P server to handle requests with a long processing time.
8093	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	This port provides an HTTPS interface for Ne3s communication NFM-P server only
8094	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	This port provides an HTTPS interface for Ne3s communication. NFM-P statistics auxiliary only
8195	TCP	None.	Tomcat shutdown port This is a local port to the host.
8196	TCP	None.	Tomcat (app1-tomcat) shutdown port This is a local port to the host.
8197	TCP	None.	Tomcat (app2-tomcat) shutdown port This is a local port to the host.
8400	TCP	None.	HTTP This port re-directs to port 443.
8443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface for XML API clients that wish to use this protocol to access the NFM-P server
8444	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface for the NFM-P client. This is a secure version of port 8085. Used only if the NFM-P client is connecting via TLS.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
8445	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface to the WebDav Server for WTA. This is a secure version of port 8086. Used only if the WTA is connecting via TLS. This port is only required on the Call Trace auxiliary.
8483	TCP	None.	JBoss RMI port for WebServices This is a local port to the host.
8543	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface for the Launchpad, Web Applications, and online help.
8544	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface for Web Applications.
8545	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) interface for RESTCONF.
8617	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	auxdb-agent Communication port from nspdctl
8889	TCP	None.	Notification port used by TAO (CORBA Notification) This is a local port to the host.
9000	TCP	None.	gRPC server used by the ts-model-app in app1-tomcat. This is a local port to the host.
9010	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	This port is used for file synchronization between redundant NFM-P servers and redundant auxiliary collectors (statistics and call trace).
9092	TCP	None. See port 9192 for secure communication.	Kafka server
9192	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Kafka server

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
9400	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port re-directs to port 443.
9443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS HTTPS port for providing access to the HSM server through swagger web interface
9736	TCP	None.	TAO Orb port This is a local port to the host.
9990	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss Management Console Used to access the JBoss management console for the main server process.
9999	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JMX Used to access the JMX console for the main server process.
10090	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss Management Console Used to access the JBoss management console for the JMS server process.
10099	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JMX Used to access the JMX console for the JMS server process.
10190	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss Management Console Used to access the JBoss management console for the auxiliary server process.
10199	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JMX Used to access the JMX console for the auxiliary server process.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
10290	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPs HTTPs interface port between the NFM-P server process and HSM server process
11800	TCP	Static Encryption Encryption provided by AES Cipher Algorithm with 128 bit Cipher Strength.	Internal system communications protocol (JBoss Clustering) This port is required to ensure that redundant NFM-P servers can monitor each other.
12010	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	This port is used for Warm standby Cache Sync communication between redundant NFM-P servers This port is not used on the NFM-P auxiliary.
12300 - 12307	TCP	None.	These ports are used for detecting communication failures between NFM-P server clusters (primary / secondary / auxiliaries)
12800	TCP	Static Encryption Encryption provided by AES Cipher Algorithm with 128 bit Cipher Strength.	Internal system communications protocol (JBoss clustering) During run-time operations, the NFM-P auxiliary uses this port to send and receive information to and from the NFM-P server. The number of required ports depends on the number of NFM-P auxiliary workstations that are installed. Note that NFM-P can be configured to use a different port for this purpose. The procedure is available from Nokia personnel.
29780	UDP	None.	Used to stream UDP PCMD data from SGW and PGW Network Elements Auxiliary PCMD collector only
47100 - 47199	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Session-Manager ignite cache communication spi Only required on the NFM-P server when hosting the nspOS components. Communication to external hosts is not required.
47500 - 47599	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Session-Manager ignite cache discovery spi Only required on the NFM-P server when hosting the nspOS components. Communication to external hosts is not required.
48500 - 48599	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Session-Manager ignite cache communication spi Only required on the NFM-P server when hosting the nspOS components. Communication to external hosts is not required.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
48600 - 48699	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	Session-Manager ignite cache discovery spi Only required on the NFM-P server when hosting the nspOS components. Communication to external hosts is not required.
NSP flow collector			
21 Ports from 1023 - 65536	TCP	None. See SCP and SFTP as secure alternatives	FTP (Passive) This port is used to enable ftp communication between the NSP flow collector and the NFM-P server or dedicated ftp server for retrieving IPDR files.
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH/SCP/SFTP This port is used to enable SSH (SFTP/SCP) communication between the NSP flow collector and the NFM-P server or dedicated ftp server for retrieving IPDR files.
2205	UDP	None.	CGNAT / IPFIX cflowd records from 7750 routers to NSP flow collector
4739	UDP	None.	cflowd records from 7750 routers to NSP flow collector
7899	TCP	None.	CPROTO
8080	TCP	None. See port 8443 for secure communications.	HTTP This port provides an HTTP Web User interface for the NSP flow collector
8083	TCP	None.	JBoss Socket for dynamic class and resource loading.
8443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTP Web User interface for the NSP flow collector This is a secure version of port 8080.
9443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) NSP flow collector management interface. This is a secure version of port 9990. Used only if the NSP flow collector is TLS secured.
9990	TCP	None. See port 9443 for secure communications.	HTTP This port provides an HTTP NSP flow collector management interface. This is a local port to the host.
9999	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JMX Used to access the JMX console. This is a local port to the host.
44444	TCP	None	RMI server port

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
NSP flow collector controller			
21 Ports from 1023 - 65536	TCP	None. See SCP and SFTP as secure alternatives	FTP (Passive) This port is used to enable ftp communication between the NSP flow collector controller and the NFM-P server or dedicated ftp server for retrieving IPDR files.
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH/SCP/SFTP This port is used to enable SSH (SFTP/SCP) communication between the NSP flow collector controller and the NFM-P server or dedicated ftp server for retrieving IPDR files.
1090	TCP	None.	JBoss RMI/JRMP socket for connecting to the JMX MBeanServer. Used for NFM-P server to NSP flow collector controller communication.
1098	TCP	None.	JBoss Socket Naming service used to receive RMI request from client proxies. Used for NFM-P server to NSP flow collector controller communication.
1099	TCP	None.	JBoss The listening socket for the Naming service. Used for Jboss communication between NFM-P and NSP flow collector controller.
4444	TCP	None.	JBoss Socket for the legacy RMI/JRMP invoker. Used for Jboss communication between NFM-P to NSP flow collector controller.
4445	TCP	None.	JBoss Socket for the legacy Pooled invoker. Used for Jboss communication between NFM-P to NSP flow collector controller.
4446	TCP	None.	JBoss Socket for the JBoss Remoting Connected used by Unified Invoker. Used for Jboss communication between NFM-P to NSP flow collector controller.
4447	TCP	None.	JBoss Socket for JBoss Remoting Connections. This is a local port to the host.
4457	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	JBoss Socket for JBoss Messaging 1.x
7879	TCP	None.	CPROTO
8080	TCP	None. See port 8443 for secure communications.	HTTP This port provides an HTTP Web User interface for the NSP flow collector controller.
8083	TCP	None.	JBoss Socket for dynamic class and resource loading.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
8443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTP Web User interface for the NSP flow collector controller. This is a secure version of port 8080.
9443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides an HTTPS (secure HTTP) NSP flow collector controller management interface. This is a secure version of port 9990. Used only if the NSP flow collector controller is TLS secured.
9990	TCP	None. See port 9443 for secure communications.	HTTP This port provides an HTTP NSP flow collector controller management interface. This is a local port to the host.
22222	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SFTP SFTP connection from NSP flow collector.
44444	TCP	None	RMI server port
NSP analytics server			
8080	TCP	None. See port 8443 for secure communications.	HTTP This port provides an HTTP Web User interface for the NSP analytics server. It's used by the NFM-P server and web based clients for HTTP requests.
8443	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS This port provides a secure HTTP Web User interface for the NSP analytics server. It's used by the NFM-P server and web based clients for HTTPS requests This is a secure version of port 8080.
10990	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	HTTPS Used to access the JMX console for the analytics process.
NFM-P auxiliary database			
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH / SFTP Vertica Administration Tools. Inter-node and inter-cluster communication
4803	TCP	None.	Spread Client connections Inter-node communication only.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
4803	UDP	None.	Spread Daemon to Daemon connections Inter-node communication only.
4804	UDP	None.	Spread Daemon to Daemon connections Inter-node communication only.
5433	TCP	None.	JDBC Client communication port (NFM-P server, statistics auxiliary, flow collector, analytics server)
5433	UDP	None.	Vertica Vertica spread monitoring Inter-node communication only.
5434	TCP	None.	Vertica Intra and inter cluster communication Inter-node communication only.
6543	TCP	None.	Spread Monitor to Daemon connections Inter-node communication only.
7299–7309	TCP	None.	RMI auxiliary database proxy port.
50000	TCP	None.	Rsync Inter-node and inter-cluster communication
32768-60999	TCP	None.	Vertica - Zygote Inter-node communication only
32768-60999	UDP	None.	Vertica - Spread Inter-node communication only
Managed devices			
21 Ports from 1023 - 65536	TCP	None.	FTP (Passive) This port is used to enable ftp communication between the NFM-P server and the managed routers. Ftp occurs to transfer information from the routers to the NFM-P server such as accounting statistics. See 7.3 “FTP” (p. 103) for a more detailed description of ftp requirements.
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH / SFTP This port used by clients to request a SSH session to a managed router.
23	TCP	None.	Telnet This port used by clients to request a telnet session to a managed router.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
80	TCP	None.	HTTP This port is required for the NFM-P client to communicate with the network element Web GUIs. See 5.8 "Network Element specific requirements" (p. 66) for the network elements that require this port.
161	UDP	Static Encryption When SNMPv3 is configured. Cipher and strength is NE dependant.	SNMP By default, NFM-P server sends SNMP messages, such as configuration requests and service deployments, to this port on the network elements.
830	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSHv2 for CMM This port is used by the CMM network elements for NetConf management.
1491	TCP	Static Encryption When SNMPv3 is configured. Cipher and strength is NE dependant.	SNMP Streaming Used for TCP Streaming during NE discovery and resync. Only applicable to 7950 XRS, 7750 SR, 7450 ESS, and 7710 SPR, 11.0R5+.
5001	TCP	None.	Proprietary Java socket connection This port is used by CPAM to communicate with the 7701 CPAA to obtain control plane information.
5010	UDP	None.	Trap Trap port used by 9500 MPR / Wavence SM devices to send traps to NFM-P clients running the NetO manager.
11500	TCP	None.	Equipment View Used while managing 9500 MPR / Wavence SM(MSS-1C, MPR-e, MSS-8) NEs using the Equipment View function as part of NetO
N/A	ICMP	N/A	ICMP Only used if the Ping Policy is enabled as part of network element mediation.
NFM-P database			
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	SSH This port is used by NFM-P for an optional rsync feature between NFM-P databases
1523	TCP	Static Encryption Encryption provided by RC4 Cipher Algorithm with 128 bit Cipher Strength.	Oracle SQL*Net Listener This port is used by the NFM-P server to connect to and communicate with the NFM-P database. When there are redundant databases, this port is also used by Oracle DataGuard to keep the databases in sync. The data on this port is encrypted.
9002	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	NFM-P database Proxy This port is used by the NFM-P server to monitor disk usage on a remote NFM-P database. When there are redundant databases, it is also allows the NFM-P server to initiate database switchovers and failovers.

Table 7-1 NFM-P firewall requirements (continued)

Default port	Type	Encryption	Description
9003	TCP	None.	Database file transfer Port This port is used by the NFM-P database workstations in a redundant workstation configuration. This port allows database transfers between the primary and standby databases. For example: when the standby database is re-instantiated, or when the standby database is installed for the first time.
NFM-P client and client delegate			
20	TCP	None.	FTP Active FTP port for 9500 MPR / Wavence SM software download from NEtO.
21 Ports from 1023 - 65535	TCP	None.	FTP 9500 MPR / Wavence SM software download from NEtO.
22	TCP	Dynamic Encryption Cipher Suite and strength as per RFC 4253	sFTP 9500 MPR / Wavence SM software download from NEtO
162	UDP	None.	Trap Trap port used by 9500 MPR / Wavence SM (MPR-e, MSS-8) devices to send traps to NFM-P clients running the NetO manager.
5010	UDP	None.	Trap Trap port used by 9500 MPR / Wavence SM devices to send traps to NFM-P clients running the NetO manager.

7.3 FTP

7.3.1 FTP between the NFM-P server and NFM-P auxiliary statistics collector and the managed network

NFM-P server and NFM-P auxiliary statistics collector may use FTP for several purposes.

The NFM-P server may use FTP, if configured, to receive backup images of managed devices, to send new software images to the managed devices and to receive accounting statistics from the managed devices.

If an NFM-P auxiliary statistics collector workstation is installed, FTP will be used, if configured, to retrieve accounting statistics from managed devices.

If STM Accounting tests are being executed, the NFM-P server will retrieve the test results from the managed devices by FTP, if configured.

The FTP communication is configured as an extended *passive* FTP connection, with the managed devices serving as the FTP servers and the NFM-P server and NFM-P auxiliary acting as the FTP client.

Extended passive FTP connections use dynamically-allocated ports on both sides of the communication channel, and are ephemeral in nature. As such, the data sent from the managed

devices will be sent from a port in the range of 1024-65536. This data will be sent to the NFM-P server on a port in the range of 1024-65536. Support for EPSV/EPRT ftp commands (commands that can replace PASV/PORT commands) must be enabled for connections to the 7x50 family of routers.

7.4 Firewall and NAT rules

7.4.1 Overview

Firewall rules are applied to the incoming network interface traffic of the NFM-P workstations. As a rule, firewall rules are not applied to the outgoing network interface traffic.

For NFM-P installations using RHEL as the Operating System, the RHEL supplied firewall can be used to filter network traffic using filter rules lists. Only experienced system administrators with extensive knowledge of the RHEL firewall should attempt to implement the filter rules lists provided with each NFM-P component. All others should disable the RHEL firewall.

The installation of each NFM-P component will include the filter rules lists to be applied for successful communication between different NFM-P components, XML API clients, and Network Elements. The table below defines the location

Table 7-2 Sample firewalld filter rules lists file locations

Component	Protocol	File location
NFM-P server	IPv4/IPv6	/opt/nsp/nfmp/server/nms/sample/firewall/
NFM-P database	IPv4/IPv6	/opt/nsp/nfmp/db/install/sample/firewall/
NFM-P Statistics/call trace/PCMD auxiliary	IPv4/IPv6	/opt/nsp/nfmp/auxserver/nms/sample/firewall/
NSP flow collector controller	IPv4/IPv6	/opt/nsp/flow/fcc/sample/firewalld/
NSP flow collector	IPv4/IPv6	/opt/nsp/flow/fc/sample/firewalld/
NFM-P auxiliary database	IPv4	/opt/nsp/nfmp/auxdb/install/config/sample/firewall/
NFM-P client	IPv4/IPv6	<base client install dir>/nms/sample/firewall/
NFM-P client delegate	IPv4/IPv6	<base client install dir>/nms/sample/firewall/

It is imperative that all rules are considered completely for the NFM-P systems to inter-operate correctly. The following tables will define the rules to be applied to each NFM-P workstation. Within the section there will be a number of conditions that indicate whether or not that particular table needs to be applied.

See 8.2 “Network Address Translation” (p. 132) for supported NAT configurations.

7.4.2 NFM-P server firewall and NAT rules

When there is a firewall at the NFM-P server(s) interface that reaches the managed network (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following firewall rules need to be applied.

Table 7-3 SNMP firewall rules for traffic between the NFM-P server(s) and the managed network

Protocol	From port	On	To port	On	Notes
UDP	Any	Managed network	162	Server(s)	SNMP trap initiated from the NE
UDP	>15000	Server(s)	161	Managed network	SNMP request
UDP	161	Managed network	>15000	Server(s)	SNMP response
TCP	>15000	Server(s)	1491	Managed network	SNMP TCP Streaming
TCP	1491	Managed network	>15000	Server(s)	SNMP TCP Streaming

Note: Due to the size of SNMP packets, IP fragmentation may occur in the network. Ensure the firewall will allow fragmented packets to reach the server(s).

Table 7-4 Telnet / FTP firewall rules for traffic between the NFM-P server(s) and the managed network

Protocol	From port	On	To port	On	Notes
TCP	>15000	Server(s)	23	Managed network	Telnet request
TCP	23	Managed network	>15000	Server(s)	Telnet response
TCP	Any	Server(s)	21	Managed network	FTP requests (example: STM, Accounting statistics, NE backups)
TCP	21	Managed network	Any	Server(s)	FTP responses
TCP	> 1023	Managed network	> 1023	Server(s)	Passive FTP ports for data transfer

Table 7-5 SSH / SFTP / SCP firewall rules for traffic between the NFM-P server(s) and the managed network

Protocol	From port	On	To port	On	Notes
TCP	Any	Server(s)	22	Managed network	NFM-P SSH request
TCP	22	Managed network	Any	Server(s)	NFM-P SSH response
TCP	>15000	Server(s)	830	Managed network	SSHv2 request for MME
TCP	830	Managed network	>15000	Server(s)	SSHv2 response for MME

Table 7-6 Other firewall rules for traffic between the NFM-P server(s) and the managed network

Protocol	From port	On	To port	On	Notes
ICMP	N/A	Managed network	N/A	NFM-P server(s)	Only used if Ping Policy is enabled.
TCP	>15000	NFM-P server(s)	5001	7701 CPAA Elements	–
TCP	>15000	Managed network	6633	NFM-P server(s)	Openflow data

Table 7-6 Other firewall rules for traffic between the NFM-P server(s) and the managed network (continued)

Protocol	From port	On	To port	On	Notes
TCP	>15000	NFM-P server(s)	57400	Managed network	Telemetry data

Table 7-7 Firewall rules for traffic between the NFM-P server(s) and OMS for Flexi MR BTS management

Protocol	From port	On	To port	On	Notes
TCP	Any	OMS	1664	NFM-P server(s)	Corba
TCP	Any	OMS	3528	NFM-P server(s)	Corba
TDP	Any	OMS	3529	NFM-P server(s)	Corba

Table 7-8 Firewall rules for traffic between the NFM-P server(s) and the LTE BTS

Protocol	From port	On	To port	On	Notes
TCP	Any	LTE BTS	8093	NFM-P server(s)	NE3S

Table 7-9 Firewall rules for traffic between the NFM-P server(s) and the 1830 SMS HSM

Protocol	From port	On	To port	On	Notes
TCP	758	NFM-P server	5552	1830 SMS HSM	Two way communication

Table 7-10 Firewall rules for remote user authentication

Protocol	From port	On	To port	On	Notes
TCP/UDP	Any	NFM-P server	49	TACACS server	For TACACS authentication
TCP/UDP	Any	NFM-P server	389	LDAP server	For LDAP authentication
TCP/UDP	Any	NFM-P server	636	LDAP server	For LDAP authentication (TLS)
UDP	Any	NFM-P server	1812	RADIUS server	For RADIUS authentication

When there is a firewall at the interface that reaches the NFM-P client(s) (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) the following rules need to be applied.

Table 7-11 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P client(s) (GUI/XML API/JMS/Web/Kafka)

Protocol	From port	On	To port	On	Notes
TCP	Any	XML API client	21	Server(s)	If FTP is required
TCP	Any	XML API client	22	Server(s)	If SFTP/SCP is required

Table 7-11 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P client(s) (GUI/XML API/JMS/Web/Kafka) (continued)

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P GUI client	443	Server(s)	HTTPS
TCP	> 1023	XML API client	> 1023	Server(s)	If (passive) FTP is required
TCP	Any	XML API/NFM-P GUI client	1097	Server(s)	JMS
TCP	Any	XML API/NFM-P GUI client	1099	Server(s)	JNDI
TCP	Any	XML API/NFM-P GUI client	4447	Server(s)	JMS
UDP	Any	NFM-P GUI client	6100-6119	Server(s)	NEM Proxy
TCP	Any	XML API client	8080	Server(s)	HTTP
TCP	Any	NFM-P GUI client	8085	Server(s)	HTTP
TCP	Any	NFM-P GUI client	8087	Server(s)	HTTP(S)
TCP	Any	NFM-P GUI client	8088	Server(s)	HTTP(S)
TCP	Any	NFM-P GUI client	8089	Server(s)	HTTP(S)
TCP	Any	XML API client	8443	Server(s)	HTTPS
TCP	Any	NFM-P GUI client	8444	Server(s)	HTTPS
TCP	Any	NFM-P GUI / Web client	8543	Server(s)	HTTPS
TCP	Any	NFM-P GUI / Web client	8544	Server(s)	HTTPS
TCP	Any	RESTCONF client	8545	Server(s)	HTTPS
TCP	Any	kafka client	9092	Server(s)	kafka
TCP	Any	kafka client	9192	Server(s)	kafka Secure
TCP	Any	Web client	9443	Server(s)	HSM

When there is a firewall configured, and there are redundant NFM-P auxiliary workstation(s), the following rules need to be applied to the appropriate interface.

Table 7-12 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P auxiliary statistics / call trace / PCMD collector(s)

Protocol	From port	On	To port	On
TCP	Any	Auxiliary server(s)	1097	Server(s)
TCP	Any	Auxiliary server(s)	1099	Server(s)
TCP	Any	Auxiliary server(s) (statistics)	2181	Server(s)

Table 7-12 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P auxiliary statistics / call trace / PCMD collector(s) (continued)

Protocol	From port	On	To port	On
TCP	Any	Auxiliary server(s) (statistics)	2281	Server(s)
TCP	Any	Auxiliary server(s)	4447	Server(s)

Table 7-13 Firewall rules for traffic coming into the NFM-P server(s) from the NSP flow collector controller(s)

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector controller	21	Server(s)
TCP	Any	NSP flow collector controller	22	Server(s)
TCP	>1023	NSP flow collector controller	>1023	Server(s)
TCP	Any	NSP flow collector controller	1099	Server(s)
TCP	Any	NSP flow collector controller	2181	Server(s)
TCP	Any	NSP flow collector controller	2281	Server(s)
TCP	Any	NSP flow collector controller	7879	Server(s)
TCP	Any	NSP flow collector controller	8080/8443	Server(s)

Table 7-14 Firewall rules for traffic coming into the NFM-P server(s) from the NSP flow collector(s)

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector controller	2181	Server(s)
TCP	Any	NSP flow collector controller	2281	Server(s)

When a firewall and NAT are configured to the NFM-P server at the NFM-P client interface (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) the following rules need to be applied to allow the XML API clients to retrieve the logToFile accounting statistics information. Services require the use of public addresses.

Table 7-15 Additional firewall rules required to allow services on the NFM-P client(s) to communicate with the NFM-P server if NAT is used

Protocol	From port	On	To port	On
TCP	Any	Server Public Address	21	Server Private Address
TCP	21	Server Public Address	Any	Server Private Address
TCP	> 1023	Server Public Address	> 1023	Server Private Address

When there is a firewall at the interface that reaches the NFM-P management network (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply.

Table 7-16 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P database server(s)

Protocol	From port	On	To port	On
TCP	1523	Database server(s)	Any	Server(s)
TCP	9002	Database server(s)	Any	Server(s)

When there is a firewall at the NFM-P management interface (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) and NFM-P server redundancy is configured, then the following rules need to be applied. Configuration needs to be in both directions to handle an activity switch.

Table 7-17 Firewall rules for setups with redundant NFM-P servers

Protocol	From port	On	To port	On
TCP	Any	Servers	22	Servers
TCP	22	Servers	Any	Servers
TCP	Any	Server	1099	Server
TCP	1099	Server	Any	Server
TCP	Any	Server	2390	Server
TCP	2390	Server	Any	Server
TCP	Any	Server	5007	Server
TCP	5007	Server	Any	Server
TCP	Any	Server	6007	Server
TCP	6007	Server	Any	Server
TCP	>15000	Server	6432	Server
TCP	6432	Server	>15000	Server
TCP	>15000	Server	7879	Server
TCP	7879	Server	>15000	Server

Table 7-17 Firewall rules for setups with redundant NFM-P servers (continued)

Protocol	From port	On	To port	On
TCP	>15000	Servers	8087	Servers
TCP	8087	Servers	>15000	Servers
TCP	>15000	Servers	8543	Servers
TCP	8543	Servers	>15000	Servers
TCP	>15000	Servers	9010	Servers
TCP	9010	Servers	>15000	Servers
TCP	>15000	Servers	11800	Servers
TCP	11800	Servers	>15000	Servers
TCP	>15000	Servers	12010	Servers
TCP	12010	Servers	>15000	Servers
TCP	>15000	Servers	12300-12307	Servers
TCP	12300-12307	Servers	>15000	Servers

When there is a firewall at the NFM-P management interface (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) and NFM-P auxiliary statistics / call trace / PCMD collectors are configured, then the following rules need to be applied:

Table 7-18 Firewall rules for traffic coming into the NFM-P server(s) from the NFM-P auxiliary statistics / call trace server(s)

Protocol	From port	On	To port	On
TCP	Any	Auxiliary server(s)	12300-12307	Server(s)
TCP	12300-12307	Auxiliary server(s)	Any	Server(s)
TCP	Any	Auxiliary server(s)	12800	Server(s)
TCP	12800	Auxiliary server(s)	Any	Server(s)

If NFM-P is not deployed with NSP, the following rules need to be applied to the NFM-P server if there is a firewall on the NFM-P management interface (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127))

Table 7-19 Firewall rules for inter-process communication on the NFM-P server(s)

Protocol	From port	On	To port	On
TCP	>15000	NFM-P server(s)	443	NFM-P server(s)
TCP	443	NFM-P server(s))	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	2181	NFM-P server(s)
TCP	2181	NFM-P server(s)	>15000	NFM-P server(s)

Table 7-19 Firewall rules for inter-process communication on the NFM-P server(s) (continued)

Protocol	From port	On	To port	On
TCP	>15000	NFM-P server(s)	2281	NFM-P server(s)
TCP	2281	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	2390	NFM-P server(s)
TCP	2390	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	5007	NFM-P server(s)
TCP	5007	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	6007	NFM-P server(s)
TCP	6007	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	6432	NFM-P server(s)
TCP	6432	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	7473	NFM-P server(s)
TCP	7473	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	7687	NFM-P server(s)
TCP	7687	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	8087	NFM-P server(s)
TCP	8087	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	8617	NFM-P server(s)
TCP	8617	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	9092	NFM-P server(s)
TCP	9092	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	9192	NFM-P server(s)
TCP	9192	NFM-P server(s)	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	10290	NFM-P server(s)
TCP	10290	NFM-P server(s)	>15000	NFM-P server(s)

If NFM-P is deployed with NSP, the following rules need to be applied to the NFM-P server if there is a firewall on the NFM-P management interface (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127))

Table 7-20 Firewall rules for communication between the NFM-P server(s) and NSP

Protocol	From port	On	To port	On
TCP	>15000	NFM-P server(s)	443	NSP
TCP	443	NSP	>15000	NFM-P server(s)

Table 7-20 Firewall rules for communication between the NFM-P server(s) and NSP (continued)

Protocol	From port	On	To port	On
TCP	>15000	NFM-P server(s)	2181	NSP
TCP	2181	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	2281	NSP
TCP	2281	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	2390	NSP
TCP	2390	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	6432	NSP
TCP	6432	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	7473	NSP
TCP	7473	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	7687	NSP
TCP	7687	NSP	>15000	NFM-P server(s)
TCP	>15000	NSP	7879	NFM-P server(s)
TCP	7879	NFM-P server(s)	>15000	NSP
TCP	>15000	NFM-P server(s)	8087	NSP
TCP	8087	NSP	>15000	NFM-P server(s)
TCP	>15000	NSP	8089	NFM-P server(s)
TCP	8089	NFM-P server(s)	>15000	NSP
TCP	>15000	NFM-P server(s)	9092	NSP
TCP	9092	NSP	>15000	NFM-P server(s)
TCP	>15000	NFM-P server(s)	9192	NSP
TCP	9192	NSP	>15000	NFM-P server(s)

Table 7-21 Firewall rules for communication between the NFM-P statistics auxiliary and NSP

Protocol	From port	On	To port	On
TCP	>15000	NFM-P statistics auxiliary	2181	NSP
TCP	2181	NSP	>15000	NFM-P statistics auxiliary
TCP	>15000	NFM-P statistics auxiliary	2281	NSP
TCP	2281	NSP	>15000	NFM-P statistics auxiliary

If NFM-P is integrated with NFM-T, without NSP, the following rules need to be applied to the

NFM-P server if there is a firewall on the NFM-P management interface (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127))

Table 7-22 Firewall rules for communication between the NFM-P server(s) and NFM-T

Protocol	From port	On	To port	On
TCP	80	NFM-T presentation server	>32768	NFM-P server(s)
TCP	>32768	NFM-P server(s)	80	NFM-T presentation server
TCP	443	NFM-P server(s)	>15000	NFM-T
TCP	>15000	NFM-T	443	NFM-P server(s)
TCP	2181	NFM-P server(s)	>15000	NFM-T
TCP	>15000	NFM-T	2181	NFM-P server(s)
TCP	2281	NFM-P server(s)	>15000	NFM-T
TCP	>15000	NFM-T	2281	NFM-P server(s)
TCP	>32768	NFM-P server(s)	8443	NFM-T OTNE server
TCP	8443	NFM-T OTNE server	>32768	NFM-P server(s)
TCP	9092	NFM-P server(s)	>15000	NFM-T
TCP	>15000	NFM-T	9092	NFM-P server(s)
TCP	9192	NFM-P server(s)	>15000	NFM-T
TCP	>15000	NFM-T	9192	NFM-P server(s)

7.4.3 NFM-P database firewall and NAT rules

When there is a firewall at the interface that reaches the NFM-P management network (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply.

Table 7-23 Firewall rules for traffic coming into the NFM-P database server(s) from the NFM-P server(s), NFM-P auxiliary statistics / call trace / PCMD collector(s), and NSP analytics server

Protocol	From port	On	To port	On
TCP	Any	Server(s), auxiliary server(s), & analytics server	1523	Database server(s)
TCP	Any	Server(s) & auxiliary server(s)	9002	Database server(s)
TCP	>15000	Server(s) & auxiliary server(s)	9003	Database server(s)

When there is a firewall at the interface that reaches the NFM-P management network (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces”

(p. 127)) and redundancy is configured, the following rules apply. Configuration needs to be in both directions to handle an activity switch.

Table 7-24 Firewall rules for traffic between the NFM-P database servers (redundant only)

Protocol	From port	On	To port	On
TCP	Any	Database servers	22	Database servers
TCP	22	Database servers	Any	Database servers
TCP	Any	Database servers	1523	Database servers
TCP	1523	Database servers	>15000	Database servers
TCP	9002	Database servers	9002	Database servers
TCP	9003	Database servers	9003	Database servers

7.4.4 NFM-P auxiliary server and NSP flow collector firewall and NAT rules

When there is a firewall at the interface that reaches the managed network (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply.

Table 7-25 SNMP firewall rules for traffic coming into the NFM-P auxiliary statistics collector server(s) from the managed network

Protocol	From port	On	To port	On	Notes
UDP	>32768	Auxiliary server(s)	161	Managed network	SNMP request
UDP	161	Managed network	>32768	Auxiliary server(s)	SNMP response

Note: Due to the size of SNMP packets, IP fragmentation may occur in the network. Ensure the firewall will allow fragmented packets to reach the server(s).

Table 7-26 SSH / Telnet firewall rules for traffic coming into the NFM-P auxiliary statistics collector server(s) from the managed network

Protocol	From port	On	To port	On	Notes
TCP	>32768	Auxiliary server(s)	22-23	Managed network	SSH/SCP/Telnet request
TCP	22-23	Managed network	> 32768	Auxiliary server(s)	SSH/SCP/Telnet response

Table 7-27 FTP firewall rules for traffic coming into the NFM-P auxiliary statistics collector(s) from the managed network

Protocol	From port	On	To port	On	Notes
TCP	Any	Auxiliary server(s)	21	Managed network	FTP requests (example: STM, accounting statistics, NE backups)
TCP	21	Managed network	Any	Auxiliary server(s)	FTP responses
TCP	> 1023	Managed network	> 1023	Auxiliary server(s)	Passive FTP ports for data transfer (See 7.3 “FTP” (p. 103))

Table 7-28 Firewall rules for traffic coming into the NFM-P auxiliary statistics collector(s) from the LTE BTS

Protocol	From port	On	To port	On	Notes
TCP	Any	LTE BTS	8094	NFM-P auxiliary server(s)	NE3S

Note: FTP access is only required for the NFM-P auxiliary statistics collector.

Table 7-29 SNMP firewall rules for traffic coming into the NFM-P auxiliary call trace collector(s) from the managed network

Protocol	From port	On	To port	On	Notes
UDP	>32768	Auxiliary server(s)	161	Managed network	SNMP request
UDP	161	Managed network	> 32768	Auxiliary server(s)	SNMP response

Note: Due to the size of SNMP packets, IP fragmentation may occur in the network. Ensure the firewall will allow fragmented packets to reach the server(s).

Table 7-30 Firewall rules for traffic coming into the NFM-P auxiliary PCMD collector(s) from the managed network

Protocol	From port	On	To port	On	Notes
UDP	Any	Managed network	29780	Auxiliary server(s)	PCMD records from SGW / PGW to NFM-P PCMD auxiliary collector

Table 7-31 Firewall rules for traffic coming into the NSP flow collector(s) from the managed network

Protocol	From port	On	To port	On	Notes
UDP	Any	Managed network	2205	NSP flow collector	CGNAT / IPFIX cflowd records

Table 7-31 Firewall rules for traffic coming into the NSP flow collector(s) from the managed network (continued)

Protocol	From port	On	To port	On	Notes
UDP	Any	Managed network	4739	NSP flow collector	cflowd records from 7750 routers to NSP flow collector

When there is a firewall at the interface that reaches the NFM-P client(s) (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply for FTP access to the NFM-P auxiliary by the XML-API client.

Table 7-32 Firewall rules for XML API client communication to the NFM-P auxiliary collector(s)

Protocol	From port	On	To port	On	Notes
TCP	Any	XML API client	21/22	auxiliary collector(s)	(S)FTP requests (logToFile statistics, and call trace information)
TCP	21/22	XML API client	Any	auxiliary collector(s)	(S)FTP responses
TCP	> 1023	XML API client	Any	auxiliary collector(s)	Passive (S)FTP ports for data transfer (See 7.3 “FTP” (p. 103))
Only for NFM-P auxiliary call trace collectors					
TCP	Any	XML API client	8086	auxiliary collector(s)	HTTP interface for WebDAV for WTA
TCP	Any	XML API client	8445	auxiliary collector(s)	HTTPS interface for WebDAV for WTA

Table 7-33 FTP/SFTP firewall rules for the NSP flow collector(s)

Protocol	From port	On	To port	On	Notes
TCP	Any	NSP flow collector(s)	21/22	Target file server	(S)FTP requests
TCP	21/22	Target file server	Any	NSP flow collector(s)	(S)FTP responses
TCP	> 1023	Target file server	Any	NSP flow collector(s)	Passive (S)FTP ports for data transfer (See 7.3 “FTP” (p. 103))

Table 7-34 FTP/SFTP firewall rules for the NSP flow collector controllers(s)

Protocol	From port	On	To port	On	Notes
TCP	Any	NSP flow collector(s)	21/22	NFM-P server	(S)FTP requests
TCP	21/22	NFM-P server	Any	NSP flow collector(s)	(S)FTP responses
TCP	> 1023	NFM-P server	Any	NSP flow collector(s)	Passive (S)FTP ports for data transfer (See 7.3 “FTP” (p. 103))

Table 7-34 FTP/SFTP firewall rules for the NSP flow collector controllers(s) (continued)

Protocol	From port	On	To port	On	Notes
TCP	Any	NSP flow collector	22222	NSP flow collector controller	SFTP requests
TCP	22222	NSP flow collector controller	Any	NSP flow collector	SFTP responses

When there is a firewall at the interface that communicates with the NFM-P servers, the following rules apply for inter process communication.

Table 7-35 Firewall rules for inter-process communication on the NFM-P auxiliary statistics / call trace collector(s)

Protocol	From port	On	To port	On
TCP	Any	Auxiliary server(s)	1095	Auxiliary server(s)
TCP	Any	Auxiliary server(s)	12300-12307	Auxiliary server(s)
TCP	12300-12307	Auxiliary server(s)	Any	Auxiliary server(s)
TCP	Any	Auxiliary server(s)	12800	Auxiliary server(s)
TCP	12800	Auxiliary server(s)	Any	Auxiliary server(s)

Table 7-36 Firewall rules for inter-process communication on the NSP flow collector controller(s)

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector(s)	1090	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	1098	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	1099	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	4444	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	4445	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	4446	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	4457	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	8083	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	9443	NSP flow collector(s)
TCP	Any	NSP flow collector(s)	44444	NSP flow collector(s)

Table 7-37 Firewall rules for inter-process communication on the NSP flow collector(s)

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector(s)	44444	NSP flow collector(s)

When there is a firewall at the interface that communicates with the NFM-P servers, the following rules apply.

Table 7-38 Firewall rules for traffic coming into the NFM-P auxiliary statistics / call trace / PCMD collector(s) from the NFM-P server(s)

Protocol	From port	On	To port	On
TCP	1097	NFM-P server(s)	Any	Auxiliary server(s)
TCP	1099	NFM-P server(s)	Any	Auxiliary server(s)
TCP	4447	NFM-P server(s)	Any	Auxiliary server(s)

Table 7-39 Firewall rules for traffic between the NSP flow collector controller(s) from the NFM-P server(s) or NSP

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector controller(s)	2181	NFM-P server(s) / NSP
TCP	Any	NSP flow collector controller(s)	2281	NFM-P server(s) / NSP
TCP	Any	NFM-P server(s) / NSP	7879	NSP flow collector controller(s)

Table 7-40 Firewall rules for traffic between the NSP flow collector(s) from the NFM-P server(s) or NSP

Protocol	From port	On	To port	On
TCP	Any	NSP flow collector(s)	2181	NFM-P server(s) / NSP
TCP	Any	NSP flow collector(s)	2281	NFM-P server(s) / NSP
TCP	Any	NFM-P server(s) / NSP	7899	NSP flow collector(s)

When there is a firewall at the interface that reaches the NFM-P client(s) (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) and NAT is used on the NFM-P auxiliary server(s), the following rules apply to allow the XML API clients to collect the logToFile accounting statistics files. Services require the use of public addresses.

Table 7-41 Additional firewall rules required to allow services on the NFM-P client(s) to communicate with the NFM-P auxiliary(s) if NAT is used on the auxiliary server(s)

Protocol	From port	On	To port	On
TCP	Any	Auxiliary server public address	21	Auxiliary server private address
TCP	21	Auxiliary server public address	Any	Auxiliary server private address
TCP	> 1023	Auxiliary server public address	> 1023	Auxiliary server private address

When there is a firewall at the interface that reaches the NFM-P management network (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply.

Table 7-42 Firewall rules for traffic coming into the NFM-P auxiliary collector(s) from the NFM-P database(s)

Protocol	From port	On	To port	On
TCP	1523	NFM-P database	Any	NFM-P auxiliary collector(s)
TCP	9002	NFM-P database	Any	NFM-P auxiliary collector(s)

When there is a firewall at the interface that reaches the NFM-P management network (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)), the following rules apply.

Table 7-43 Firewall rules for traffic coming into the NFM-P auxiliary collector(s) from the NFM-P server(s)

Protocol	From port	On	To port	On
TCP	Any	NFM-P server(s)	12300-12307	NFM-P auxiliary collector(s)
TCP	12300-12307	NFM-P server(s)	Any	NFM-P auxiliary collector(s)
TCP	Any	NFM-P server(s)	12800	NFM-P auxiliary collector(s)
TCP	12800	NFM-P server(s)	Any	NFM-P auxiliary collector(s)

Table 7-44 Firewall rules for traffic between redundant NFM-P auxiliary statistics collectors

Protocol	From port	On	To port	On
TCP	Any	NFM-P auxiliary statistics collector	22	NFM-P auxiliary statistics collector
TCP	Any	NFM-P auxiliary statistics collector	9010	NFM-P auxiliary statistics collector

Table 7-45 Firewall rules for traffic between redundant NFM-P auxiliary call trace collectors

Protocol	From port	On	To port	On
TCP	Any	NFM-P auxiliary call trace collector	22	NFM-P auxiliary call trace collector
TCP	Any	NFM-P auxiliary call trace collector	9010	NFM-P auxiliary call trace collector

7.4.5 NFM-P auxiliary database firewall rules

Apply the following firewall rules to the connection between the NFM-P auxiliary database and various NFM-P components. Note that all connections are bi-directional. Since the inter-node communication should traverse a private LAN, it is not recommended to implement a firewall on this interface.

Table 7-46 Firewall rules for traffic between the NFM-P server and the NFM-P auxiliary database

Protocol	From port	On	To port	On	Notes
TCP	>32768	NFM-P server(s)	5433	NFM-P auxiliary database(s)	JDBC
TCP	>32768	NFM-P server(s)	7299–7309	NFM-P auxiliary database(s)	RMI

Table 7-47 Firewall rules for traffic between the NFM-P auxiliary database(s) and NSP

Protocol	From port	On	To port	On
TCP	>32768	NSP	5433	NFM-P auxiliary database(s)
TCP	>32768	NSP	7299-7309	NFM-P auxiliary database(s)

Table 7-48 Firewall rules for traffic between the NFM-P auxiliary statistics collector and the NFM-P auxiliary database

Protocol	From port	On	To port	On	Notes
TCP	>32768	Statistics auxiliary(s)	5433	NFM-P auxiliary database(s)	JDBC

Table 7-49 Firewall rules for traffic between the NSP flow collector and the NFM-P auxiliary database

Protocol	From port	On	To port	On	Notes
TCP	>32768	NSP flow collector(s)	5433	NFM-P auxiliary database(s)	JDBC

Table 7-50 Firewall rules for traffic between the NSP analytics server and the NFM-P auxiliary database

Protocol	From port	On	To port	On	Notes
TCP	>32768	NSP analytics server	5433	NFM-P auxiliary database(s)	JDBC

Table 7-51 Firewall rules for traffic between NFM-P auxiliary database clusters

Protocol	From port	On	To port	On	Notes
TCP	>32768	NFM-P auxiliary database	22	NFM-P auxiliary database(s)	SFTP
TCP	>32768	NFM-P auxiliary database	50000	NFM-P auxiliary database(s)	Rsync

7.4.6 NSP analytics server firewall rules

Apply the following firewall rules to the connection between the NSP analytics server and NFM-P server / NFM-P client. Note that all connections are bi-directional.

Table 7-52 Firewall rules for traffic between the NFM-P server(s) and NSP analytics server

Protocol	From port	On	To port	On	Notes
TCP	>32768	NSP analytics server(s)	2181	NFM-P server(s)	ZooKeeper
TCP	>32768	NSP analytics server(s)	2281	NFM-P server(s)	ZooKeeper
TCP	>32768	NFM-P server(s)	8080	NSP analytics server	HTTP
TCP	>32768	NFM-P server(s)	8443	NSP analytics server	HTTPS

Table 7-53 Firewall rules for traffic between the NFM-P client and NSP analytics server

Protocol	From port	On	To port	On	Notes
TCP	>32768	Client	8080	NSP analytics server(s)	HTTP
TCP	>32768	Client	8443	NSP analytics server(s)	HTTPS

Apply the following firewall rules to the connection between the NSP analytics server and NSP server when NFM-P is integrated with NSP. Note that all connections are bi-directional.

Table 7-54 Firewall rules for traffic between the NSP analytics server and NSP

Protocol	From port	On	To port	On	Notes
TCP	>32768	NSP analytics server(s)	2181	NSP	ZooKeeper
TCP	>32768	NSP analytics server(s)	2281	NSP	ZooKeeper Secure

7.4.7 NFM-P server to NFM-P client delegate

Ensure that ICMP protocol traffic from the NFM-P server can reach the NFM-P client delegate.

7.4.8 NFM-P client to managed network communications

Apply the following changes to the connection between the NFM-P client and the managed network. Note that all connections are bi-directional.

Table 7-55 Firewall rules for traffic between the NFM-P client and GNEs

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P client(s)	80	Managed network	HTTP (See GNE vendor for specifics)
TCP	Any	NFM-P client(s)	443	Managed network	HTTPS (See GNE vendor for specifics)

Table 7-56 Firewall rules for traffic between the NFM-P client (NEtO) and 9500 MPR / Wavence SM (MSS-8/4/1)

Protocol	From port	On	To port	On	Notes
TCP	20	NFM-P client(s)	Any	Managed network	Active FTP
TCP	Any	NFM-P client(s)	21	Managed network	FTP
TCP	21	NFM-P client	Any	Managed network	FTP
TCP	22	NFM-P client	Any	Managed network	sFTP
TCP	Any	NFM-P client	22	Managed network	sFTP
TCP	Any	NFM-P client(s)	23	Managed network	Telnet
TCP	Any	NFM-P client(s)	80	Managed network	HTTP
UDP	Any	NFM-P client(s)	161	Managed network	SNMP
TCP	>1023	NFM-P client(s)	>1023	Managed network	Passive FTP
UDP	5010	NFM-P client(s)	5010	Managed network	SNMP

Table 7-57 Firewall rules for traffic between the NFM-P client (NEtO) and 9500 MPR / Wavence SM (MSS-1C / MPR-e / MSS-8)

Protocol	From port	On	To port	On	Notes
TCP	20	NFM-P client(s)	Any	Managed network	Active FTP
TCP	21	NFM-P client	Any	Managed network	FTP
TCP	Any	NFM-P client(s)	23	Managed network	Telnet
UDP	Any	NFM-P client(s)	161	Managed network	SNMP
TCP	>1023	NFM-P client(s)	>1023	Managed network	Passive FTP
UDP	5010	NFM-P client	Any	Managed network	SNMP
UDP	Any	NFM-P client	11500	Managed network	Equipment View (GUI)

Table 7-58 Firewall rules for traffic between the NFM-P client (NEtO) and 9400 AWY

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P client(s)	21	Managed network	FTP
TCP	21	NFM-P client	Any	Managed network	FTP
TCP	Any	NFM-P client(s)	23	Managed network	Telnet
TCP	Any	NFM-P client(s)	80	Managed network	HTTP
UDP	Any	NFM-P client(s)	161	Managed network	SNMP
TCP	>1023	NFM-P client(s)	>1023	Managed network	Passive FTP
UDP	5010	NFM-P client	Any	Managed network	SNMP

Table 7-59 Firewall rules for traffic between the NFM-P client and OmniSwitches

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P client(s)	80	Managed network	HTTP
TCP	Any	NFM-P client(s)	443	Managed network	HTTPS

7.4.9 NFM-P client to NSP

If NFM-P is integrated with NSP, the following firewall rules must be implemented on the NFM-P client. Note that all connections are bi-directional.

Table 7-60 Firewall rules for traffic between the NFM-P client and NSP

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P client(s)	80	NSP	HTTP
TCP	Any	NFM-P client(s)	443	NSP	HTTPS

7.4.10 NFM-P to pki-server

If the pki-server is used for TLS certificate generation, the NFM-P server, NFM-P auxiliary servers, NSP analytics server, and NSP flow collector require the following firewall rules:

Table 7-61 Firewall rules for traffic between NFM-P and the server hosting the pki-server

Protocol	From port	On	To port	On	Notes
TCP	Any	NFM-P	2391	Server hosting pki-server	default port
TCP	2391	Server hosting pki-server	Any	NFM-P	default port

8 Multiple network interface NFM-P deployments

8.1 Multihoming

8.1.1 Overview

The NFM-P server and NFM-P auxiliary collector components of the application communicate with very different entities: a managed network, a collection of clients (GUIs and XML API), and between each other. Since these entities usually exist in very different spaces, Nokia recognizes the importance of separating these different types of traffic. Nokia therefore supports configuring the NFM-P server and NFM-P auxiliary such that it uses different network interfaces (IP addresses) to manage the network and to service the requirements of the NFM-P clients.

The NFM-P server uses an internal communications system (JGroups/JMS) to handle bi-directional access to the NFM-P server for the NFM-P clients and the NFM-P auxiliary collectors. In NFM-P, this communication system can be configured to allow the NFM-P clients and NFM-P auxiliary collectors to communicate using different network interfaces on the NFM-P server. This adds significant flexibility when isolating the different types of traffic to the NFM-P server. If using this mode, special attention must be paid to the firewall rules on the network interfaces on the NFM-P server and NFM-P auxiliary collectors (NICs 1 and NICs 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)).

It is a security requirement that all IP communications from an NFM-P auxiliary collector to the NFM-P main server use only one IP address. This IP Address must be the same IP address as the auxiliary collector IP address configured when installing the main server. Any other IP communications originating from a different IP address on the auxiliary collector will be rejected by the NFM-P main server.

When installing NFM-P components on workstations with multiple interfaces, each interface must reside on a separate subnet, with the exception of interfaces that are to be used in IP Bonding.

Figure 8-1, “Collocated NFM-P server/database deployment with multiple network interfaces” (p. 126) illustrates a collocated NFM-P server/database deployment where the NFM-P is configured to actively use more than one network interface.

It is not necessary to use the first network interface on the NFM-P server workstation (i.e. ce0, bge0) to communicate with the NFM-P GUI clients.

Figure 8-1 Collocated NFM-P server/database deployment with multiple network interfaces

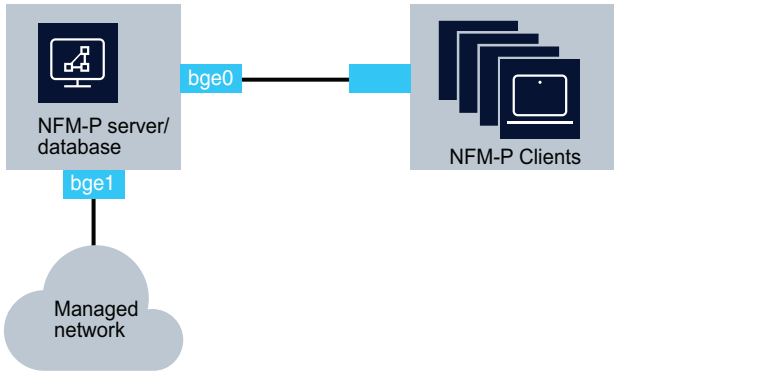

Figure 8-2 Distributed NFM-P server/database deployment with multiple network interfaces

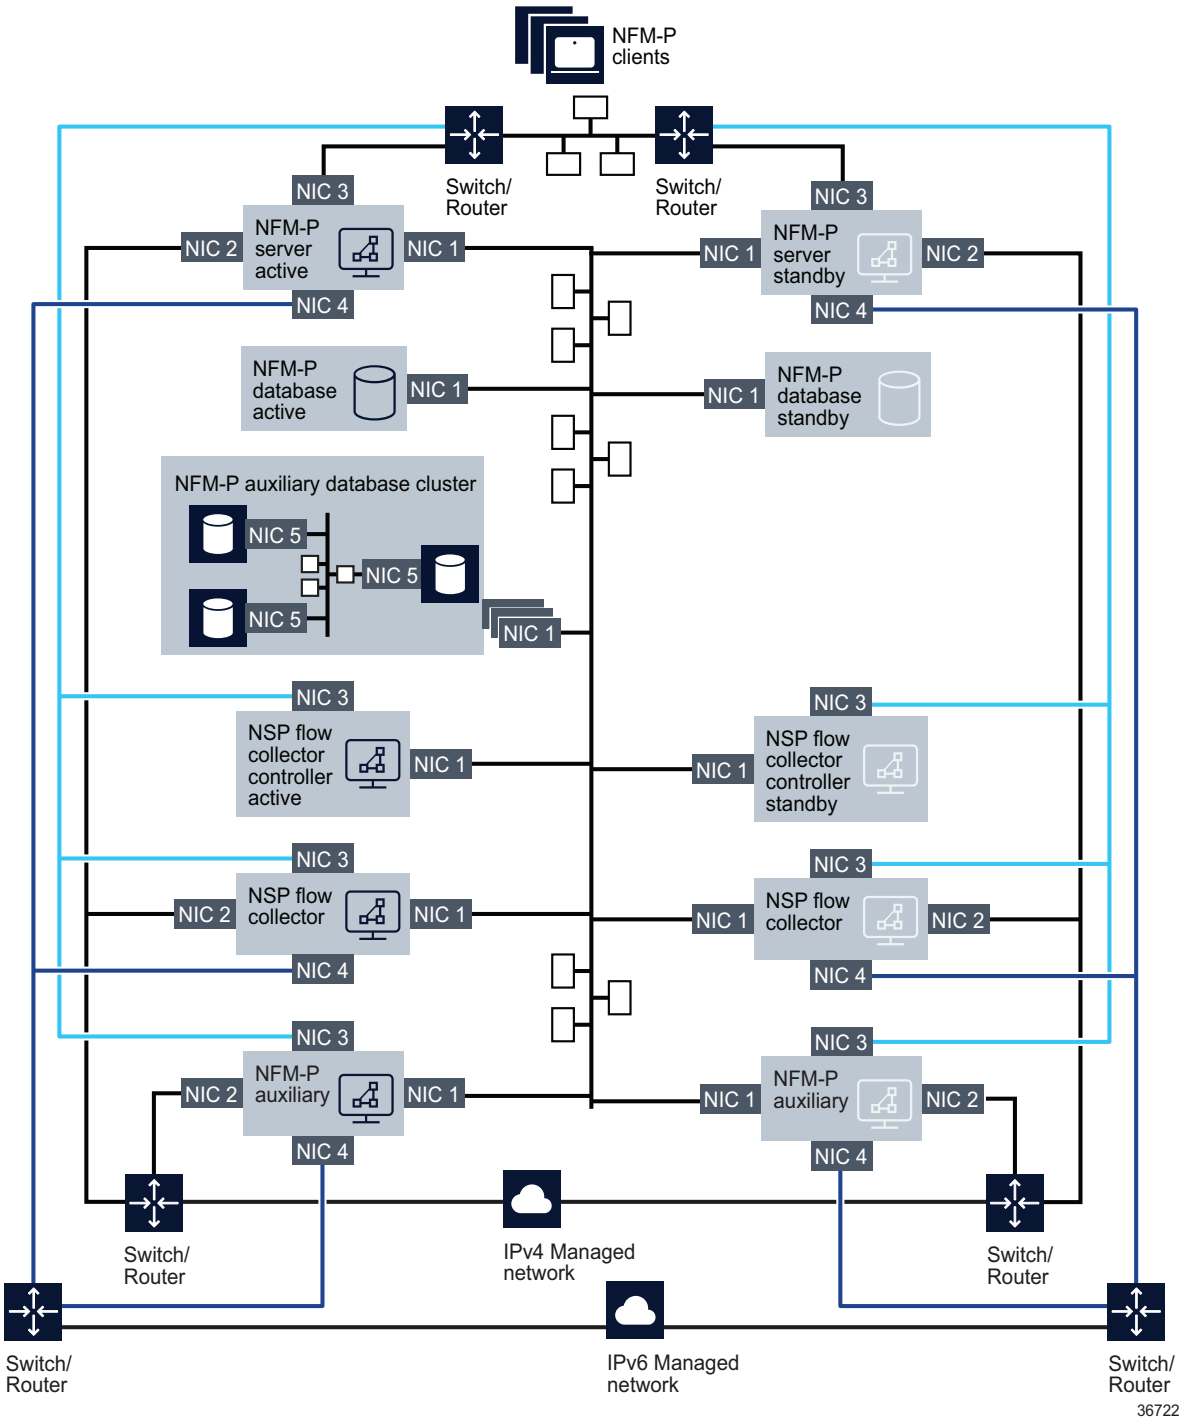

Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127) illustrates a distributed, redundant NFM-P deployment where the NFM-P components are configured to actively use more than one network interface.

Due to limitations with the inter-process and inter-workstation communication mechanisms, a specific network topology and the use of hostnames is required (see 8.3 “Use of hostnames for the NFM-P client” (p. 134)). Contact an Nokia representative to obtain further details.

8.1.2 NFM-P server multiple IP addresses deployment scenarios

The NFM-P server supports the configuration of different IP addresses for the following purposes:

- One or multiple network interfaces can be used to manage the network. (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the managed devices will use to communicate with the NFM-P server and NFM-P auxiliary. If managing a network element with both an in-band and out-of-band connection, the same network interface on the NFM-P server must be used for both communication types.
- One network interface can be used to service the requirements of the NFM-P clients (GUIs and XML API) (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that all clients (GUIs and XML API) will use to communicate with the NFM-P server. All clients (GUIs and XML API) must be configured to use the same IP address to communicate to the NFM-P server. This IP address can be different from the one used by the managed devices to communicate with the NFM-P server. Each client can use the hostname to communicate with the NFM-P server, where the hostname could map to different IP addresses on the NFM-P server - i.e. some clients could connect over IPv4 and some over IPv6. In this scenario, the NFM-P server must be configured for clients to use hostname and not IP.
- One network interface can be used to communicate with the NFM-P database, NFM-P auxiliary database, and NFM-P auxiliary collectors as well as any redundant NFM-P components should they be present (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that the NFM-P database, NFM-P auxiliary database, and redundant NFM-P components will use to communicate with the NFM-P server. This IP address can be different from the addresses used by the NFM-P clients and the managed devices to communicate with the NFM-P server.
- In a redundant NFM-P installation, the NFM-P servers and NFM-P auxiliary collectors must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NFM-P server workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.
- IPv4 and IPv6 network elements can be managed from the same interface or from separate interfaces. (NIC2 and/or NIC4 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)).

8.1.3 NFM-P auxiliary statistics collector multiple IP addresses deployment scenarios

The NFM-P auxiliary statistics collector supports the configuration of different IP addresses for the following purposes:

- One or multiple network interfaces can be used to retrieve information from the managed network. (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the managed devices will use to retrieve the accounting statistics files, and performance statistics from the network elements.
- One network interface can be used to service the requirements of the XML API clients (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that all XML API clients will use to communicate with the NFM-P auxiliary statistics collector. XML API clients will use this IP address to retrieve the logToFile statistics collection data from the NFM-P auxiliary statistics collector.
- One network interface can be used to communicate with the NFM-P server, NFM-P database, NFM-P auxiliary database cluster as well as any redundant NFM-P components should they be present (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that the NFM-P server, NFM-P database, NFM-P auxiliary database, and redundant NFM-P components will use to communicate with the NFM-P auxiliary statistics collector. This IP address can be different from the addresses used by the NFM-P XML API clients and the managed devices to communicate with the NFM-P auxiliary statistics collector.
- In a redundant NFM-P installation, the NFM-P auxiliary statistics collector must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NFM-P auxiliary statistics collector workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.
- IPv4 and IPv6 network elements can be managed from the same interface or from separate interfaces. (NIC2 and/or NIC4 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)).

8.1.4 NFM-P auxiliary call trace collector multiple IP addresses deployment scenarios

The NFM-P auxiliary call trace collector supports the configuration of different IP addresses for the following purposes:

- One network interface can be used to retrieve information from the managed network. (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the managed devices will use to send the call trace messages from the network elements.
- One network interface can be used to service the requirements of the 9958 WTA client (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that all clients will use to communicate

with the NFM-P auxiliary call trace collector. 9958 WTA will use this IP address to retrieve the call trace data from the NFM-P auxiliary call trace collector.

- One network interface can be used to communicate with the NFM-P management complex as well as any redundant NFM-P components should they be present (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that the NFM-P management complex components will use to communicate with the NFM-P auxiliary call trace collector. If a redundant NFM-P auxiliary call trace collector is present, this network interface will also be used to sync call trace and debug trace data collected from the network, with the peer NFM-P auxiliary call trace collector. This IP address can be different from the addresses used by the 9958 WTA clients and the managed devices to communicate with the NFM-P server.
- In a redundant NFM-P installation, the NFM-P auxiliary call trace collector must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NFM-P auxiliary call trace collector workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.

8.1.5 NSP flow collector controller multiple IP addresses deployment scenarios

The NSP flow collector supports the configuration of different IP addresses for the following purposes:

- One network interface can be used to communicate with the NFM-P management complex as well as any redundant NFM-P components, should they be present (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the NFM-P management complex components will use to communicate with the NSP flow collector controller. This IP address can be different from the addresses used by the clients and the managed devices to communicate with the NFM-P server. If the NSP deployment includes NSP, this is the network interface that would be used for communication.
- One network interface can be used to communicate with the clients (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the user will connect to with the web management interface.
- In a redundant NFM-P installation, the NSP flow collector controller must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NSP flow collector controller workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.

8.1.6 NSP flow collector multiple IP addresses deployment scenarios

The NSP flow collector supports the configuration of different IP addresses for the following purposes:

- One network interface can be used to communicate with the NSP flow collector controller (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the NSP flow collector controller and NFM-P server will use to communicate with the NSP flow collector(s). This IP address can

be different from the addresses used by the clients and the managed devices to communicate with the NFM-P server. If the NSP deployment includes either NSP, this is the network interface that would be used for communication.

- One network interface can be used to retrieve information from the managed network. (NIC 2 and/or NIC 4 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that the managed devices will use to send the cflowd flow data from the network elements.
- One network interface can be used to communicate with the clients (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the user will connect to with the web management interface.
- One network interface can be used to send the formatted IPDR files to the target file server (NIC 4 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that all clients will use to communicate with the NSP flow collector.
- In a redundant NFM-P installation, the NSP flow collector must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NSP flow collector workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.

8.1.7 NFM-P auxiliary PCMD collector multiple IP addresses deployment scenarios

The NFM-P auxiliary PCMD collector supports the configuration of different IP addresses for the following purposes. To meet scaling targets a minimum of two separate interfaces must be used — one for management traffic and one for PCMD data collection:

- One network interface can be used to retrieve information from the managed network. (NIC 2 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)) This network interface contains the IP address that the managed devices will use to send the PCMD data from the network elements.
- One network interface can be used for retrieval of the formatted PCMD files by the target file server (NIC 3 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that all clients will use to communicate with the NFM-P auxiliary PCMD collector.
- One network interface can be used to communicate with the NFM-P management complex as well as any redundant NFM-P components should they be present (NIC 1 on Figure 8-2, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 127)). This network interface contains the IP address that the NFM-P management complex components will use to communicate with the NFM-P auxiliary PCMD collector. This IP address can be different from the addresses used by the clients and the managed devices to communicate with the NFM-P server.
- In a redundant NFM-P installation, the NFM-P auxiliary PCMD collector must have IP connectivity to the NFM-P server peer.
- Additional network interfaces may be configured on the NFM-P auxiliary PCMD collector workstation, at the customer’s discretion, to perform maintenance operations such as workstation backups.

8.2 Network Address Translation

8.2.1 Network Address Translation deployment scenarios

NFM-P supports the use of Network Address Translation (NAT) between the following components:

- The NFM-P server and NFM-P clients (GUIs or XML API)
- The NFM-P auxiliary server and NFM-P XML API clients
- The NFM-P server and the managed network
- The NFM-P auxiliary statistics collector and the managed network
- The NFM-P auxiliary PCMD collector and the managed network

The figure below illustrates a deployment of NFM-P where NAT is used between the NFM-P server and the managed network.

Figure 8-3 NFM-P server deployments with NAT between the server and the managed network

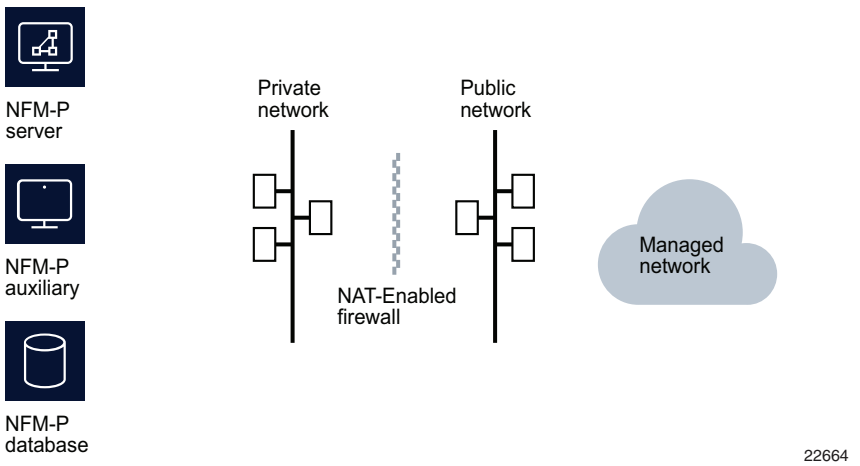

Note: Network Address Translation is not supported between the NFM-P auxiliary call trace collector and the managed network.

The following two figures illustrates a deployment of NFM-P where NAT is used between the NFM-P server and the NFM-P clients (GUIs, XML API or client delegates). In Figure 8-4, “NFM-P server deployment using NAT with IP Address communication” (p. 133), NFM-P clients on the private side and public side of the NAT-Enabled Firewall must connect to the public IP address of the NFM-P server. A routing loopback from the NFM-P server private IP address to the NFM-P server public IP address must be configured in this scenario as all NFM-P clients must communicate to the NFM-P server through the NFM-P server public IP address.

The NFM-P auxiliary will need to be able to connect to the public IP address of the NFM-P server.

Figure 8-4 NFM-P server deployment using NAT with IP Address communication

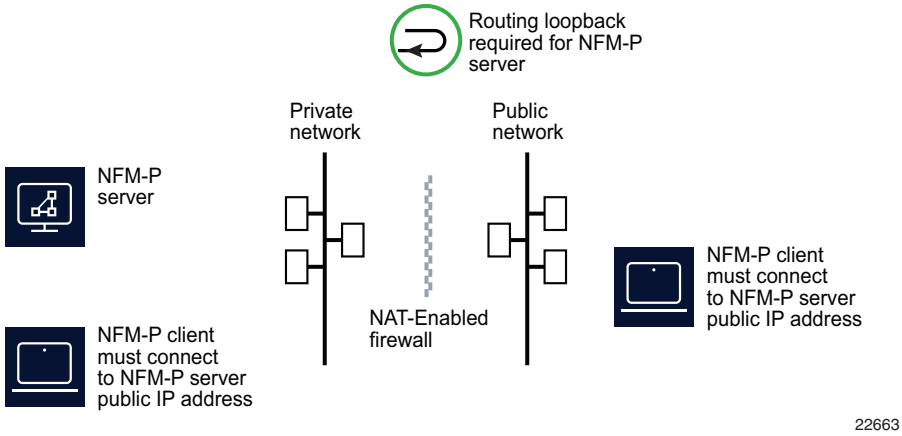

Figure 8-5 NFM-P server deployment using NAT with name resolution based communication

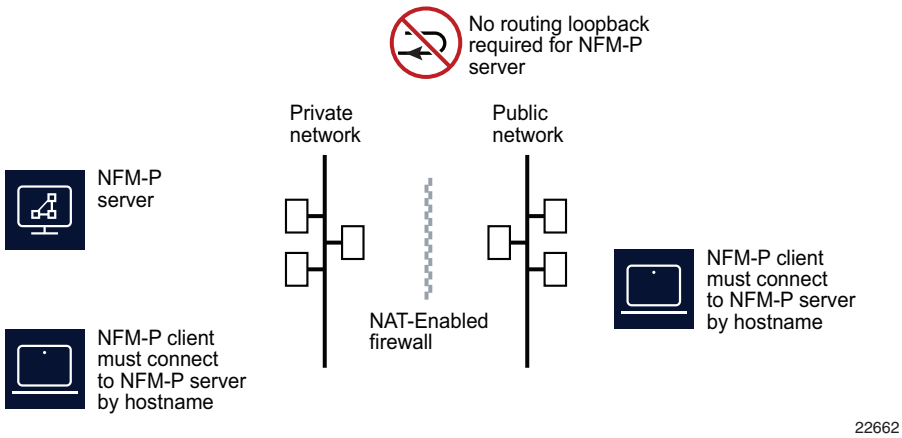

In Figure 8-5, “NFM-P server deployment using NAT with name resolution based communication” (p. 133), a name resolution service on the public side of the NAT-Enabled Firewall is configured to resolve the NFM-P server hostname to the public IP address of the NFM-P server. Name resolution service on the private side of the NAT-Enabled Firewall is configured to resolve the NFM-P server hostname to the private IP address of the NFM-P server. clients on both sides of the NAT-Enabled Firewall are configured to communicate with the NFM-P server via hostname where the NFM-P server hostname must be the same on both sides of the NAT-Enabled Firewall.

The figure below illustrates a deployment of NFM-P where NAT is used between the NFM-P complex, NFM-P clients, and the managed network.

Figure 8-6 NFM-P deployment with NAT

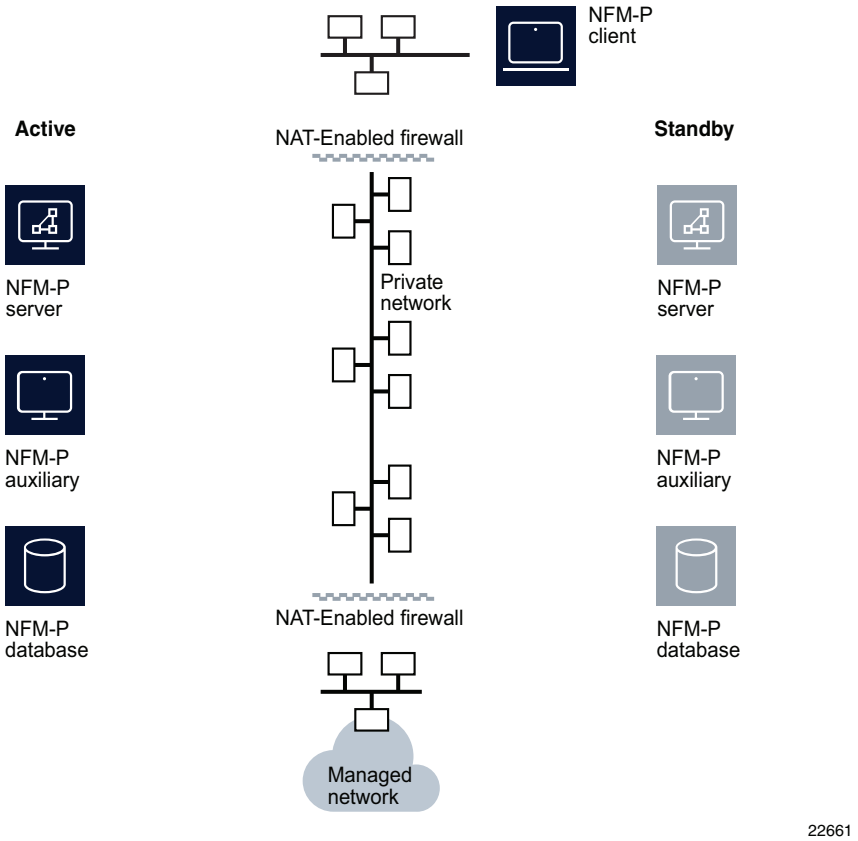

For installations using NAT between the NFM-P server and NFM-P client, a reverse DNS look-up mechanism must be used for the client, to allow proper startup.

NAT rules must be in place before NFM-P installation can occur, since the installation scripts will access other systems for configuration purposes.

Note: Network Address Translation is not supported between the NFM-P auxiliary call trace collector and the managed network.

8.3 Use of hostnames for the NFM-P client

8.3.1 Hostnames usage scenarios

The following scenarios identify situations where it is necessary for the NFM-P client to be configured to use a hostname rather than a fixed IP address to reach the NFM-P server:

- When CA signed TLS certificates are used, the FQDN must be used for client communication.

- When NFM-P clients can connect to the NFM-P server over multiple interfaces on the NFM-P server. For example, when clients can connect over both IPv4 and IPv6 interfaces.
- When NAT is used between NFM-P clients and the NFM-P server.
- For situations where the NFM-P client and the NFM-P auxiliary (and/or NFM-P peer server) are using different network interfaces to the NFM-P server, the NFM-P client must use a hostname to reach the NFM-P server.