



NSP

Network Services Platform

Release 25.11

Planning Guide

3HE-21462-AAAC-TQZZA
Issue 3
March 2026

© 2026 Nokia.

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

Legal notice

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

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

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

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

No part of this document may be modified.

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

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

© 2026 Nokia.

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.....7**
- 1 NSP product overview9**
 - 1.1 The Network Services Platform9
 - 1.2 Core NSP technologies.....9
 - 1.3 NSP deployment12
 - 1.4 Including the NFM-P in an NSP deployment.....16
- 2 NSP system requirements.....25**
 - 2.1 NSP container environment requirements25
 - 2.2 NSP cluster requirements30
 - 2.3 Other NSP components hardware platform requirements33
 - 2.4 Other NSP components hardware platform and resource requirements using virtualization.....34
 - 2.5 Other NSP components minimum platform requirements.....38
 - 2.6 NFM-P platform requirements for larger networks46
 - 2.7 Storage47
- 3 Operating system specifications51**
 - 3.1 Red Hat Enterprise Linux (RHEL).....51
 - 3.2 NFM-P on Microsoft Windows.....52
 - 3.3 NFM-P on Apple macOS.....53
 - 3.4 NSP operating system summary.....53
 - 3.5 Third-party software for NFM-P single-user client or client delegate server53
- 4 Network requirements55**
 - 4.1 NSP network requirements55
 - 4.2 NSP deployment network addressing requirements55
 - 4.3 Network requirements between NSP and other components57
 - 4.4 Network requirements for NSP redundancy and communications within an NSP cluster57
 - 4.5 NSP Support for RHEL IP Bonding58
 - 4.6 NFM-P network requirements58
 - 4.7 Network elements58
 - 4.8 Bandwidth requirements59
 - 4.9 Contributors to bandwidth requirements63
 - 4.10 Network bandwidth65
 - 4.11 Network latency68
 - 4.12 Network reliability.....70

4.13	Network element specific requirements	71
4.14	Mechanism to maintain the current state of network elements	72
5	Scaling and performance	75
5.1	NSP scaling and performance	75
5.2	Scale limits for NSP deployments	75
5.3	Scale limits for functions	79
5.4	NSP application downtime in HA and redundant deployments	87
5.5	NFM-P Scalability guidelines	90
5.6	Scaling guidelines for NFM-P XML API clients	94
5.7	Scaling guidelines for statistics collection	94
5.8	Scaling guidelines for service assurance tests	100
5.9	cflowd statistics collection	106
5.10	CPAM and vCPAA.....	109
6	Security	111
6.1	Introduction.....	111
6.2	Securing the NSP	111
6.3	Operating system security for NSP cluster components	112
6.4	Communication between the NSP and external systems	112
6.5	NSP Port Communications.....	114
6.6	NSP Kubernetes Platform Communications	133
6.7	Securing the NFM-P	135
6.8	Additional NSP component port information	138
6.9	FTP	146
6.10	Additional NSP component firewall and NAT rules	146
7	NSP deployment with multiple network interfaces and IP addresses	155
7.1	Support for multiple network interfaces.....	155
7.2	NSP Network Address Translation.....	158
7.3	NFM-P multihoming	158
7.4	NFM-P Network Address Translation.....	163
7.5	Use of hostnames for the NFM-P client	165
8	Appendix A	167
8.1	Storage-layer I/O performance tests	167

About this document

Purpose

The *NSP Planning Guide* is intended for technology officers, network planners, and system administrators who need the information required to plan a successful deployment of the Nokia Network Services Platform, or NSP. The reader is encouraged to become familiar with the NSP architecture, the relevant components for both IP and optical networks, and the virtualization, system, and network requirements.

Scope

The scope of this document is limited to the planning of an NSP deployment. The guide provides an overview of the NSP product and the components that comprise an NSP deployment, including the NFM-P. The guide also describes operating-system specifications, and system resource and network requirements for successful NSP system deployment. Scale limits and general information about platform security recommendations are also included.

NFM-P only deployments

Starting in Release 24.4, NSP and NFM-P components must be deployed in tandem; independent NFM-P installations are no longer supported.

Safety information

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

Document support

Customer documentation and product support URLs:

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

How to comment

Please send your feedback to [Documentation Feedback](#).

1 NSP product overview

1.1 The Network Services Platform

1.1.1 NSP system overview

The Network Services Platform, or NSP, provides role-based network equipment, infrastructure, service rollout, resource control, and FCAPS management within and among multiple network domains.

The NSP management scope includes Nokia and multi-vendor network elements, or NEs, using a variety of protocols and element-management mechanisms. The NSP has interfaces for functions such as programmable multi-layer service provisioning, rollout, and activation.

1.2 Core NSP technologies

1.2.1 Java Virtual Machine

NSP employs Java technology. The NSP installation software includes a Java Virtual Machine, or JVM that is dedicated to the NSP and does not conflict with other JVMs that may be installed on the same station. NSP uses OpenJDK versions 8, 11 and 17.

NFM-P uses OpenJDK 17, and a version of the Oracle JVM is also embedded with the NFM-P database. This embedded versions of the Oracle JVM is fully licensed through Nokia and can only be used for their embedded purpose.

1.2.2 Kubernetes container environment

An NSP system is deployed as cluster of one or more VMs in the Kubernetes container environment and consists of a deployer host (which holds the container image and Helm repositories and deploys the containerization environment for the NSP cluster), and the NSP cluster VMs (which host the main NSP functions and load balance as needed). Within the NSP cluster VMs, the NSP cluster host performs configuration operations. Kubernetes is the orchestration system that hosts the NSP VM clusters and is included in the NSP software bundle. For more information about the NSP components described here, see the *NSP Installation and Upgrade Guide*.

1.2.3 Databases

An NSP deployment has multiple databases. NSP PCE has a Neo4j database for network topology information. The nspOS component has a PostgreSQL database for policy management and common applications data, and a Neo4j database for topology data for the Map Server. NSP's IP/optical coordination function contains a Neo4j database for network topology information.

A Neo4j database contains a representation of the network topology and alarm KPI information in a multi-layer graph. The installation of the Neo4j database is customized for, and must be dedicated to, the NSP. Data redundancy and replication within an HA cluster is managed within the neo4j instances.

The PostgreSQL database contains non-topological NSP information, including policies, templates, and nspOS common model data. PostgreSQL is an open source database application.

PostgreSQL database redundancy is managed by the role-manager. In a redundant configuration of the NSP, the active NSP cluster hosts the primary PostgreSQL database. The standby NSP cluster hosts the standby PostgreSQL database.

In an HA NSP cluster deployment, the primary/active site contains both a primary PostgreSQL database (read/write), as well as a replica (read only). The standby site contains a standby leader, and a replica (replicating from the standby leader). When the primary database goes down, a qualified replica may be promoted to primary. When this occurs, the database that previously served as primary is demoted to replica. A qualified replica is a functioning replica that is part of the cluster and is sufficiently synchronized with the primary (this threshold of synchronization can be configured; the default is within 50MB). If no qualified replica is identified - but a replica exists - the user must manually promote that replica, with the understanding that some data loss will occur.

i **Note:** Nokia does not support any PostgreSQL database configuration that deviates from the NSP installation procedure.

i **Note:** PostgreSQL cluster state is kept in configmaps. If PostgreSQL data is deleted manually, the cluster state must be deleted as well, or the cluster will not initialize.

i **Note:** Nokia does not support direct customer access to the Neo4j and PostgreSQL databases.

NFM-P 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 must be aware of and must accept this risk when choosing to run NFM-P in a RHEL KVM virtualized environment.

NSP auxiliary database

Embedded within the NSP 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 NSP License Agreement Violation. Modifying the installation can impact system performance, stability and upgrades. Customer support agreements may be violated.

1.2.4 Network mediation

The NSP application has southbound interfaces that consist of plug-ins that interact with the NFM-P using CPROTO and HTTP protocols secured with TLS. The NFM-P manages IP network elements using SNMP.

The NSP communicates with MDM using gRPC, and MDM communicates with network elements using gRPC/gNMI, NETCONF, SNMP and CLI over SSH or Telnet.

For LSP management (as enabled by NSP's path control function), a VSR-NRC communicates with the PCC network elements via PCEP, IGP, and BGP.

The WS-RC is installed with WS-NOC to provide a TLS secured REST API for optical network discovery and service provisioning. The WS-NOC uses TL-1 and SNMP to manage optical switches.

1.2.5 Browser applications

NOTICE

View of network can be affected

Browser applications' view of the network can be affected whenever activities are drawing heavily on CPU and memory usage.

This can happen when a large number of services are being created, modified, or deleted via the NSP REST APIs.

The NSP provides functionality using browser-based NSP applications. These applications require that WebGL be enabled and use standard REST security mechanisms for authentication and authorization. All NSP applications are HTML-5 based and are supported on the following web browsers:

- Latest version of Google Chrome
- Latest version of Chromium Edge for most applications (see the NSP Release Notice for restrictions)
- Latest version of Mozilla Firefox
- Latest version of Apple Safari

Additional Internet browsers and older versions may function with NSP but are not supported by Nokia.

i **Note:** You cannot switch browsers between clients. You must always use the system default browser.

i **Note:** In order for the Safari web browser to open NSP's analytics views, you must ensure that the following Safari privacy settings are configured, if present in your browser version:

- Safari Preferences page, Cookies And Website Data—Always Allow
- Prevent cross-site tracking—disabled

i **Note:** If you are using Chrome or Firefox on Windows 8.1, it is recommended that you enable ClearType Text for optimal viewing of fonts. To enable, open the Display settings in Windows Control Panel and enable the Turn on ClearType parameter under the Adjust ClearType text settings.

Localized language support

NSP supports localized language display. Localized language display, also known as internationalization, displays GUI text in a specified language. The localized language setting applies to most GUI objects, except system components and database objects. Contact Nokia technical support for more information about localized language support.

i **Note:** The NSP components support localized language settings using predefined strings, and do not translate data to different languages.

1.2.6 APIs

NSP provides APIs to expose the capabilities of its various modules and functions. NSP APIs are designed with REST, RESTCONF and Kafka to access and manipulate resources. The [Network Developer Portal](#) contains updated information for each API as well as examples and documentation.

1.3 NSP deployment

1.3.1 NSP deployment types

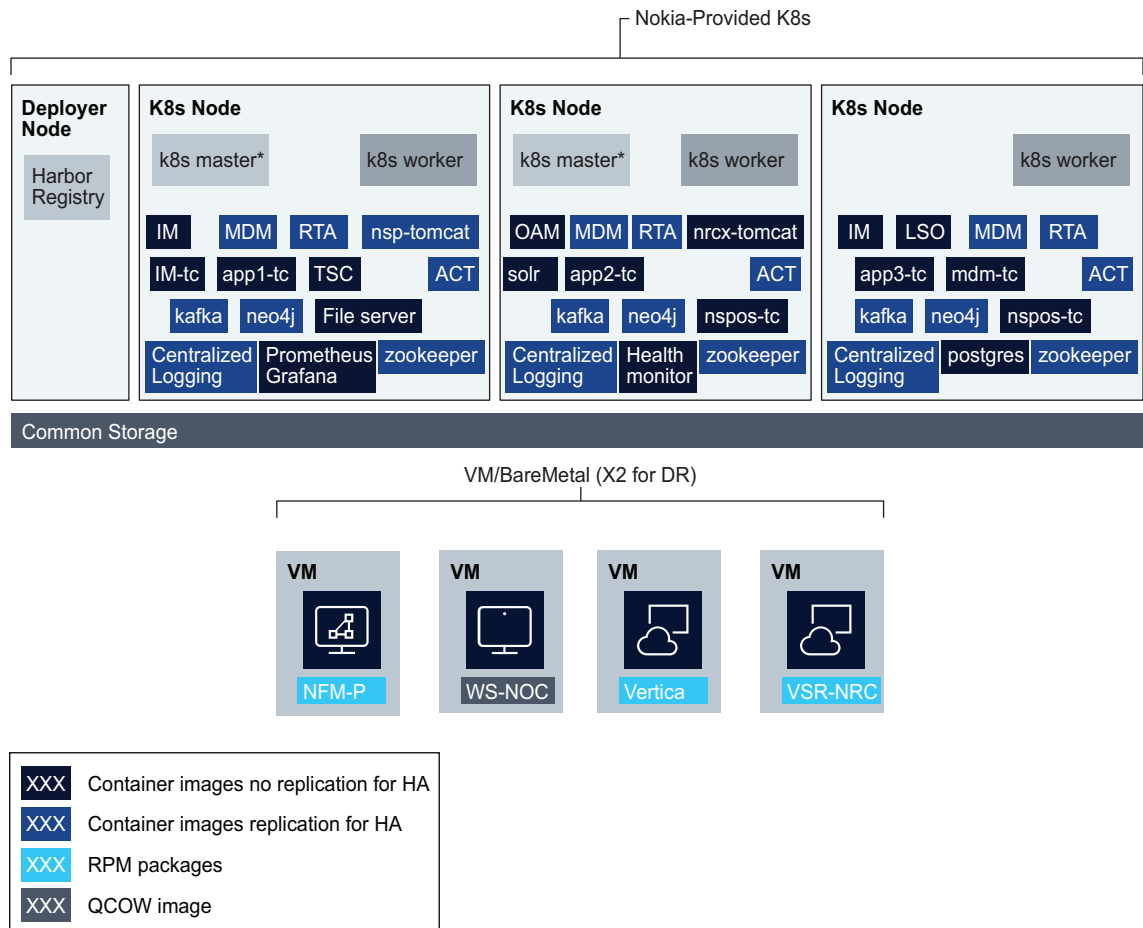
An NSP deployment in a data center consists of the following:

- container environment in which a Kubernetes orchestration layer co-ordinates NSP service deployment in a group of virtual machines called the NSP cluster. Each VM in the cluster is deployed as a Kubernetes node.
- RPM-based components on separate stations or VMs outside the NSP cluster

Production deployment

[Figure 1-1, “NSP system, production deployment” \(p. 13\)](#) shows an example three-node NSP cluster and a number of external RPM-based components. Depending on the specified NSP deployment profile and installation options, a cluster may have additional nodes.

Figure 1-1 NSP system, production deployment




38107

Lab deployment

A one-node NSP cluster is used in a lab deployment, as shown in [Figure 1-2, “NSP system, lab deployment”](#) (p. 14).

Figure 1-2 NSP system, lab deployment



38444

1.3.2 NSP system redundancy

NSP installations can provide redundancy through two mechanisms:

- **High-availability (HA)** is a local fault recovery mechanism within a multi-node NSP cluster. HA uses Kubernetes pod replicas to ensure minimal downtime in the event of the failure of a pod that provides critical services to the cluster, avoiding a system switchover to a redundant DR data center.
- **Disaster recovery (DR)** is a redundancy option available to any deployment size where two identical NSP installations exist in geographically separate data centers, providing fundamental network resiliency through geographical redundancy. Failover times for DR deployments are significantly longer than HA fault recovery and result in a brief interruption of services. In a DR deployment, the geographically dispersed installations must be identical in terms of NSP cluster components and redundancy models.

You must ensure that your NSP deployment has the redundancy that your use case requires. Nokia recommends combining HA and DR for the most robust redundancy solution.


An NSP deployment can include the NFM-P and WS-NOC, which also support redundant deployment. An NSP deployment can have a mixed redundancy model when the NSP cluster is deployed as an enhanced (HA) deployment type. In this scenario, the NSP cluster can be standalone, while the NFM-P or WS-NOC components can be deployed with DR redundancy.

Note: A redundant NSP deployment supports classic HA and fast-HA WS-NOC deployment; see the *NSP Release Notice* for compatibility information.

In a DR deployment of NSP that includes the NFM-P and WS-NOC, the primary NSP cluster operates independently of the primary NFM-P and WS-NOC; if the NSP cluster undergoes an activity switch, the new primary cluster connects to the primary NFM-P and WS-NOC. Similarly, if the NFM-P or WS-NOC undergoes an activity switch, the primary NSP cluster connects to the new primary NFM-P or WS-NOC.

Figure 1-3, “DR NSP deployment including NFM-P and WS-NOC” (p. 14) shows a fully redundant NSP deployment that includes the NFM-P and WS-NOC .

Figure 1-3 DR NSP deployment including NFM-P and WS-NOC




38428

VSR-NRC redundancy

The VSR-NRC also supports redundant deployment. A standalone NSP can be deployed with a standalone or redundant VSR-NRC. When NSP is installed with a redundant VSR-NRC, if the communication channel between NSP and primary VSR-NRC fails, then the NSP switches communication to the standby VSR-NRC.

In a DR deployment of NSP, a redundant deployment of VSR-NRC is required. The primary NSP will communicate to NEs through the primary VSR-NRC, but if the communication channel to primary VSR-NRC fails, then the primary NSP can switch to the standby VSR-NRC.

Figure 1-4 Redundant NSP deployment with redundant VSR-NRC



27498

When an activity switch takes place between redundant NSP clusters, the new active NSP cluster communicates with IP NEs through its corresponding VSR-NRC instance.

In a redundant deployment, the two VSR-NRCs synchronize their TE and LSP databases using a cproto sync channel running between them. This allows for faster re-synchronization after a VSR-NRC failure.

i **Note:** VSR-NRC can only connect to one NSP system at a time, regardless of whether that system is standalone or redundant. In order to connect to another, separate NSP system, the VSR-NRC's connection to the previous NSP system would first have to be terminated.

MDM redundancy and HA

The MDM is deployed within the NSP cluster. There is no limit on the total replica count of MDM instances in a NSP cluster. The MDM instances will be evenly distributed among the k8s nodes. In a DR NSP cluster deployment, each NSP cluster will have the same number of nodes and MDM instances.


In an HA NSP cluster, the MDM pods are deployed in a N+M deployment, with N active MDM instances and M standby instances. If an active MDM instance fails, a standby MDM instance in the active cluster will take over management of the nodes that were managed by the failed MDM instance. When the failed instance recovers, it becomes a standby instance (it will not automatically revert to active). When more than M active MDM instances fail simultaneously, a manual activity switch to the standby NSP cluster will be required. Each NSP cluster must have the same N+M configuration of MDM.

1.4 Including the NFM-P in an NSP deployment

1.4.1 NFM-P deployment types

Figure 1-5, "Collocated standalone NFM-P deployment" (p. 17) is an example of a standalone NFM-P deployment in which the NFM-P main server and main database are collocated on one station.


Figure 1-5 Collocated standalone NFM-P deployment



22675

Figure 1-6, “Distributed standalone NFM-P deployment” (p. 17) is an example of a standalone NFM-P deployment in which the NFM-P main server and main database are hosted on separate stations.


Figure 1-6 Distributed standalone NFM-P deployment



22674

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 servers are used. In this configuration there can be up to three active NFM-P auxiliary statistics servers or it could be configured redundant.

Figure 1-7 NFM-P standalone deployment - distributed NFM-P server and NFM-P database configuration and NFM-P auxiliary servers



38329

For bare metal installations, the NFM-P server, NFM-P auxiliary server, NSP Flow server, NSP Flow Collector Controller, NSP auxiliary database, and NFM-P database are supported on specific Intel x86 based HP and Nokia AirFrame stations.

The NFM-P client and client delegate server software may be installed on stations running different operating systems from the NFM-P server, NFM-P auxiliary, NSP auxiliary database, NSP Flow Collector, NSP Flow Collector Controller, and NFM-P database. The NFM-P client can be installed on RHEL 8 server x86-64, Windows, or Mac OS where the NFM-P client delegate server can be installed on RHEL 8 server x86-64, Windows Server 2016, Windows Server 2019, or Windows Server 2022. See [3.5 “Third-party software for NFM-P single-user client or client delegate server” \(p. 53\)](#) for Operating System specifics.

All NFM-P stations in the NFM-P management complex must maintain consistent and accurate time. The RHEL chronyd service is strongly recommended as the time-synchronization mechanism.

1.4.2 NFM-P system redundancy

NFM-P supports redundancy of the NFM-P server, NFM-P database, and NFM-P auxiliary server stations. The NFM-P auxiliary statistics server supports 3+1 redundancy.

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 station 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.

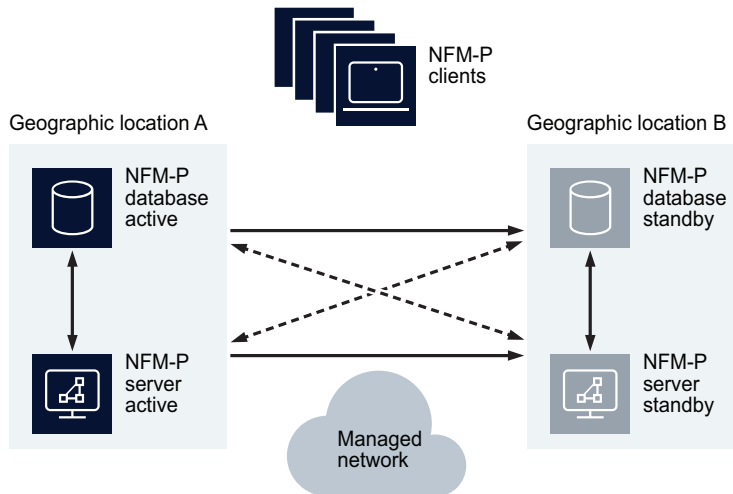
Figure 1-8 NFM-P collocated server/database redundancy deployment



22671

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

Figure 1-9 NFM-P distributed server/database redundancy deployment in a geographically redundant setup.



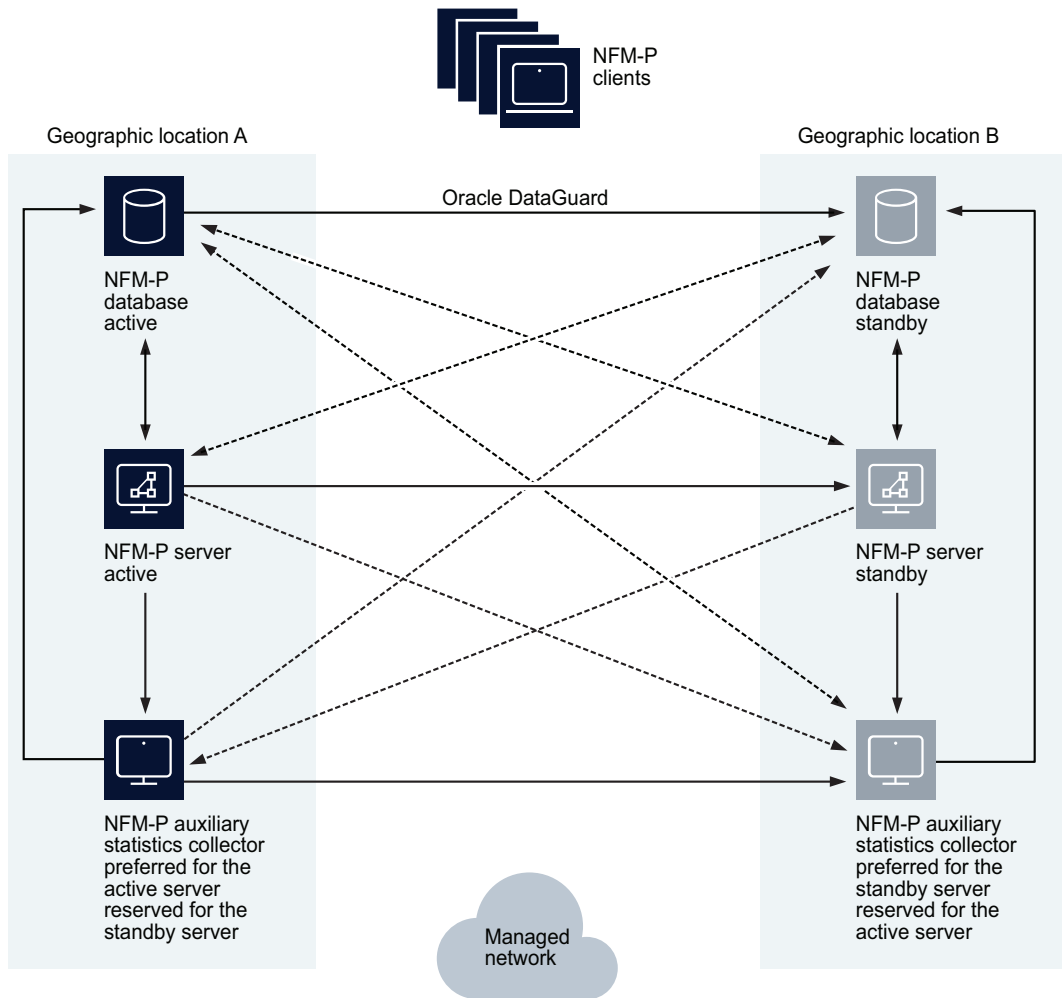
22670

Redundancy and NFM-P auxiliaries

In customer networks, where the statistics collection requirements exceed the scalability capabilities of an NFM-P server, the NFM-P auxiliary statistics server can be used. As with other HA components, the NFM-P auxiliary statistics server 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 server. When collecting statistics using the NSP 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 server.

In [Figure 1-10, “NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries that crosses geographic boundaries” \(p. 19\)](#) there are NFM-P auxiliary servers configured. In the example where redundancy is geographic, there can be up to four NFM-P auxiliary statistics servers configured in each geographic location. The Preferred/Reserved/Remote role of the NFM-P auxiliary statistics server is dependent upon and configured, based on the NFM-P server that is active. When there are more than one active auxiliary statistics server, local redundancy (Preferred/Reserved) of the auxiliary statistics server must be used in conjunction with geographic redundancy, where the same number of auxiliary statistics servers will be deployed in each geographic site. The NFM-P auxiliary statistics servers in the opposite geographic location are configured to be Remote. In this scenario, if one of the NFM-P auxiliary statistics servers for the active NFM-P server were no longer available, the active NFM-P server would use the reserved NFM-P auxiliary statistics server in the same geographic location to collect statistics. If using the same redundant configuration but with statistics collection using a geographically redundant NSP auxiliary database the latency between geographic sites must be less than 200 ms.

Figure 1-10 NFM-P distributed server/database redundant deployment with redundant NFM-P auxiliaries that crosses geographic boundaries



38331

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

Redundancy deployment considerations for NFM-P

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

It is a best practice to keep the NFM-P server, NFM-P database, and NFM-P auxiliary servers 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 servers are all aligned in the same geographic location. If the auto-align functionality is not enabled, a manual switch of the stations is desirable.

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 stations hosting the NFM-P software should be connected in a way to prevent a single physical failure from isolating the two stations from each other.
- Stations that host the NFM-P server/database software must be configured to perform name service lookups on the local station 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.

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 stations to be used as active NFM-P server and NFM-P database should be connected to the pair of stations 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 station pairs from each other.
- Stations that host the NFM-P server and NFM-P database software must be configured to perform name service database lookups on the local station 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.

Redundancy with distributed NFM-P server and NFM-P database and NFM-P auxiliary servers

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 servers must be of the same versions and patch levels as the NFM-P server and NFM-P database stations.
- If collecting statistics using the NSP auxiliary database, the operating systems installed on the

NSP auxiliary database stations must be of the same versions and patch levels as the NFM-P server, NFM-P database, and NFM-P auxiliary statistics server stations.

- NFM-P auxiliary servers are intended to be on the same HA network as the NFM-P server and NFM-P database. NFM-P auxiliary statistics servers 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 the NSP 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 NSP auxiliary database nodes in the same cluster must be deployed in the same geographic site, with less than 1 ms 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.
- When using more than one active NFM-P auxiliary statistics server in a geographic (greater than 1 ms latency) configuration, the active and reserved servers for a give NFM-P server must reside in the same geographic site. The auxiliary statistics servers in the opposite geographic site would be configured as Remote.
- Stations that host the NFM-P auxiliary server software must be configured to perform name service database lookups on the local station 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 NSP system requirements

2.1 NSP container environment requirements

2.1.1 Supported NSP cluster environments

An NSP cluster must be deployed in a Nokia provided container environment.

Nokia supports and recommends installing NSP components on VMs using VMware ESXi or RHEL KVM, including OpenStack. The guest OS must be a supported RHEL version.

2.1.2 Supported container software versions

The NSP cluster is validated against a container environment that uses the following software versions:

Table 2-1 Supported container software versions

Containerization element	Supported version
Kubernetes core	v1.33.5
calico	v3.30.3
cni	v1.8.0
containerd	v2.1.4
etcd	v3.5.23
Helm	v3.18.4
coredns	v1.12.0
kubespray	2.29.0
k9s	v0.50.15
metallb	v0.15.2
harbor app	v2.12.2
nerdctl	v2.1.6

2.1.3 KVM virtualization

The NSP supports using RHEL6, RHEL 7, RHEL 8, and RHEL 9 based KVM on x86 based servers natively supported by KVM. (Please refer to Red Hat documentation for hypervisor compatibility with guest OS.) See the *Host Environment Compatibility Reference for NSP and CLM* for the current KVM compatibility level, requirements, and restrictions. See the RHEL Hardware Compatibility List (HCL) for information about specific hardware support. Please refer to *SR OS Software Release Notes* for information on supported hypervisors for deployment of VSR-NRC.

Not all features offered by KVM are supported when using the NSP. For example Snapshots, Live Migration and HA are not supported. Contact Nokia to determine if a specific KVM feature is supported with an installation of the NSP cluster.

KVM CPU and Memory

The required memory resources must be reserved and dedicated to each guest OS, and cannot be shared or oversubscribed. CPU resources may be oversubscribed using the guidelines in the following table.

Table 2-2 KVM CPU oversubscription guidelines

NSP cluster deployment type	KVM cpu_allocation_ratio
enhanced	2.0
standard	1.5
medium	1.25

i **Note:** CPU oversubscription is only supported on NSP clusters with three or more nodes. Smaller NSP deployments must have reserved and dedicated CPU resources for each node. CPU oversubscription is only supported on Intel based processors using the Skylake or newer microarchitecture. Only the NSP Cluster VMs support oversubscription. Other NSP VMs require all resources to be dedicated.

KVM configuration

When you configure the KVM, set the parameters listed in the following table to the required values.

Table 2-3 KVM configuration parameters

Parameter	Value
Disk Controller type	virtio
Storage format	raw
Cache mode	none
I/O mode	native
NIC device model	virtio
Hypervisor type	kvm
Video model	Virtio

CPU pinning may be used for NSP virtual machines (deployer and cluster nodes) but may restrict other KVM functions. Please consult KVM documentation for more information.

2.1.4 OpenStack virtualization

The NSP cluster is tested in an open-source OpenStack environment. Nokia supports NSP cluster

deployment on VMs provided in any OpenStack distribution that is based on the tested version. Any product issues deemed related to a specific distribution must be pursued by the customer and OpenStack vendor. The minimum supported OpenStack version for NSP cluster deployment is Newton. Nokia has not tested all supported Openstack versions with the NSP cluster. Nokia recommends that customers test and validate proper operation and performance of the NSP cluster as a whole before going to production.

To ensure NSP cluster compatibility with an OpenStack environment, you must follow the requirements in the following topics.

Hypervisor

KVM is the only supported hypervisor for an OpenStack environment. For information about the supported KVM hypervisor versions, see [2.1.3 “KVM virtualization” \(p. 25\)](#).

CPU and memory resources

The required memory resources must be reserved and dedicated to each guest OS, and cannot be shared or oversubscribed. You must set the `ram_allocation_ratio` parameter to 1.0 in the OpenStack Nova configuration on the control NE, or on each individual compute node that may host an NSP cluster VM. CPU resources may be oversubscribed using the guidelines for KVM CPU oversubscription in table 2-2 above.

Simultaneous Multi-threading (SMT)

The usage of CPUs with enabled SMT must be consistent across all compute nodes. If the CPUs do not support SMT, you must disable SMT at the hardware level on each compute node that may host an NSP cluster VM.

CPU pinning

CPU pinning is supported, but may restrict some OpenStack functions such as migration.

Availability zones/affinity/placement

Nokia does not provide recommendations for configuring VM placement in OpenStack.

Migration

The OpenStack environment supports only regular migration; live migration is not supported.

Networking

Basic Neutron functions using Open vSwitch with the ML2 plugin are supported in an NSP cluster deployment, as is the use of OpenStack floating IP addresses.

Storage

All storage must meet the throughput and latency performance criteria in the response to your NSP Platform Sizing Request.

VM storage

The VM storage must be persistent block (Cinder) storage, and not ephemeral. In order to deploy each VM, you must create a bootable Cinder volume. The volume size is indicated in the response to your NSP Platform Sizing Request.

Flavors

Flavors must be created for each Station Type described in the response to your NSP Platform Sizing Request.

Firewalls

You can enable firewalls using OpenStack Security Groups, or on the VMs using the firewall service, except as noted. If firewalld is enabled, an OpenStack Security Group that allows all incoming and outgoing traffic must be used.

ARP spoofing

OpenStack ARP spoofing protection must be disabled. MetalLB load balancer looks like an ARP spoofing attempt to OpenStack.

2.1.5 VMware virtualization

The NSP cluster supports using VMware vSphere ESXi on x86 based servers natively supported by ESXi. See the *Host Environment Compatibility Reference for NSP and CLM* for the current VMware compatibility level, requirements, and restrictions. See the VMware Hardware Compatibility List (HCL) to determine specific hardware support.

Not all features offered by ESXi are supported when using the NSP cluster. For example, Fault Tolerant, High Availability (HA), Memory Compression, Distributed Resource Scheduler (DRS), and vMotion features are not supported. VM snapshots are not supported when using the NSP cluster. Contact Nokia to determine if a specific ESXi feature is supported with an NSP installation.

CPU pinning is supported on NSP virtual machines (deployer and cluster nodes) but may restrict some VMware functions. Please consult VMware documentation for more information.

i **Note:** The system clocks of the NSP components must always be closely synchronized. The RHEL chronyd service is required as the time-synchronization mechanism to engage on each NSP component during deployment.

Caution: Some components - members of an etcd cluster, for example - will not trust the integrity of data if a time difference is detected. As such, failure to closely synchronize system clocks can complicate debugging and cause outages or other unexpected behavior.

Only one time-synchronization mechanism can be active in an NSP system. Before you enable a service such as chronyd on an NSP component, you must ensure that no other time-synchronization mechanism, for example, the VMware Tools synchronization utility, is enabled.

Virtual Machine Version 15, or 17 and above must be used.

See the following table for additional Virtual Machine setting requirements:

Table 2-4 Additional Virtual Machine setting requirements

Resource type	Parameter	Setting
CPU	Shares	Set to High
	Reservation	See VMWare CPU Reservation below
	Limit	Check box checked for unlimited
Memory	Shares	Set to High
	Reservation	Reserve all guest memory
	Limit	Check box checked for unlimited
Disk	Shares (moved to Storage Policies in 8.0)	Set to High
	Limit - IOPs (moved to Storage Policies in 8.0)	set to Unlimited
	Type	Thick Provision Eager Zeroed
SCSI controller	Type	VMware Paravirtual
Network Adapter	Type	VMXNET 3

VMWare CPU Reservation

CPU resources may be oversubscribed when the NSP cluster is deployed on VMWare. In the following table, the CPU Reservation is determined by: factor X * the number of CPUs * the CPU frequency. For example, with a factor of 0.5 on a 2.4 GHz 8 vCPU configuration, the reservation must be set to $(0.5 * 8 * 2400) = 9600$ MHz.

Table 2-5 CPU Reservation

NSP cluster deployment type	VMWare factor X
enhanced	0.25
standard	0.33
medium	0.4

i **Note:** CPU oversubscription is only supported on NSP clusters with three or more nodes. Smaller NSP cluster deployments must have reserved and dedicated CPU resources for each node. CPU oversubscription is only supported on Intel based processors using the Skylake or newer microarchitecture. Only the NSP Cluster VMs support oversubscription. Other NSP VMs require all resources to be dedicated.

2.2 NSP cluster requirements

2.2.1 Overview

The exact hardware and software requirements of an NSP deployment will vary based on a number of factors, including licensed feature packages and network size. An NSP cluster should be sized according to the platform requirements provided in your NSP Sizing Response, which you can request from your Nokia account representative. These platform requirements will adhere to the network dimensions described in [Chapter 5, “Scaling and performance”](#).

i **Note:** The NSP deployer is a crucial element of an NSP system that holds the required container images and Helm repositories for deployment to each NSP cluster VM. If the deployer host is unavailable, NSP cluster recovery in the event of a failure may be compromised. Ensure that the NSP deployer host remains operational and reachable by the NSP cluster after the cluster deployment.

i **Note:** Only NSP workloads are supported on the NSP Kubernetes Platform. Third-party workloads cannot be deployed on the NSP Kubernetes Platform.

2.2.2 Cluster sizing criteria

The platform requirements for NSP cluster VMs and VMs that host additional components depend on, but are not limited to, the following factors:

- number of managed LSPs and services
- number of managed NEs
- number of MDM learned services
- number of simultaneous operator and API sessions
- expected numbers of:
 - flows
 - monitored NEs
 - AS
 - ports with real-time statistics collection

2.2.3 NSP cluster VM hardware requirements

NSP deployments are server- and vendor-agnostic, but must meet all NSP component hardware criteria and performance targets. Server-class hardware is required; desktop hardware is inadequate. Processor support is limited to specific Intel Xeon-based x86-64 and AMD Epyc x86-64 CPUs that have the required minimum CPU core speed listed in [Table 2-6, “VM processor requirements”](#) (p. 30).

Table 2-6 VM processor requirements

Processor microarchitecture	Minimum CPU core speed	Supported deployments
Intel Xeon Haswell or Broadwell	2.4 GHz	Supported for all NSP deployments

Table 2-6 VM processor requirements (continued)

Processor microarchitecture	Minimum CPU core speed	Supported deployments
Intel Skylake or newer	2.0 GHz	Supported for all NSP deployments
AMD Epyc Zen 3 or newer	2.0 GHz	Supported on all NSP components except the following: <ul style="list-style-type: none"> • vCPAA • VSR-NRC

i **Note:** In newer Intel CPUs which support P-cores and E-cores, NSP only supports the use of P-cores.

Provisioned CPU resources are based upon threaded CPUs. The NSP Platform Requirements will specify a minimum number of vCPUs to be assigned to the VM. VMs are recommended to be configured with all vCPUs on one virtual socket.

A host system requires CPU, memory and disk resources after resources for NSP cluster VMs have been allocated. Contact the hypervisor provider for requirements and best practices related to the hosting environment.

You must provide information about the provisioned VMs to Nokia support. You can provide the information through read-only hypervisor access, or make the information available upon request. Failure to provide the information may adversely affect NSP support.

i **Note:** A virtual machine running a VSR-NRC instance requires CPU pinning and isolation from other virtual machines running on the host system. For platform requirements of the VSR-NRC, refer to the *VSR Installation and Setup Guide*.

2.2.4 NSP cluster storage-layer requirements

The storage layer of an NSP cluster requires a minimum read/write IOPS based on deployment type, network size and storage type. NSP cluster deployments with customer provided storage also have storage capacity requirements for database backups. See the tables below for IOPS and storage capacity requirements.

i **Note:** IOPS performance can vary in a cloud environment and may degrade over time. Degraded IOPS performance may affect NSP performance. Customers are advised to monitor IOPS performance over time.

Table 2-7 Minimum NSP cluster IOPS requirements - local storage

Deployment Type and Network Size	Minimum storage-layer read/write IOPS
Lab/trial	2000
Production: Fewer than 2000 NEs	2500

Table 2-7 Minimum NSP cluster IOPS requirements - local storage (continued)

Deployment Type and Network Size	Minimum storage-layer read/write IOPS
Production: More than 2000 NEs (maximum of 10 000 NEs)	3000

i **Note:** If the number of NEs in your network exceeds 10 000, please contact your Nokia account representative or refer to your NSP Sizing Response.

Table 2-8 Minimum NSP cluster IOPS requirements - customer provided storage

Deployment Type and Network Size	Minimum storage-layer read/write IOPS
All	10 000

Table 2-9 Customer provided storage capacity requirements for database backup storage

Component	Lab	Basic/Small	Medium	Standard	Enhanced
nspos-backup-storage	10GB	20GB	50GB	100GB	100GB

Refer to the *NSP Installation and Upgrade Guide* for file system requirements and disk partitioning recommendations.

i **Note:** The disk layout and partitioning of each VM in a multi-node NSP cluster, including DR deployments, must be identical.

2.2.5 VM memory

The virtual memory configuration of each NSP cluster VM requires a parameter change to support Centralized Logging.

The following command entered as the root user displays the current setting:

```
sysctl -a | grep "vm.max_map_count"
```

If the setting is not at least 262144, you must enter the following command before you deploy the NSP:

```
sysctl -w vm.max_map_count=262144
```

2.2.6 Time synchronization

The system clocks of the NSP components must always be closely synchronized. The RHEL chronyd service is required as the time-synchronization mechanism to engage on each NSP component during deployment.

Caution: Some components - members of an etcd cluster, for example - will not trust the integrity of data if a time difference is detected. As such, failure to closely synchronize system clocks can complicate debugging and cause outages or other unexpected behavior.

i **Note:** Only one time-synchronization mechanism can be active in an NSP system. Before you enable a service such as chronyd on an NSP component, you must ensure that no other time-synchronization mechanism, for example, the VMware Tools synchronization utility, is enabled.

2.2.7 nsp system user

The NSP cluster VMs require sufficient available user IDs to create system users for NSP applications.

NSP deployer nodes and Kubernetes worker nodes require a system user account that must have uid 1000. Often the "nsp" user is assigned to uid 1000, however the name for user ID 1000 is configurable.

2.2.8 Using hostnames

NSP deployments can use hostnames or FQDNs rather than IP addresses, but the hostname resolution must be through DNS.

i **Note:** Where NSP is deployed with hostnames or FQDNs in its configuration, reliable name resolution services is required for consistent and reliable NSP operations. Nameservers should be reachable, accept incoming requests from the NSP cluster and return valid responses for all cases.

i **Note:** Using multiple nameservers in `/etc/resolv.conf` is only intended for redundant nameservers and not for resolving different domains.

The hostname of an NSP component station must meet the following criteria:

- include only alphanumeric ASCII characters and hyphens
- not begin or end with a hyphen
- if an FQDN, FQDN components are delimited using periods
- hostname FQDN does not exceed 63 characters

2.3 Other NSP components hardware platform requirements

2.3.1 Overview

For all bare metal installations of the NFM-P server, NFM-P database, NFM-P statistics auxiliary, NSP flow collector controller, NSP flow collector controller, and NSP auxiliary database, Nokia requires the use of supported HP or Nokia AirFrame Intel based x86 stations 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. Contact the server vendor to determine the correct HBA controller for creation of the correct number of volumes, and to enable the read and write caches.

Redundant installations of NFM-P can use different stations 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.

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.

2.4 Other NSP components hardware platform and resource requirements using virtualization

2.4.1 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, NSP 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 a supported version of RHEL 8 server x86-64. For installations of the NFM-P client and NFM-P client delegate server on a Guest Operating System of a virtualized installation, the Guest Operating System can be either a supported version of RHEL 8 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 NSP.

Virtualized installations are server vendor agnostic but must meet specific hardware criteria and performance targets to be used. Server class hardware must be used, not desktops. Processor support is limited to specific Intel Xeon based x86-64 and AMD Epyc based x86-64 CPUs with a minimum CPU core speed as outlined in the proceeding tables. The Intel CPUs must be from the Haswell microarchitecture, or newer, where the CPU microarchitecture determines the minimum supported CPU speed. The AMD CPUs must from the Zen 3 microarchitecture, or newer. In newer Intel CPUs which support P-cores and E-cores, NSP only supports the use of P-cores.

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 NSP platform sizing responses. Performance must meet the documented requirements for both throughput and latency.

You must provide information about the provisioned VM to Nokia support. You can provide the information through read-only hypervisor access, or make the information available upon request. Failure to provide the information may adversely affect NSP support.

2.4.2 VMware virtualization

NSP supports using VMware vSphere ESXi on x86 based servers natively supported by ESXi. See the *Host Environment Compatibility Reference for NSP and CLM* for the current VMware compatibility level, requirements, and restrictions. See the VMware Hardware Compatibility List (HCL) to determine specific hardware support. Not all features offered by ESXi are supported when using NSP. 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 NSP 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 13, and 19 have been tested with the NFM-P server, NFM-P database, NFM-P statistics auxiliary, NSP auxiliary database, NSP flow collector controller, NSP flow collector, NFM-P client, and NFM-P client delegate, where the minimum supported Virtual Machine Hardware version is 10 and the latest supported version is 19.

See the following table for additional VM setting requirements.

i **Note:** The system clocks of the NSP components must always be closely synchronized. The RHEL chronyd service is required as the time-synchronization mechanism to engage on each NSP component during deployment.

Only one time-synchronization mechanism can be active in an NSP system. Before you enable a service such as chronyd on an NSP component, you must ensure that no other time-synchronization mechanism, for example, the VMware Tools synchronization utility, is enabled.

Table 2-10 VMware VM settings

Resource Type	Parameter	Setting
CPU	Shares	Set to High
	Reservation	Must be set to 1/2 the number of vCPUs * the CPU frequency. For example, on a 2.4GHz 8 vCPU configuration, the reservation must be set to (0.5*8*2400) 9600MHz
	Limit	Unlimited
Memory	Shares	set to High
	Reservation	Reserve all guest memory
	Limit	Unlimited

Table 2-10 VMware VM settings (continued)

Resource Type	Parameter	Setting
Disk	Shares (moved to Storage Policies in 8.0)	set to High
	Limit - IOPs (moved to Storage Policies in 8.0)	set to Unlimited
	Type	Thick Provision Eager Zeroed
	Sharing	No Sharing
	Disk Mode	Dependent
SCSI Controller	Type	VMware Paravirtual
Network Adapter	Type	VMXNET 3

2.4.3 VMware features

The following VMware features have been tested with the NFM-P server, NFM-P database, NFM-P statistics auxiliary, NSP auxiliary database. To ensure NSP 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 NSP 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
- VMs should be shutdown before taking snapshots
- NSP performance can be degraded by as much as 30% when a snapshot exists and therefore 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
- Do not consolidate snapshots when the server is active
- NFM-P database failover will occur when VMs are reverted to snapshots, requiring a re-instantiation of the standby database

-
- Supported on the NFM-P server, NFM-P database, and NFM-P statistics auxiliary

Distributed Resource Scheduler (DRS)

- Always follow VMware best practices
- Manual and partially automated DRS is supported
- Fully Automated DRS is not supported
- Supported on the NFM-P server, NFM-P database, and NFM-P statistics auxiliary

2.4.4 KVM virtualization

NSP supports using RHEL 6, RHEL 7, RHEL 8, and RHEL 9 based KVM on x86 based servers natively supported by KVM. The Host Environment Compatibility Reference for NSP and CLM should be consulted for up-to-date KVM compatibility along with requirements and restrictions. See the RHEL Hardware Compatibility List (HCL) 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.

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”, and the I/O mode set to “native”. The NIC device model must be “virtio”. The hypervisor type must be set to “kvm”.

2.4.5 OpenStack

NSP 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 NSP 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 NSP could reside.

Simultaneous Multi-threading (SMT)

- SMT CPU usage must be consistent across all compute nodes. If there are CPUs that do not

support SMT, SMT 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. See the 7701 CPAA Installation Guide for vCPAA CPU Pinning requirements

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 NSP deployment. OpenStack floating IP address functionality can be used on specific interfaces used by NSP 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 NSP Platform Sizing 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 your NSP Sizing Response.

Flavors

- Flavors should be created for each “Station Type” indicated in your NSP 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.

2.5 Other NSP components minimum platform requirements

2.5.1 Minimum hardware platform requirements

The following tables specify the minimum hardware platform requirements necessary to successfully operate the various NSP components in a bare metal configuration.

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

2.5.2 Bare metal hardware configurations

Table 2-11 NFM-P bare metal minimum collocated platform

For networks not exceeding: <ul style="list-style-type: none"> • 675 equivalent MDAs • 1000 GNEs • 5 simultaneous NFM-P clients (GUI or XML-API) • 3000 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* Intel x86 CPU Cores (12 Hyper-threads), minimum 2.0GHz ¹ 46 GB RAM recommended (37 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 2-12 NFM-P bare metal minimum distributed platform

For networks not exceeding: <ul style="list-style-type: none"> • 1,875 equivalent MDAs and 5 simultaneous NFM-P clients (GUI or XML-API) OR <ul style="list-style-type: none"> • 1,275 equivalent MDAs and 25 simultaneous NFM-P clients (GUI or XML-API) • Maximum of 5000 GNEs • 6000 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* Intel x86 CPU Cores (12 Hyper-threads), minimum 2.0GHz ¹ 41 GB RAM recommended (33 GB RAM minimum). 2*10K RPM SAS disk drives of at least 300 GB each in size
NFM-P database	4* Intel 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.

2.5.3 Minimum VM resource requirements

The following tables list the minimum VM resource requirements for deployments of various NSP components using VMware vSphere ESXi or RHEL KVM.

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

2.5.4 VM minimum collocated resource requirements

Table 2-13 NFM-P VM minimum collocated configuration

For networks not exceeding: <ul style="list-style-type: none"> • 675 equivalent MDAs • 1000 GNEs • 5 simultaneous NFM-P clients (GUI or XML-API) • 3000 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 46 GB RAM recommended (37 GB RAM minimum) 800 GB disk space I/O requirements found in 2.7 "Storage" (p. 47)

Notes:

1. 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture and AMD Epyc Zen 3 microarchitecture. For Intel CPUs, the 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.

2.5.5 VM minimum distributed resource requirements

Table 2-14 NFM-P VM minimum distributed configuration

For networks not exceeding: <ul style="list-style-type: none"> • 1,875 equivalent MDAs and 5 simultaneous NFM-P clients (GUI or XML-API) OR <ul style="list-style-type: none"> • 1,275 equivalent MDAs and 25 simultaneous NFM-P clients (GUI or XML-API) • Maximum of 5000 GNEs • 6000 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 41 GB RAM recommended (33 GB minimum) 500 GB disk space I/O throughput and latency as provided in NFM-P sizing response
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 Intel Xeon Skylake and newer CPU microarchitecture and AMD Epyc Zen 3 microarchitecture. For Intel CPUs, the 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.

2.5.6 Scaling limits for collocated configurations

Collocated NFM-P 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 2-15 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	1000
Performance statistics per 15 minute interval	50 000

Table 2-15 Scaling limits for collocated configurations (continued)

Scaling parameter	Maximum
Accounting statistics per 15 minute interval	200 000
TCA's	50 000

2.5.7 Minimum platform requirements for NFM-P auxiliary collectors

Table 2-16 NFM-P auxiliary platforms - Bare Metal

Architecture	Supported NFM-P auxiliary type	Configuration
Bare Metal x86	statistics collector	4 * Intel x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz ¹ 12 GB RAM minimum. 16 GB RAM is recommended. 4*10K RPM SAS disk drives of at least 300 GB each in size

Notes:

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

Table 2-17 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. 16 GB RAM is recommended. 600 GB disk space I/O throughput and latency as provided in NFM-P sizing response

Notes:

- 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture and AMD Epyc Zen 3 microarchitecture. For Intel CPUs, the minimum speed on CPUs older than Skylake is 2.4GHz.

2.5.8 Minimum platform requirements for NSP auxiliary database

Table 2-18 NSP auxiliary database platform - single node cluster

Architecture	Configuration
Bare Metal x86	4 * Intel x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz ¹ 64 GB RAM minimum. 2 SAS 10K RPM disk drives of at least 300 GB 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)

Table 2-18 NSP auxiliary database platform - single node cluster (continued)

Architecture	Configuration
VMware/KVM	8 vCPUs, minimum 2.0GHz ² 64 GB RAM minimum. 500+ GB disk space I/O throughput as measured with the vioperf utility

Notes:

- 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture. Minimum speed on CPUs older than Skylake is 2.4GHz.
- 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture and AMD Epyc Zen 3 microarchitecture . Minimum speed on Intel CPUs older than Skylake is 2.4GHz.

Table 2-19 NSP auxiliary database platform - three+ node cluster

Architecture	Configuration (for each node of the auxiliary database cluster ¹)
Bare Metal x86	12 * Intel x86 CPU Cores (24 Hyper-threads), minimum 2.0GHz ² 128 GB RAM minimum. 2 SAS 10K RPM disk drives of at least 300 GB each in size (RAID 1) + 12 SAS 10K RPM disk drives of at least 600 GB 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 2.0GHz ³ 128 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:

- Minimum of three nodes required in the cluster.
- 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture Minimum speed on CPUs older than Skylake is 2.6GHz.
- 2.0GHz only supported on Intel Xeon Skylake and newer CPU microarchitecture and AMD Epyc Zen 3 microarchitecture. Minimum speed on Intel CPUs older than Skylake is 2.6GHz.

2.5.9 Platform requirements for NFM-P client delegate server stations

NFM-P allows multiple clients to be installed on a single HP x86 or Nokia AirFrame station running RHEL 8 server x86-64, or specific versions of Windows. This option enables customers to launch multiple NFM-P clients from a single station. 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 server platform provides an option to consolidate multiple installations of the NFM-P client onto a single station where one instance of the NFM-P client is run by many users (with unique Operating System accounts).

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 [4.13 “Network element specific requirements” \(p. 71\)](#).

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 2-22, “Element manager operating system support summary” \(p. 45\)](#) lists these element managers and their operating system support. See the element-manager documentation to determine current operating system support.

The NFM-P client delegate server configuration is only supported on specific HP or Nokia AirFrame x86 stations running RHEL server x86-64, or specific versions of Windows. Additionally, the NFM-P client delegate server 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. [Table 2-20, “Minimum NFM-P client delegate server resource requirements” \(p. 43\)](#) describes resource requirements for this type of station.

Table 2-20 Minimum NFM-P client delegate server resource requirements

Architecture	Configuration
Bare Metal x86	4* Intel x86 CPU Cores (8 Hyper-threads), minimum 2.0GHz 28 GB RAM minimum, 36 GB for networks with greater than 15 000 NEs 1*10K RPM SAS disk drive, 146 GB in size
VMware/KVM	8 vCPUs, minimum 2.0GHz 28 GB RAM minimum, 36 GB for networks with greater than 15 000 NEs 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 2-21, “Additional client NFM-P client delegate server resource requirements” \(p. 44\)](#) are added to the platform.

Table 2-21 Additional client NFM-P client delegate server resource requirements

Architecture	Additional resources per client
Bare Metal x86	1/4 * Intel 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 1 GB Disk Space
VMware/KVM	1/2 vCPU, minimum 2.0GHz 1.75 GB RAM, 2.25 GB for networks with greater than 15.000 NEs 1 GB Disk Space

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

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 requires extra system resources that will need to be added to those that are required by the NFM-P delegate. See the Citrix documentation to determine the additional Citrix resource requirements.

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

- 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 server 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.

2.5.10 NFM-P client platform requirements

Nokia recommends having 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 [4.13 “Network element specific requirements” \(p. 71\)](#) .

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. See the element-manager documentation to determine current operating system support.

Table 2-22 Element manager operating system support summary

Element manager	Node type	RHEL 8 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 2-23 NFM-P client hardware platform requirements

NFM-P client hardware platform requirements	
RHEL platforms	Microsoft Windows
1 Intel 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 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 Intel 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 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.

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 2-24 Virtualized NFM-P client resource requirements

VM 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)

2.6 NFM-P platform requirements for larger networks

2.6.1 Determining platform requirements for larger networks

NFM-P may require larger stations in order to successfully manage networks that exceed any of the dimensions supported by the minimum hardware platforms. In order to determine station 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 your Nokia account 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).

2.7 Storage

2.7.1 Storage overview

This section provides information about configuring stations that will host NSP software.

Specific partition sizes and configuration procedures are available in the *NSP 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 NSP. 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 file systems, including Samba are not supported.

While Nokia identifies areas of the storage that are not specifically required for NSP and are partitionable for customer use, station resources are expected to be dedicated for NSP. As such, these “Remainder” portions of the storage 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 stations running the NFM-P database in a bare metal configuration. 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 stations 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.

In NSP upgrade scenarios, previous disk configurations may still be valid.

2.7.2 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 Installation and Upgrade Guide*. The increased disk I/O performance offered by RAID 0 is required for all NSP deployments. The *NSP Installation and Upgrade Guide* provides details of these configurations. A RAID 0 stripe size of 512 Kbytes is required for optimal NSP 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 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 using 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 NSP auxiliary database backup partition is supported with RAID 5.

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

2.7.3 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 NSP without interruption in a low latency environment.

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

i **Note:** Nokia is not responsible for installation, administration or recovery of SANs on an NSP platform.

2.7.4 Virtualization I/O requirements

When using NSP 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, SAS or NVMe based SSDs), 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 NSP Platform sizing responses. In VMware environments, storage I/O shares must be set to “High” and IOPs set to “Unlimited” for best performance and low latency.

See [Table 2-25, “Minimum collocated configuration throughput and latency” \(p. 49\)](#) 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 your NSP Sizing Response that is required for every NSP deployment.

Specific NSP components 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 NSP. 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 2-25 Minimum collocated configuration throughput and latency

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

See the *NSP Installation and Upgrade Guide* for the recommended partition sizes.

3 Operating system specifications

3.1 Red Hat Enterprise Linux (RHEL)

3.1.1 NSP deployment on RHEL

The following topics define the operating system (OS) requirements for deploying an NSP cluster and optional NSP components.

3.1.2 RHEL version support

NSP supports the following RHEL server x86-64 versions in an NSP cluster deployment (cluster node and deployer node):

- RHEL server 9 x86-64 - Update 6 (9.6)

Other NSP components, including NFM-P, are supported on RHEL server x86-64 versions 8.6 through 8.10. (See NSP operating system summary below.)

Previous releases, or other variants of Red Hat, and other Linux variants are not supported.

Customers must ensure that NSP supports each update applied to RHEL OS in an NSP deployment. An automated update by a subscription manager may deploy an unsupported RHEL version. In order to prevent the accidental deployment of an unsupported RHEL version, Nokia recommends to lock the supported release in RHEL subscription manager.

For the latest compatibility information, see the *Host Environment Compatibility Reference for NSP and CLM*.

Nokia provides RHEL OS images based upon RHEL 9.6 and 8.10 and can only be used for the deployment of NSP, and not for the deployment of any other Nokia or third-party product. VMs created from the Nokia-provided NSP RHEL OS images can only be updated with the Nokia-provided NSP RHEL OS updates.

The RHEL OS must be installed in English.

The NSP does not necessarily support all functionality provided in RHEL. The NSP does not support the `requiretty` option in `/etc/sudoers`. The RHEL `chronyd` service is required as the time-synchronization mechanism on NSP cluster nodes. The NSP also requires that the server hostname is configured in the `/etc/hosts` file. RHEL must be installed in 64-bit mode where the NSP will be installed. The NSP product team does not support the configuration of OS services not enabled by default in the NSP RHEL OS image and/or required by NSP applications.

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

With the exception of documented operating system parameter changes for NSP, all other settings must be left at the RHEL default configuration.

The *NSP Installation and Upgrade Guide* provides detailed instructions for the RHEL OS installation.

3.1.3 Red Hat support

For customers using the NSP RHEL OS 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 the RHEL server with NSP. It is strongly recommended to purchase a support package from Red Hat that provides 24x7 support.

The NSP RHEL OS 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.

3.1.4 Third-party applications

Applications that are not sanctioned by Nokia must not be running on any virtual instance running NSP components. Nokia reserves the right to remove any applications that are suspected of causing issues from stations running NSP components.

3.2 NFM-P on Microsoft Windows

3.2.1 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.

i **Note:** The following table documents the basic support as of the current release of NSP. For the latest compatibility information, which could include off-cycle updates, see the *Host Environment Compatibility Reference for NSP and CLM*.

Table 3-1 Windows operating system support summary

Microsoft Windows version	NFM-P client	NFM-P client delegate server
Windows 8 / 8.1 Enterprise (64-bit)	Supported	Not supported
Windows 10 Professional	Supported	Not supported
Windows 11 Professional	Supported	Not supported
Windows Server 2016	Supported	Supported
Windows Server 2019	Supported	Supported
Windows Server 2022	Supported	Supported

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

3.2.2 Microsoft support

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

3.3 NFM-P on Apple macOS

3.3.1 NFM-P support

The macOS operating system is only supported for NFM-P clients. macOS 15 (Sequoia) has been tested with NFM-P.

3.3.2 Apple support

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

3.4 NSP operating system summary

3.4.1 Operating system summary

The following table summarizes the supported configurations for each of the operating systems supported by NSP.

Table 3-2 NSP operating system support summary

NSP Component	RHEL 8 server x86-64	RHEL 9 server x86-64	Microsoft Windows	macOS
NSP Cluster	Not supported	9.6	Not supported	Not supported
Deployer Node	Not supported	9.6	Not supported	Not supported
NFM-P server	8.6 through 8.10	Not supported	Not supported	Not supported
NFM-P database	8.6 through 8.10	Not supported	Not supported	Not supported
Collocated NFM-P server/database	8.6 through 8.10	Not supported	Not supported	Not supported
NFM-P client	8.6 through 8.10	Not supported	Supported	Supported
NFM-P auxiliary	8.6 through 8.10	Not supported	Not supported	Not supported
NSP auxiliary database	8.6 through 8.10	Not supported	Not supported	Not supported
NFM-P client delegate server	8.6 through 8.10	Not supported	Supported	Not supported

3.5 Third-party software for NFM-P single-user client or client delegate server

3.5.1 NFM-P client or client delegate server software requirements

NFM-P clients are launched, installed, and uninstalled by the direct installer download. 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 server can only support a single NEtO instance running among all clients connected to a client delegate server 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.

4 Network requirements

4.1 NSP network requirements

4.1.1 Introduction

This chapter describes networking requirements for an NSP deployment.

4.2 NSP deployment network addressing requirements

4.2.1 NSP MetalLB load balancer

MetalLB is a load balancer for Kubernetes clusters that supports three ingress applications: ingress controller, trap forwarder, flow forwarder. These ingress applications require a pool of IP addresses.

- All NSP clusters require an ingress controller client VIP for client traffic. An internal VIP is used when the client and internal traffic are on different network interfaces. A mediation VIP is used when there is a dedicated mediation interface and network elements require access to NSP server applications such as file server app
- an NSP cluster with MDM deployed requires a trap forwarder mediation VIP for mediation traffic. (See Note below.)
- an NSP cluster with Flow Collector deployed requires a flow forwarder VIP for mediation traffic.

i **Note:** For SR OS releases 23.10 R1 or later, SNMP is not required. Refer to *Nokia SR OS Artifact Guide* for further information. Older SR releases and some multi-vendor NEs will require a trap forwarder mediation VIP for SNMP traps.

i **Note:** Each VIP specified must be distinct from any other VIP or IP addresses on the cluster nodes.

When installing the NSP Kubernetes cluster, you must specify which NSP cluster nodes will act as load balancer endpoints. Such nodes are called ingress gateways, and can accept ingress application traffic, e.g. GUI client connections or SNMP traps from network elements. There should be at least two ingress gateway nodes for redundancy and they should belong to different host stations. Each ingress application will accept traffic on one ingress gateway node. Different ingress applications could accept traffic on different ingress gateway nodes.

Each NSP cluster node that is an ingress gateway must have a network interface address that belongs to the same subnet for each VIP configured to the cluster.

4.2.2 NSP cluster nodes in same subnet

All NSP cluster nodes must belong to the same subnet when configuring the NSP kubernetes cluster. The `nodeIP` parameters in `k8s-deployer.yml` must be within the same subnet. This applies to NSP cluster deployments that use a single network interface, or multiple network interfaces. This requirement is independent of any restrictions around virtual IP addresses used in an NSP cluster deployment.

4.2.3 Using IPv4 and IPv6 in NSP deployments

The NSP supports IPv4 and IPv6 network connectivity with client applications and other components in the NSP architecture.

NSP can be deployed on VMs with dual stack network interfaces within any restrictions described in this document.

The following limitations and restrictions apply to deploying the NSP with IPv6 network communications:

- The deployer host for an NSP cluster must have IPv4 connectivity to NSP cluster nodes. The NSP cluster can be configured for IPv6 communications for NSP applications, but must have IPv4 connectivity to the deployer node.
- Common web browser applications have security policies that may prevent the use of bracketed IPv6 addresses in the URL browser bar. Customers that use IPv6 networking for client communications to the NSP must use the hostname configuration for NSP components. (The hostname of an NSP cluster is defined in `nsp-config.yml` with parameter `platform.ingressApplications.ingressController.clientAddresses.advertised`. The hostname of a NFM-P server is defined in `samconfig` with parameter `client.hostname`. See [2.2.8 “Using hostnames”](#) (p. 33) for more information.
- All NSP components in an NSP deployment must use IPv4 or IPv6 for inter-component communications. Different integrated components in an NSP deployment cannot communicate between IPv4 and IPv6 interchangeably (for example, if the NSP is deployed with IPv6, then the NFM-P also needs to be deployed with IPv6).
- The NSP Kubernetes cluster communications uses internal addressing in 10.233.0.0/18 subnet. Customers should avoid using this subnet in their NSP deployment on VM network interfaces.
- WS-NOC does not support IPv6 deployment. An integrated deployment of NSP with WS-NOC must be deployed with IPv4 addressing.
- When there is an IPv4 interface and an IPv6 interface on the same switch, each interface needs its own VPLS.

The NSP can be deployed with multiple network interfaces using IPv4 and IPv6 addressing. Chapter 7 of this guide documents the requirements and limitations of a multiple network interface NSP deployment.

i **Note:** When a hostname resolves to multiple IP addresses (eg IPv4 and IPv6), any IP addresses that are not used by the NSP deployment must not be defined in the system `/etc/hosts` file.

VSR-NRC

For NSP to VSR-NRC communications, both IPv4 and IPv6 are supported. These protocols are also supported in communications between VSR-NRC and PCCs.

4.3 Network requirements between NSP and other components

4.3.1 NSP cluster and OSS clients

The bandwidth requirements between NSP and OSS clients depend on the number of concurrent connections and on the type of transactions that are performed. For a single provisioning thread, Nokia recommends providing 50 Kbps of bandwidth from the OSS client to the NSP cluster. An OSS client that performs frequent query operations (for example, port or service inventory) must be provided with additional bandwidth.

4.3.2 NSP cluster and GUI clients

The bandwidth requirements between the NSP cluster and GUI clients mostly depend on the size of the network. A larger network with additional nodes and services requires more data to download to GUI clients. Optimal GUI performance is achieved with 10 Mbps of bandwidth with minimal network latency. Nokia recommends providing a minimum of 2.5 Mbps of bandwidth.

High network latency between the NSP cluster and GUI clients slows GUI performance. Nokia recommends limiting the round-trip network latency time to 100 ms.

4.3.3 NSP cluster and NFM-P

The bandwidth requirements between the NSP cluster and NFM-P depend on the following factors:

- the number of NEs, LSPs, and services configured on the NFM-P
- the frequency of NE updates to the NSP cluster

When the NSP cluster resynchronizes with NFM-P, optimal performance is achieved with 50 Mbps of bandwidth between the NSP cluster and NFM-P. Nokia recommends providing a minimum of 25 Mbps of bandwidth.

Network latency impacts the time it takes for the NSP cluster to resynchronize a large amount of data from the NFM-P. Nokia recommends limiting the round-trip network latency time to 100 ms.

4.3.4 NSP cluster and network elements

When NSP or NFM-P is managing nodes via SNMP, the mediation network needs to support SNMP packet size of 9216 bytes. Refer to the *NSP Troubleshooting Guide* for more details on packet size and fragmentation in the mediation network.

When NSP is only managing MD managed nodes with other mediation protocols, there are no specific IP MTU requirements.

4.4 Network requirements for NSP redundancy and communications within an NSP cluster

4.4.1 Communication requirements between redundant NSP deployments

The network requirements between active/standby NSP clusters depend on the network size (number of NEs and configured services) and the rate of service provisioning activities. The peak

bandwidth requirement between redundant servers is 50 Mbps, with a sustained bandwidth of 25 Mbps. Round-trip network latency between the pair of redundant servers must be limited to 100 ms.

4.4.2 Communication between nodes in an NSP cluster

In a multi-node deployment of an NSP cluster, it is recommended that the nodes within the cluster have 1Gbps of Ethernet connectivity with less than 1 ms round-trip latency.

4.5 NSP Support for RHEL IP Bonding

4.5.1 RHEL IP Bonding

Nokia supports the deployment of NSP using the RHEL IP Bonding feature. Support for IP Bonding is intended only to provide network interface redundancy that is 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. See the RHEL documentation for information about configuring IP Bonding.

4.6 NFM-P network requirements

4.6.1 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 the NFM-P stations and the connection to the network being managed. Nokia recommends making sufficient bandwidth available to the NFM-P stations within the Data Communication Network.

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

i **Note:** Ensure that you consider 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 [7.3 “NFM-P multihoming” \(p. 158\)](#) for information about configuring the NFM-P components with multiple interfaces.

4.7 Network elements

4.7.1 Network element connectivity support

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

- 9500 MPR / Wavence SM
- 7950 XRS
- 7750 SR

- 7705 SAR
- 7705 SAR-Hm
- 7450 ESS
- 7250 IXR
- 7210 SAS
- OmniSwitch 6350, 6465, 6560, 6865
- vCPAA

The 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.

4.8 Bandwidth requirements

4.8.1 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, and the number of network elements being managed.

Table 4-1 NFM-P collocated server/database bandwidth requirements

Available bandwidth required from a primary NFM-P server/database station	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 a primary and a standby NFM-P server/database station 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)

4.8.2 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, and the number of network elements being managed.

Table 4-2 NFM-P distributed server/database bandwidth requirements

Available bandwidth requirements for NFM-P	Recommended bandwidth, excluding statistics bandwidth requirements
NFM-P server to an NFM-P database Note: This depends on GUI changes and lists, the number of changes occurring in the network, and network objects being managed.	5-10 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 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 4-3 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 station	1.5 Mbps
Between the NFM-P database stations – 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)

4.8.3 Additional bandwidth requirements for NFM-P 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 stations 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 requires additional CPU requirements on the station that is collecting the statistics (NFM-P server or NFM-P auxiliary statistics collector). In this case, the bandwidth required will be reduced.

Table 4-4 Additional bandwidth requirements for accounting statistics collection

Record storage location	Bandwidth between the collecting server and NFM-P database	Bandwidth between the collecting server and NSP auxiliary database	Bandwidth between the NFM-P databases	Bandwidth between the NSP auxiliary database clusters	Bandwidth between the 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
NSP auxiliary database	N/A	2.2 Mbps	N/A	0.8 Mbps per NSP 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 4-5 Additional bandwidth requirements for application assurance accounting statistics collection

Record storage location	Bandwidth between the collecting server and NFM-P database	Bandwidth between the collecting server and NSP auxiliary database	Bandwidth between the NFM-P databases	Bandwidth between the NSP auxiliary database clusters	Bandwidth between the NFM-P databases during re-instantiation	Bandwidth between the collecting server and XML API client
NFM-P database	3.1 Mbps	N/A	4.2 Mbps	N/A	20 Mbps	N/A
NSP auxiliary database	N/A	3.1 Mbps	N/A	0.8 Mbps per NSP 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 4-6 Additional bandwidth requirements for performance statistics collection

Record storage location	Bandwidth between the collecting server and NFM-P database	Bandwidth between the collecting server and NSP auxiliary database	Bandwidth between the NFM-P databases	Bandwidth between the NSP auxiliary database clusters	Bandwidth between the NFM-P databases during re-instantiation	Bandwidth between the collecting server and XML API client
NFM-P database	5.4 Mbps	N/A	14.4 Mbps	N/A	72 Mbps	N/A
NSP auxiliary database	5.4 Mbps	5.4 Mbps	14.4 Mbps	0.8 Mbps per NSP 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

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 locations A and B must be the sum of the bandwidth requirements between the NFM-P auxiliary statistics collector to the NFM-P server and the NFM-P server to the 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 within the same geographic site.

4.8.4 NSP auxiliary database

When an NSP 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 NSP auxiliary database for statistics collection.

When the auxiliary database cluster is deployed with a minimum of three nodes, the NSP 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 NSP auxiliary database servers in the cluster. The interface for internal data communication needs to be dedicated with a minimum interface speed of 1 Gbps and part of a private network.

Table 4-7 Additional bandwidth requirements for the NSP auxiliary database

Bandwidth requirements for installations with the NSP auxiliary database	Bandwidth usage characterization
NFM-P auxiliary statistics collector to the NSP auxiliary database cluster	Higher bandwidth to write statistics data into the NSP auxiliary database cluster

Table 4-7 Additional bandwidth requirements for the NSP auxiliary database (continued)

Bandwidth requirements for installations with the NSP auxiliary database	Bandwidth usage characterization
NSP analytics server to NSP auxiliary database cluster Note: A higher bandwidth may be desirable.	Higher bandwidth to generate reports based upon raw and aggregated data
NSP auxiliary database node to the NSP auxiliary database node (intra-cluster)	High - must use a dedicated, minimum 1 Gbps interface
NSP auxiliary database to the NSP auxiliary database (redundant cluster)	High - up to 500 Mbps

4.9 Contributors to bandwidth requirements

4.9.1 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 the information visible to the user based on recent changes in the network.

A few examples of network changes which will be reported to the 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, and configuration changes of routing protocols.

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.

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

Additional bandwidth may be required to support the network elements described in [4.13 “Network element specific requirements” \(p. 71\)](#)

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

4.9.3 XML API clients

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

- 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 the NFM-P XML API or applications which retrieve large pieces of information via the API, such as statistics information or network inventory information, require access to a dedicated bandwidth from the machine hosting the application to the NFM-P server according to the tables provided above. Applications which do not require real-time event and alarm notifications 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.

4.9.4 NFM-P auxiliary statistics collector

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

- 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 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 requires more significant bandwidth utilization between the NFM-P auxiliary statistics collector and the NFM-P database stations. The bandwidth requirements are not dependent on the network activity.

4.9.5 NSP Flow Collector

The main factors impacting communication to and from the NSP Flow Collector are the following:

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

Table 4-8 Additional bandwidth requirements for the NSP Flow Collector

Bandwidth requirements for an NSP Flow Collector	Bandwidth usage characterization
Managed Network to the 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 the 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 station where IPDR files are being transferred.	Use the information on the left to calculate the amount of data generated for the expected statistics and to calculate the time to transfer at a given bandwidth. The total time must be less than 50% of the collection interval. For example, if 1 GB of IPDR files is expected per interval, and the collection interval is 5 minutes, a 45 Mbps connection will take 3minutes and 2 seconds to transfer. This is more than 50% and a larger network connection is required.
NFM-P server 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 remain in sync. The amount of bandwidth required is dependent on the network size.	Bandwidth requirement will depend upon the network size, which determines the network extraction file size, and the desired time to complete the file transfer from the NFM-P server to the NSP Flow Collector Controller.

4.10 Network bandwidth

4.10.1 Bandwidth requirements

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

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

Table 4-9 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)	200-400 Kbps

Table 4-9 NFM-P server to network bandwidth requirements (continued)

Network element example	Bandwidth requirement from NFM-P server(s) to the network element
7250 IXR-6 / 7250 IXR-R4 / 7250 IXR-R6 / 7250 IXR-R6d / 7250 IXR-R6dl / 7250 IXR-x	800-1000 Kbps
7250 IXR-e / 7250 IXR-e2	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, 6865, 9000 Series	300 Kbps
OmniSwitch 6860, 6860E, 6860N, 6900, 10K	400 Kbps
1830 VWM OSU	400 Kbps

4.10.2 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 the 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 the 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 the 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. However, several of the SNMP traps do not contain all of the information required for the NFM-P to completely represent the new status of the network element. As a result, the 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. However, in the worst case, this bandwidth utilization will be less than that utilized during a network element discovery.
- SNMP polling: It is possible to configure the NFM-P to poll the SNMP MIBs on the network elements at various intervals. By default, the 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 the 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 the 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 the 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 the NFM-P to request a backup of the network element at a specified interval. During the NE backup cycle, the amount of data transferred between the 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 the 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 4-9, “NFM-P server to network bandwidth requirements” \(p. 65\)](#). 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, the management of the NE includes methods other than the standard MIB/SNMP management, for example, web-based tools. These network elements may require additional bandwidth above the bandwidth levels stated in [Table 4-9, “NFM-P server to network bandwidth requirements” \(p. 65\)](#).

4.10.3 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 an 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 the NFM-P because of congestion will increase. This is significant since the NFM-P will detect it has missed traps from the network element and will result in the 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.

4.10.4 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 4-9, “NFM-P server to network bandwidth requirements” \(p. 65\)](#).

The total amount of bandwidth that is required for the 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 the perspective of NFM-P, there must be sufficient bandwidth (according to [Table 4-9, “NFM-P server to network bandwidth requirements”](#) (p. 65)) between itself and each of the NEs it manages.

In cases where the management traffic is carried over physical point-to-point links between the NFM-P server and the 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
- retrieving STM test results, also via (s)FTP
- NE reachability checks and NE trap gap checks

It would be rare for all of the operations listed above to be performed simultaneously; therefore, in normal circumstances, the NFM-P can communicate with at least 20 to 30 NEs simultaneously. This minimum can increase to 60 or 70 NEs on a 16 CPU core NFM-P server station. For networks of over 1000 NEs or where an NFM-P auxiliary statistics collector is being used, that number should be increased by 20 to 30 NEs. A higher bandwidth may be required under special cases where above average data is attempted to be transferred between the NFM-P and the network elements. For example, large statistics files, NE backups, or software images.

4.11 Network latency

4.11.1 Network latency considerations

Network latency can potentially impact the performance of the 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 [5.7 “Scaling guidelines for statistics collection”](#) (p. 94) for more information about the impact of latency on SNMP stats collection).
- NFM-P server and NFM-P auxiliary collectors to the 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, the overall NFM-P performance will be significantly impacted. The NSP 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, if a specific operation required the NFM-P to send 1000 SNMP gets to a single network element, the NFM-P will spend a total of 5 seconds sending and receiving packets when the round-trip latency to the network element is 5 ms. The time that NFM-P spends sending and receiving the same packets would increase to 50 seconds if the round-trip latency was increased to 50 ms.

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

The 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, on the releases that support streaming:

- 7950 XRS
- 7750 SR
- 7450 ESS
- 7250 IXR

4.11.2 Geographical redundancy of NFM-P components

It is ideal to ensure that all NFM-P stations 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 stations (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 1 ms. When an NFM-P component (server, auxiliary, or database) switchover or failover occurs, manual intervention may be required to align the stations 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 the NFM-P.

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

4.11.3 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 10 Mbps and round-trip latency is higher than 25 ms.

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

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

See the RHEL documentation for information about 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 that the current amount of system memory can support the expected number of network connections with the new buffer size.

4.11.4 Additional NFM-P database throughput optimizations

In addition to the optimizations above, the NFM-P database station 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.

4.12 Network reliability

4.12.1 Network reliability considerations

This section describes network reliability considerations.

4.12.2 Reliability between NSP components

NSP requires reliable network communications between all the NSP components. The performance and operation of NSP can be significantly impacted if there is any measurable packet loss between the NSP components. A significant packet loss can cause NSP reliability issues.

Nokia supports the deployment of NSP using the RHEL IP Bonding feature. The support for IP Bonding is intended only to provide network interface redundancy that is 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. See the RHEL documentation for information about configuring IP Bonding.

4.12.3 NFM-P server to NE network reliability

The NFM-P server requires reliable network connectivity between the NFM-P server/auxiliary and 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 the performance of the NFM-P's ability to manage the network elements. The loss of packets between the NFM-P and NE will have an impact on (but not limited to) the following:

- 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 the 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 timeout of 10 seconds, and an average round-trip latency of 50 ms between the 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 1000), the NFM-P server will wait for the retry timeout (10 seconds) to expire before re-transmitting. 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 a 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%.

4.13 Network element specific requirements

4.13.1 GNE, Nokia OmniSwitch, and 9500 / Wavence considerations

NFM-P clients support the web-based WebView functionality on the 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.

The 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. See the NE 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.

4.14 Mechanism to maintain the current state of network elements

4.14.1 Overview

The 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 XRS, 7750 SR, 7450 ESS, 7450 SR, 7210 SAS, and 7705 SAR network elements and the NFM-P network element interaction layers.

4.14.2 IP connectivity (ping) verification

The 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.

4.14.3 SNMP connectivity verification

The NFM-P performs an SNMP communication check every 4 minutes. If the NFM-P cannot communicate via the SNMP with a network element, it will raise a communications alarm against that network element. The NFM-P will also color the network element as red on the map to indicate the communication problem. Then, the NFM-P will clear the alarm and color the network element as green once the NFM-P detects that 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.

4.14.4 SNMP traps

The NFM-P listens to SNMP traps to receive changes from the network elements. The 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 the 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.

4.14.5 SNMP trap sequence verification

The 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. The 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, the NFM-P will request only the missing traps from the network element. If at any point the 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, the 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.

4.14.6 Scheduled SNMP MIB polling

The 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.

4.14.7 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. See [4.14.5 “SNMP trap sequence verification” \(p. 73\)](#) for a description of the NFM-P behavior with regards to trap handling. When a network outage occurs, the network element configuration in the NFM-P will be made consistent with the network element, but any event notifications, such as SNMP traps, that have occurred during the network outage will not be processed. This will cause intermediate state change alarms to not be reflected in the NFM-P during the network outage.

5 Scaling and performance

5.1 NSP scaling and performance

5.1.1 Introduction

The following sections present the network dimension parameters for NSP platforms that follow system requirements as described in this guide and any NSP Sizing Responses that you've requested.

i **Note:** NSP deployments managing a large number of NEs may need several hours (potentially a day for larger networks) to resync the network following an NSP upgrade. Customers should be prepared for post-upgrade NE resyncs when planning NSP upgrade activities.

5.2 Scale limits for NSP deployments

5.2.1 NSP deployment with LSP management

The following tables present key dimension details for an NSP deployment with LSP management. These scale numbers apply to all functions of NSP unless otherwise specified.

Table 5-1 Scale limits for an NSP deployment with LSP management

Key dimension	Lab platform	Production platform
Number of IP services managed	3000	10 000 000
Number of intent aware services	1500	6 000 000
Number of service endpoints	6000	20 000 000
Number of intent aware service endpoints	3000	12 000 000
Combined total of RSVP-TE and SR-TE LSPs (non PCE controlled)	400	40 000
Number of service tunnels	600	60 000
Number of supported NEs	500	50 000

5.2.2 NSP deployment with path control

The following table presents key dimension details for an NSP deployment with path control.

Table 5-2 Scale limits for an NSP deployment with path control

Key dimension	Lab	Medium	Large
Total Number of LSPs, of which	5000	10 000	200 000
Number of delegated RSVP-TE or SR-TE LSPs or both (PCE-Control, PCE-Compute and PCE-Report), of which	2000	5000	90 000
Utilization based LSPs, of which		2500	45 000
Flow Collector LSP Paths		1000	30 000
Latency based LSPs		2500	45 000
SR Policy Segment Lists	1000	1000	20 000
Number of IP NEs (PCEP)	1000	2000	10 000 (see Note)
Number of IP NEs (BGP-LS)	1000	2000	10 000
Number of IP links, of which	2000	7000	80 000
Telemetry enabled links		7000	50 000

Notes:

1. Customers wishing to deploy more than 6000 PCEP IP NEs should contact Nokia for details on timer configuration.

5.2.3 NSP deployment with path simulation

The following table presents key dimension details for an NSP deployment with path simulation.

Table 5-3 Scale limits for an NSP deployment with path simulation

Key dimension	Lab platform	Production platform
Total Number of LSPs	5000	144 000
Number of IP NEs	1000	10 000
Number of IP links	3000	50 000

5.2.4 Scale limits for Workflows deployment

NSP supports up to 64 asynchronous tasks. When the limit is exceeded, additional tasks are queued and will get executed as any of the executing workflows completes.

NSP supports up to 15 synchronous workflows. Synchronous workflows consume more resources and must be limited in parallel executions.

The following table presents scaling dimensions of stored executions for deployment of the NSP's Workflows function.

Table 5-4 Scale limits for Workflows deployment

Key dimension	Lab platform	Production platform
Number of stored executions (results)	500 000	1 000 000

5.2.5 Scale limits for Model Driven Mediation

The following tables present scaling dimensions for Model Driven Mediation application.

Table 5-5 Scaling dimensions – Model Driven Mediation application

Key dimension	Scale Limit
Network Infrastructure Management (nim) LLDP link objects	40 000

i **Note:** The nim module is used to identify and stitch MDM link endpoints. The above limit applies to any LLDP link with at least one MDM endpoint.

Table 5-6 MDM Authentication Session limits with NEs

Protocol	Session Limits
NETCONF	Up to three concurrent channels per NE during discovery. (One RW with notifications/keep-alive, two RO with 10-min idle timeout.) One channel during normal operation.
gNMI/gRPC	Two channels (one channel from MDM, one channel from CN-Telemetry).
CLI/SSH	One active during database backup or upgrade operations.
SFTP/SSH	One for accounting file transfers and one for NE load transfer.

5.2.6 Scale limits for IGP Topology

The following table presents scaling dimensions for IGP Topology.

Table 5-7 Scaling dimensions – IGP Topology

Key dimension	Lab	Medium	Large
Number of IP NEs	1000	2000	10 000
Number of IP Links	2000	7000	80 000

5.2.7 Scale limits for Device Configuration

The following table presents scaling dimensions for intent-based device configuration in the Device Management views.

Table 5-8 Scaling dimensions – intent-based device configuration

Key dimension	Scale Limit
Configuration templates within Device Management	500
Total number of configuration instances in Device Management	500 000

i **Note:** Customers should use caution when invoking bulk actions at the template level with many thousands of configuration instances as this may take many hours to complete. Bulk operations will be optimized in a future NSP release.

5.2.8 Scale limits for intent type artifacts

NSP provides a wide variety of artifact bundles, some of which contain intent type artifacts. It is recommended that artifact bundles are installed selectively to ensure that no more than 200 intent type artifacts are deployed on a single NSP system.

If your NSP system requires more than 200 intent type artifacts, or if you have unknowingly exceeded 200 intent type artifacts and your pod won't come back up after a restart, please contact Nokia support.

5.2.9 IP/optical coordination scaling within NSP classic and model-driven IP + optical deployment

The following table presents key dimension details for NSP's IP/optical coordination function in an NSP classic and model-driven IP + optical deployment as described in the *NSP Installation and Upgrade Guide*.

Table 5-9 Scale limits for IP/optical coordination in an NSP classic and model-driven IP + optical deployment

Key dimension	Lab platform	Production platform
IP NEs	100	8000
Optical NEs	100	3000
Ports ¹	2400	490 000
Links	250	25 000
CDLs	100	10 000
Optical services	200	20 000

Table 5-9 Scale limits for IP/optical coordination in an NSP classic and model-driven IP + optical deployment (continued)

Key dimension	Lab platform	Production platform
LLI	50	5000
IP-optical correlation	100	10 000

Notes:

1. Ports include IP physical ports, LAG ports, and Optical physical ports

5.2.10 Scale limits for SNMP adapters

The NSP SNMP adapter artifacts, which provide SNMP adaptation for standard SNMP devices, can manage up to 50,000 devices in a large-scale NSP deployment. Devices are discovered at /24 subnet (250NEs at once), however, multiple discovery rules can be initiated in sequence. For large deployments, the user should stagger the discovery rule execution to limit the impact on the system. SNMP adapters utilize the MDM framework to manage basic functions such as discovery, fault management, and port inventory.

i **Note:** The above figures represent the discovery of 7705 SAR NEs with minimal configurations, and do not account for more complex configurations such as services.

5.2.11 Scale limits for REST tokens

The supported maximum rate for granting REST tokens is one token every 2.5 seconds.

5.3 Scale limits for functions

5.3.1 Concurrent session limits

NSP supports a combined concurrent session limit of 125. This limit applies to GUI and API sessions but does not include NBI subscriptions. The following table defines any concurrent session limits of NSP functions within the global NSP wide session limit.

Table 5-10 Concurrent session limits for NSP functions

Function	Maximum number of concurrent sessions
Analytics reports	10
Concurrent NE sessions on NFMP managed nodes (see Note)	100
Concurrent NE sessions on MDM managed nodes (see Note)	100

Notes:

1. NSP can open a SSH or Telnet session with a NE in a browser window. NE sessions can be opened with NEs that are managed using NFM-P or MDM. Concurrent NE sessions refers to NE management (SSH or Telnet) sessions, not NE mediation communications with NSP and NFMP.

5.3.2 Scale limits for Kafka event notifications

Kafka event notifications supports a maximum of 200 OSS subscriptions through NBI notification. Within that global limit, Alarm Management supports a maximum of 5 OSS subscriptions.

5.3.3 Scale limits for Telemetry

Scale Limits for MDM Telemetry (multi-vendor SNMP)

MDM Telemetry data collection is limited by the maximum OSS subscriptions supported by Kafka. The maximum number of Telemetry notifications per second is 1500 per active MDM instance, where one Telemetry record (a collection of statistics counters) update equals one Telemetry notification.

Scale Limits for Cloud Native Telemetry (gNMI)

Cloud Native Telemetry data collection is limited by the maximum OSS subscriptions supported by Kafka. The maximum number of Telemetry notifications per second is 1500 per CN gNMI Collector instance, where one Telemetry record (a collection of statistics counters) update equals one Telemetry notification.

Telemetry data collection is also limited to a maximum of 2,000 NEs per CN gNMI Collector instance.

Scale Limits for Cloud Native Telemetry (accounting file processing)

NSP cluster processing of NE accounting files is performed by Accounting Processor instances. Multiple Accounting Processors can be deployed to gain higher throughput based on network size. Each Accounting Processor instance can support up to 2000 NEs, and can support a file output rate of 6200 records/s with database and kafka rates of 1500 records/s.

Typically local storage can support up to two Accounting Processor instances depending on the available system resources. To gain higher throughput, customer provided network storage should be used.

Maximum scale of 140 million accounting records every 15 minutes is achieved with 25 Accounting Processor instances subject to NSP system capacity.

Output rates limit will remain same, irrespective of which combinations of outputs are enabled (database/Kafka/file).

Scale Limits for combined MDM and Cloud Native Telemetry persistence

The maximum number of rows uploaded to the database per minute is 90 000 per combined active MDM Instance and CN gNMI Collector instance, where one row equals one Telemetry record. This limit applies for Postgres and Auxiliary database storage.

If NSP is deployed with multiple active MDM and/or CN Telemetry Instances, the maximum collective upload rate to a Postgres database is 180 000 rows per minute, regardless of the number of collectors present in the system.

When Telemetry data is stored in the Auxiliary Database, the upload rate scales horizontally with more active MDM and/or CN Telemetry Instances. network activity, database activity and latency can also affect database upload rates.

5.3.4 Event timeline limits for managed NEs and services

Some applications make use of historical data for managed NEs and services. The amount of historical data is limited according to the mediation component and database storage.

NFM-P managed NEs and services have a default event timeline of 1 week for oracle database storage and can be configured to a maximum of 1 month. For Auxiliary Database storage the event timeline can be increased to a maximum of 1 year.

For NSP, MDM managed NEs and services have an event timeline of 1 week for Postgres database storage. Auxiliary Database storage is not supported for MDM managed NEs and services.

5.3.5 File Server Service

The NSP File Server Service App can sustain a file generation rate of 22 files per second in File Server storage with an average size of 10KB. File transfer rates depend on 30 parallel API calls for downloading and deleting these files. However, users should fine-tune the number of parallel API calls based on factors such as API rate limits, network bandwidth, server capacity, and file processing time to ensure optimal performance and avoid bottlenecks.

5.3.6 Scale limits for alarms

The following table defines the alarm limits for NSP:

Table 5-11 Alarm limits

Key dimension	Maximum number of alarms
Historical alarms from non-NFM-P systems (eg. WS-NOC , MDM, NSP)	10 million
Active alarms from NFM-P, WS-NOC , and/or MDM-managed nodes	500 thousand

i **Note:** Alarm limits describe the aggregate number of alarms that can be handled by NSP but do not supersede individual datasource limits.

The following table defines the performance limits for alarms:

Table 5-12 Alarm performance limits

Key dimension	Rate
Sustained alarm rate (combined from all sources) (see Note)	200/second
Concurrent event notification subscriptions limit	5

i **Note:** Alarm rate describes the aggregate volume that can be handled by NSP but does not supersede individual datasource limits.

The following table defines the squelching limits for alarms:

Table 5-13 Alarm squelching limits

Key dimension	Maximum number of objects
Port squelching	1000 ports
Network element squelching	1000 network elements
Resource group squelching	250 000 ports and/or network elements combined

i **Note:** Resource groups are collections of network objects that NSP users are granted permission to access based on their role. Because the maximum size for a port group is currently 100k (100 000) ports, multiple resource groups are needed to achieve the 250k squelching limit.

5.3.7 Network Health Overview

Network Map

The number of NEs and links managed in the network may affect performance and topology rendering time.

Multi-layer maps support a recommended maximum of 4000 objects. Users should expect the following multi-layer map loading times with different numbers of NEs.

- For 250 NEs (125 physical links), approximately six seconds for the initial page loading and four seconds to reload.
- For 500 NEs (250 physical links), approximately nine seconds for the initial page loading and six seconds to reload.
- For 2000 NEs (1000 physical links), approximately 50 seconds for the initial page loading and 28 seconds to reload.

Link Utilization Map

The Link Utilization map view has limits on the number of supported endpoints and links that can subscribe for stats simultaneously. The following table lists the recommended maximum number of links on the current operational view for different NE types. These are not absolute maximum values but safe recommended limits based on product testing.

Table 5-14 Recommended maximum number of links

Link Type	Maximum
7750 SR physical link	500
7705 SAR / 7210 SAS physical link	200

Table 5-14 Recommended maximum number of links (continued)

Link Type	Maximum
7750 SR LAG link	160
7705 SAR / 7210 SAS LAG link	60

5.3.8 MD-OAM Scale and Performance

Scale limits and performance metrics are based on testing with an enhanced NSP deployment.

Table 5-15 Scale and performance limits - isolated per metric type

Performance Metric	Scale Limit (records/time)	Scale Limit (tests/time interval)
gNMI collection rate	2.7m records/15min	30 000 OAM tests @ 1 metric/10s
oam-pm accounting collection rate	2.4m records/15min	36 000 OAM tests @ 22 bins, 15min measured interval
combined CRUD rate ¹	48 000 CRUD actions/15min	24 000 service validation workflows/hour

Notes:

1. Combined CRUD is a term used to define an aggregation of CRUD (create/update/delete) operations in the setup/execution/retrieval/deletion of service validation tests.

Table 5-16 Scale and performance limits - combined activity

Performance Metric	Scale Limit (records/time)	Scale Limit (tests/time interval)
gNMI collection rate + oam-pm accounting collection rate	3.8m records/15min	20 000 OAM tests @ 1metric/10s + 30 000 tests @ 22bins, 15m measured interval
gNMI collection rate + oam-pm accounting collection rate + combined CRUD rate ¹	2.2m records/15min + 24 000 CRUD actions/15min	12 000 OAM tests @ 1metric/10s + 16 000 tests @ 22bins, 15m measured interval + 12 000 service validation workflows/hr

Notes:

1. Combined CRUD is a term used to define an aggregation of CRUD (create/update/delete) operations in the setup/execution/retrieval/deletion of service validation tests.

The calculation of records per test type can be determined with the following table.

Table 5-17 Records per Test Result

Test Type	Record Count	Results class info
cfm-dmm (accounting)	24 to 93 per measurement interval	3: /base/oampm-accounting/cfm-dmm-session-acc-stats; 3 * # of bins (7-30): /base/oampm-accounting/cfm-dmm-bin-acc-stats
cfm-dmm (gNMI) ¹	1 to 6 per sample window	1 * # of metrics: /base/oam-pm/eth-cfm-delay-streaming
cfm-lmm (accounting)	1 per measurement interval	1: /base/oampm-accounting/cfm-lmm-session-acc-stats
cfm-loopback	1 per test execution	1: /base/oam-result/loopback-result
cfm-slm (accounting)	1 per measurement interval	1: /base/oampm-accounting/cfm-slm-session-acc-stats
twamp-light (gNMI) ¹	1 to 6 per sample window	1 * # of metrics: /base/oam-pm/twamp-light-delay-streaming
twamp-light delay (accounting)	24 to 93 per measurement interval	3: /base/oampm-accounting/twl-session-acc-stats; 3 * # of bins (7-30): /base/oampm-accounting/twl-bin-acc-stats
twamp-light loss (accounting)	1 per measurement interval	1: /base/oampm-accounting/twl-session-loss-acc-stats

Notes:

1. Metrics for gNMI collection records include: fd-average (directions forward, backward, round-trip) and ifdv-average (directions forward, backward, round-trip).

5.3.9 Device Management

NSP supports up to 50,000 nodes in the Device Management application. Up to 8000 MDM managed nodes are supported.

The following limitations apply when using NSP's Device Management views.

While the NE inventory tree is manually traversed, 500 objects are loaded at a time, to a maximum of 2000 objects.

If any logical tree level shows more than 500 objects, that entire logical tree level is refreshed after a manual refresh is clicked, but only the first 500 objects are displayed.

5.3.10 Scale limits of Map Layout and Group Directories

Nokia recommends a maximum of 2000 NEs per region for the Operational map view.

Where Path Control and/or IP/Optical Coordination is deployed, the following scaling limits for Map Layout will apply:

- Maximum number of nodes per region is 250.
- Maximum number of links per region is 1200.

Group directories have the following scaling limits.

- There is no limit on the number of directories for each directory type (network element, port, service, analytics resource).
- Maximum number of groups per directory is 5000.
- Maximum number of objects per group is 100 000.

5.3.11 Scale limits for NSP Baseline Analytics

The NSP Baseline Analytics can support collection storage in the Postgres database or in the Auxiliary database. Baselines are supported on NFM-P and MDM managed nodes.

Table 5-18 Scale limits – NSP Baseline Analytics

Key dimension	Postgres database storage	AuxDB storage
Number of baselines	10 000	100 000
Retention time	35 days	403 days
Collection Interval	300 seconds	300 seconds
Window Duration	15 minutes	15 minutes
Season	1 week	1 week

i **Note:** Reducing the Collection Interval or Window Duration will result in a reduced number of Baselines that can be supported.

5.3.12 Scale limits for NSP Indicators

The NSP Indicators can support collection storage in the Postgres database or in the Auxiliary database. Indicators are supported on NFM-P and MDM managed nodes.

The NSP Indicators can only support up to a total of 20 Indicator rules. The recommended maximum number of resources that can feed an Indicator rule is 2500.

Table 5-19 Scale limits – NSP Indicators

Key dimension	Postgres database storage	AuxDB storage
Number of resources (number of incoming entities into NSP Indicators)	10 000	50 000
Retention time	35 days	403 days
Collection Interval (Complex Indicators)	300 seconds	300 seconds
Collection Interval (Simple Indicators)	900 seconds	900 seconds
Window Duration (Complex Indicators)	15 minutes	15 minutes

i **Note:** Reducing the Collection Interval or Window Duration will result in a reduced number of resources that can be supported.

5.3.13 Flow Collector scale for NAT collection

The Flow Collector BB NAT collection limit is 350,000 records/s per flow collector instance when the customer retrieves files with the native s/ftp application. If using the ftp offload option, the BB NAT collection limit is 250,000 records/s, per flow collector instance.

5.3.14 Scale limits for Large Scale Operations

The Large Scale Operations feature has scaling limits for framework and for device operations.

The following table summarizes the framework limits.

Table 5-20 Framework limits – Large Scale Operations

Key dimension	Maximum
Number of concurrent LSO executions	20
Number of stored operations (historical and running)	500
Number of operation types	100
Number of targets per operation	10 000
Number of phases per operation type	10

The following table summarizes the NE device operation limits.

Key dimension	Maximum
Number of nodes for NE backup	10 000

i **Note:** Numbers are based on using enhanced profile disk availability for File Service.

i **Note:** Role Based Access Control will not apply to LSO app user operations in this release.

5.3.15 Scale limits for Zero Touch Provisioning

The following limits apply to Day 0 Zero Touch Provisioning (ZTP):

Table 5-21 Scale limits – Zero Touch Provisioning

Key dimension	Maximum
NE instances created per second	5
Simultaneous downloads from file server	10
ZTP instances in various provisioning states	1000

5.3.16 Scale limits for Generic Mediator

The Generic Mediator application has the following scaling limits:

Table 5-22 Scale limits – Generic Mediator

Key dimension	Maximum
Concurrent threads	10
Request queue size	50

5.3.17 User Access Control Performance

In an NSP deployment with User Access Control enabled, and more than 20 user groups are defined, in large networks (> 2000 NEs), NSP GUI performance may be affected if the resource groups contain a very large number of equipment and/or service objects.

5.4 NSP application downtime in HA and redundant deployments

5.4.1 Overview

NSP components in a redundant configuration will experience application down time during an activity switch. NSP cluster components that support HA deployment may experience application down time during pod re-selection. Other NSP components that do not support HA deployment may incur down time for a pod restart.

Time estimates in this section are based on testing in a lab environment with a small number of nodes. Customer production networks may experience different downtime intervals based on, but not limited to, deployment type, managed network size and installation options. Estimates provided here are intended to provide guidance to network engineers and administrators.

i **Note:** The first DR switchover in an NSP cluster deployment after an install or upgrade will initialize adaptors on the redundant NSP cluster. As a result, slower recovery of NSP services is expected on the first DR switchover. Subsequent DR switchovers will not have this performance impact.

5.4.2 NSP GUI and Access Token

Users and client applications that need access to NSP GUI and APIs will experience downtime during an activity switch. When an activity switch is initiated, all active sessions will terminate. Once the new active is up and initialized, new GUI and API sessions can be opened.

The NSP authentication module does not support high availability. If the nspos-keycloak pod experiences a restart, then users and client applications will experience a period of time where logins cannot be processed but existing sessions will continue to operate.

Table 5-23 Down time – User and client applications

NSP GUI and API	
Login down time for an activity switch in a redundant NSP deployment	8 minutes
Login down time for a nspos-keycloak pod restart	2 minutes

5.4.3 Service management

Service provisioning activities will be impacted by a DR switchover from active to standby NSP cluster.

Table 5-24 Down time – Activity switch in a redundant deployment

Service management	
Down time for an activity switch from active to standby in a redundant deployment	10 minutes

5.4.4 PCE operations

PCE operations will be affected by HA switchover in an HA deployment and during an activity switch. PCE operations do support HA functionality through replica pod.

Table 5-25 Down time – PCE operations

PCE operations	
Down time for pod reselection in an HA deployment	30 seconds
Down time for an activity switch from active to standby in a redundant deployment	15 minutes

5.4.5 NRC-X operations

NRC-X operations will be affected by a pod restart in an active NSP deployment and during an activity switch. NRC-X operations do not support HA functionality through a replica pod.

Table 5-26 Down time - NRC-X operations

NRC-X operations	
Down time for pod restart on active NSP system	3 minutes
Down time for an activity switch from active to standby in a redundant deployment	20 minutes

5.4.6 Alarms

Alarm events and updates will be affected during HA switchovers and DR activity switches.

Table 5-27 Down time – Alarm events and updates

Alarms	
Down time for alarm event notifications due to HA switchover	up to 4 minutes
Down time for alarm updates due to HA switchover	up to 4 minutes
Down time for alarm event notifications due to a DR activity switch	up to 10 minutes
Down time for alarm updates due to a DR activity switch	up to 10 minutes

5.4.7 Telemetry Collection

In a N+M MDM deployment, telemetry collection will be impacted during HA and DR switchovers.

Table 5-28 Down time – MDM server pods

Active NSP cluster MDM pod restart	Telemetry Collection Down Time
Reset 1 active MDM server pod (switch to protection pod)	15 seconds
Reset all active MDM server pods	90 seconds

When the new active NSP is starting up following a DR switchover, the telemetry application must wait for all MDM servers and NFM-P main server to be active, as well as Postgres DB to be up and running before enabling previously persisted telemetry subscriptions.

Table 5-29 DR switchover times

DR switchover type	Restconf replays telemetry subscription	MDM server completes subscriptions to NEs	Total Down Time
Manual switchover	7 minutes	5 minutes	12 minutes
Automatic switchover	4 minutes	5 minutes	9 minutes

For deployments with CN Telemetry the following manual DR switchover times can be anticipated:

Test environment dimensions: 4000 NEs under management with throughput factor 2 configured in nsp-config.yml

Table 5-30 gNMI collection with output to DB and Kafka

Event following DR switchover	Elapsed time
Collector pods up	120 seconds
First NE record processed	7 minutes
Ramp to 4000 NEs and 3000 records/s	37 minutes

Table 5-31 Accounting collection with output to DB, Kafka and file

Event following DR switchover	Elapsed time
Accounting processor pods up	160 seconds
First NE file processed	7 minutes
Ramp to 4000 NEs and 3.88 million records / 15 minutes	27 minutes

5.5 NFM-P Scalability guidelines

5.5.1 Scalability limits

Table 5-32, “NFM-P scalability limits” (p. 91) represents the scalability limits for NFM-P. 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.

2.3 “Other NSP components hardware platform requirements” (p. 33) contains information about 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 station.

Contact Nokia to ensure that you have the correct platform and configuration for your network size.

Table 5-32 NFM-P 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	125
Maximum number of simultaneous active XML API HTTP applications	30
Maximum number of simultaneous active XML API JMS applications	20
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

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 5-33 Network element maximums and equivalency

Network element type	Maximum number of network elements supported	MDA equivalency
7750 SR, 7450 ESS, 7450 SR	50 000	1 MDA == 1 equivalent MDA ^{1 2}

Table 5-33 Network element maximums and equivalency (continued)

Network element type	Maximum number of network elements supported	MDA equivalency
7705 SAR	50 000	1 NE == 1 equivalent MDA
7250 IXR-6 / 7250 IXR-10 / 7250 IXR-R4 / 7250 IXR-R6 / 7250 IXR-R6d / 7250 IXR-R6dl	50 000	1 MDA == 1 equivalent MDA
7250 IXR-s / 7250 IXR-e / 7250 IXR-e2	25 000	1 NE == 2 equivalent MDAs
7210 SAS	50 000	1 NE == 1 equivalent MDA
OMNISwitch 6250, 6400, 6450, 6850 (each shelf in the stackable chassis)	50 000	1 NE == 1 equivalent MDA
OMNISwitch 6350, 6465, 6560, 6865 (each shelf in the stackable chassis)	5000	1 NE == 1 equivalent MDA
OMNISwitch 6860, 6860E, 6860N	5000	1 NE == 1 equivalent MDA
OMNISwitch 6900	800	1 NE == 1 equivalent MDA
OMNISwitch 9600, 9700, 9700E, 9800, 9800E (each NI)	1000	1 NI == 1 equivalent MDA
OMNISwitch 10K (each NI)	400	1 NI == 1 equivalent MDA
9500 MPR / Wavence SM	15 000	1 NE == 1 equivalent MDA
1830 VWM OSU	2000	³
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 1830 VWM OSU Card Slot has an MDA equivalency of 1/4 MDA per card to a maximum MDA equivalency of 30 000

5.5.2 NFM-P performance targets

Table 5-34, “NFM-P performance targets” (p. 93) 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 5-34 NFM-P 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
Functionality downtime from failover	
Alarms / SNMP Traps	2- 30 minutes (subject to network dimensions)
Statistics Collection	2- 30 minutes (subject to network dimensions)
Functionality downtime from server activity switch	
Alarms / SNMP Traps	10 - 30 minutes (subject to network dimensions)
Statistics Colleciton	10 - 30 minutes (subject to network dimensions)

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.

5.6 Scaling guidelines for NFM-P XML API clients

5.6.1 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 station 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.

5.6.2 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.

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

Table 5-35 JMS durable messaging rates

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

5.7 Scaling guidelines for statistics collection

5.7.1 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.

5.7.2 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.

5.7.3 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 station for statistics collection. The following sections should be used to determine the maximum performance and accounting statistics for different hardware setups.

5.7.4 Performance statistics

See 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 1000 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 5-36 Maximum number of performance statistics records processed by an NFM-P server

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

Table 5-37 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, 32 GB RAM	12 CPU Cores, 32 GB RAM	8 CPU Cores, 32 GB RAM	12 CPU Cores, 32 GB RAM	12 CPU Cores, 32 GB 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.

5.7.5 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 100 ms, 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).

5.7.6 Accounting statistics

See 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 station and whether the NFM-P database software is collocated with the NFM-P server software on the same station.

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 station, the NFM-P database station requires a customized configuration that can be obtained from Nokia personnel.

Table 5-38 Maximum number of accounting statistics records processed by an NFM-P server station

Number of CPU cores on NFM-P server stations	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 5-39 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, 32 GB RAM	12 CPU cores, 32 GB RAM	8 CPU cores, 32 GB RAM	12 CPU cores, 32 GB RAM	12 CPU cores, 32 GB 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.

5.7.7 Application assurance accounting statistics

See 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 station and whether the NFM-P database software is collocated with the NFM-P server software on the same station.

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 5-40 Maximum number of application assurance accounting objects configured for collection by an NFM-P server station

Number of CPU cores on NFM-P server stations	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 5-41 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, 32 GB RAM	12 CPU cores, 32 GB RAM	8 CPU cores, 32 GB RAM	12 CPU cores, 32 GB 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.

5.7.8 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.

5.7.9 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 station. The table below highlights these requirements.

Table 5-42 NFM-P database station 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 station 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

5.7.10 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 5-42, “NFM-P database station hardware requirements for a distributed configuration” \(p. 99\)](#) shows an example of simultaneous collection.

5.7.11 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 backup times.

5.7.12 Statistics record retention

The tables below shows the retention that is achievable depending upon the total number of records to retain, the statistic type, and the database used to retain the records:

Table 5-43 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 5-44 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).

5.8 Scaling guidelines for service assurance tests

5.8.1 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 5-45 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	6000	22 500	60 000

Table 5-45 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	3000	1500	15 000

Notes:

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

5.8.2 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 5-46, "OAM test weight" \(p. 103\)](#) 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 5-46 OAM test weight

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

5.8.3 Accounting file STM test configuration

Accounting file collection of STM test results requires 7750 SR and 7450 ESS 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. See [Table 5-45, “Maximum number of STM elemental test results” \(p. 101\)](#) for specific scaling limits. The logToFile method for file accounting STM results supports a maximum of two JMS clients.

5.8.4 Examples of STM test configuration

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

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.

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

5.8.5 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 (for example, 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; in general, 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 station used to perform the tests will dictate how many physical resources NFM-P can dedicate to executing elemental tests. On the minimum supported station (collocated NFM-P server and NFM-P database on a single server), the number of concurrent tests must be limited to 3 000.

5.8.6 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:

- 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.
- 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.

5.8.7 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 80 GB. 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 600 GB 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.

5.8.8 Scaling guidelines for OAM PM test results

See the NSP NFM-P Classic Management User Guide for details on OAM PM test configuration and result retrieval.

The quantity of resources which are allocated to the retrieval and processing of OAM PM test results 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 station and whether the NFM-P database software is collocated with the NFM-P server software on the same station

The following tables provide the maximum number of OAM PM test results that can be retrieved and processed by the NFM-P server or NFM-P statistics auxiliary in various configurations.

Table 5-47 Maximum number of OAM PM test results processed by an NFM-P server

Number of CPU cores on the NFM-P server	Maximum number of OAM PM test results per 15-minute interval	
	Collocated configuration	Distributed configuration
6	100 000	200 000
8 or greater	200 000	400 000

Table 5-48 Maximum number of OAM PM test results processed by an NFM-P statistics auxiliary

Number of active NFM-P statistics auxiliaries	Maximum number of OAM PM test results per 15-minute interval				
	OAM PM test result collection with NFM-P database		OAM PM test result collection with single auxiliary database	OAM PM test result collection with three+ auxiliary database cluster	logToFile only
	8 CPU cores, 32 GB RAM	12 CPU cores, 32 GB RAM	8 CPU cores, 32 GB RAM	12 CPU cores, 32 GB RAM	12 CPU cores, 32 GB 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

The table below shows the retention that is achievable depending upon the total number of test results to retain and the database used to retain the records:

Table 5-49 Maximum OAM PM test result retention

Database to retain records	Total number of OAM PM test results to be stored in the database	Maximum number of retention intervals
NFM-P database	<40M	672
	>40M	96
NSP auxiliary database	N/A	35,040

5.9 cflowd statistics collection

5.9.1 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 5-50 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, that is 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 (for example, <5% total flows TCP/VOL/RTP)

Table 5-51 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

5.9.2 AA Comprehensive Raw Aggregation Limits

Flow Collector performance is impacted by large cflowd message sizes. The larger the cflowd message, the lower the Flow Collector output performance.

Table 5-52 Flow Collector scale limits by cflowd message size

cflowd message size	Maximum flows per second
less than 548 bytes	100,000
less than 864 bytes	70,000
less than 1000 bytes	50,000

Notes:

1. The Flow Collector can handle increased flow rates with fine-tuning, MTU changes and additional resources. Please contact Nokia Support for details.

5.10 CPAM and vCPAA

5.10.1 Scaling guidelines for CPAM and vCPAA

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

i **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 5-53 7701 CPAA Hardware revision 2 and vCPAA scalability

Item description	Scaling
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)	2000 configured 2000 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 5-54 CPAM scalability

Item description	Scale Limit
Maximum number of routers per admin domain	4000
Maximum Active 7701 CPAAs	40
Maximum Number of IGP Routers	12 000
Maximum Number of OSPF Domains/Areas	150

Table 5-54 CPAM scalability (continued)

Item description	Scale Limit
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
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 Security

6.1 Introduction

6.1.1 Overview

This chapter provides general information about platform security for a deployment of NSP and NFM-P. Recommendations in this section apply to NSP and its optional components except where indicated.

The NSP implements a number of safeguards to ensure the protection of private data. Additional information can be found in the NSP Data Privacy section of the *NSP System Architecture Guide*.

6.2 Securing the NSP

6.2.1 Overview

Nokia recommends performing the following steps to achieve station security for the NSP cluster:

- Install the latest recommended patch cluster for RHEL. For customers using the Nokia provided NSP RHEL OS image, only the NSP RHEL OS update can be used for applying OS patches. For customer sourced and manually deployed RHEL OS instances, the patches must be obtained from Red Hat.
- NSP has no ingress or egress requirements to access the public internet and should be isolated with properly configured firewalls.
- Implement traffic management policies to control access to ports on NSP systems, as detailed in this section
- Use SSL certificates with strong hashing algorithms.
- Enforce minimum password requirements and password renewal policies on user accounts that access the NSP applications.
- Configure a warning message in the NSP GUI login page Security Statement.
- OAUTH2 authentication module provides login protection mechanisms to prevent denial of service attacks, lockout users for consecutive failed logins and configure maximum sessions for GUI and OSS users. See *NSP Installation and Upgrade Guide* for details.
- When using custom TLS certificates for NSP deployment, ensure that the server private key file is protected when not in use by nsp configurator.
- Enable SELinux in permissive/enforcing mode for components that support it
- Optional: Revoke world permission on compiler executables (see *NSP Installation and Upgrade Guide*).

See the *NSP Security Hardening Guide* for NSP RHEL OS recommendations and best practices. These best practices are already implemented on NSP RHEL OS images.

6.2.2 TLS communications

Communications of the NSP cluster is secured using TLS. The NSP supports TLS version TLSv1.2.

The NSP supports the use of custom TLS certificates for client communications with NSP applications. Internal communications between NSP components is secured with internal TLS certificates signed by a local PKI server. The NSP cluster software package provides a PKI server that can be used to simplify the TLS certificate distribution to NSP components.

An NSP cluster will check the expiry date of TLS certificates every 24h and raise an alarm in Network Map and Health dashboard if the certificate is expired or nearing expiry. See the *NSP System Administrator Guide* for further information.

See the *NSP Installation and Upgrade Guide* for instructions on the configuration of custom TLS certificates and the provided PKI server application.

6.3 Operating system security for NSP cluster components

6.3.1 RHEL patches

For customer sourced and manually deployed RHEL OS instances, Nokia supports customers applying RHEL patches provided by Red Hat which include security fixes as well as functional fixes. If a patch is found to be incompatible with the NSP, the patch may need to be removed until a solution to the incompatibility is provided by Red Hat or Nokia. See the NSP Release Notice for up-to-date information about the recommended RHEL maintenance update and patch levels.

For customers using the Nokia provided NSP RHEL OS images, only the NSP RHEL OS update provided by Nokia can be applied.

NSP is supported on RHEL installed with the list of required RHEL Packages documented in the NSP Installation and Upgrade Guide. SELinux is supported in both permissive and enforcing mode for most NSP components. Auxiliary databases support SELinux in permissive mode only

6.3.2 Platform hardening

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


6.4 Communication between the NSP and external systems

6.4.1 Overview

The following diagrams illustrate the various components of the NSP and its internal communications, as well as communications with external systems.

The following figure shows a standalone NSP deployment and its communications with external systems.

Figure 6-1 Standalone NSP deployment



38443

Table 6-1 NSP systems or environment connections

Connection	Usage
1	Web Client/REST API client connections. REST over HTTPS secured with TLS
2	SSO authentication (secure), zookeeper registration (secure), neo4j database (non-secure), kafka (secure), NFM-P API (secure), Data connection – CPROTO protocol secured with TLS, PKI server communications
3	NE mediation using SNMP and FTP/SCP
4	NE mediation using gRPC/gNMI, SNMP, NETCONF, SSH, ICMP

Table 6-1 NSP systems or environment connections (continued)

Connection	Usage
5	Data connection – CPROTO (non-secure)
6	BGP-LS (supports GTSM), PCEP (secured by TLS), OpenFlow communications (secured by TLS) * Note
7	SSO authentication (secure), zookeeper registration (secure), REST over HTTPS secured with TLS, proprietary HTTP with WS-NOC, PKI server communications
8	NE mediation with SNMP and TL-1
9	syslog notifications secured with TLS
10	Mediator communications with external controller, REST/RESTCONF secured with TLS
11	LDAP, RADIUS, TACACS communications with remote authentication servers (ICMP connectivity required)
12	SMTP, SMTPS, STARTTLS communication with email server

i **Note:** VSR-NRC supports secure PCEP and OpenFlow communications in specific releases. See the SR OS documentation for details.

6.5 NSP Port Communications

6.5.1 Overview

This section will document network communications between components in an NSP deployment. These tables can be used by customers to design traffic management policies based on their NSP deployment.

A complete listing of network communications for additional NSP components can be found in section 6.10 of this guide.

The following port changes are reported for NSP in Release 25.11

- Remove port 111 from NSP deployer and cluster nodes.

Table notes:

- Each table identifies network communications based on the destination component.
- Each communication link defines traffic from the originating component/port to the destination component/port. When traffic policies are applied in both directions of communication, the return path must also be permitted.
- In a multi-node NSP cluster deployment, communications originating from NSP to a destination must allow traffic from each node of that NSP cluster to the destination component. Traffic destined to a multi-node NSP cluster will require communications to the virtual IP address of the NSP cluster.

- For NSP deployments with multiple network interfaces, the communications matrix will define on which network interface the communications will be received.
- Where multiple components may be communicating with a destination component and port, each source component with source port range is listed.
- A system administrator will require SSH access to components in the NSP deployment for installation and maintenance purposes. For this purpose, tables will list a source component of System administration server.

i **Note:** The ephemeral port range of different server types may vary. Many Linux kernels use the port range 32768 - 61000. To determine the ephemeral port range of a server, execute `cat /proc/sys/net/ipv4/ip_local_port_range`

i **Note:** Some NSP operations require idle TCP ports to remain open for long periods of time. Therefore, customers that implement a network traffic policy that closes idle TCP connections should adjust operating system TCP keep-alives to ensure that NSP communications is not impacted (ie. set OS TCP keep-alives to be less than idle TCP timeout within network traffic policies).

i **Note:** The use of firewalld is not supported on NSP cluster virtual machines. Nokia recommends using Calico policies to control traffic to an NSP cluster deployment. (Kubernetes networking relies on calico rules added to iptables. Using firewalld changes the order of those calico rules and can disrupt traffic flow in the NSP cluster.) Nokia does not provide installation, configuration, or troubleshooting support for Calico policies.

i **Note:** The 'node exporter' referred to in this section is used by Prometheus to pull performance metrics from a VM.

Table 6-2 NSP Kubernetes virtual machine communications

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
System administration server	any	22	TCP	Dynamic Encryption	any	Administrator SSH access, software installation
remote DR NSP cluster	>32768					
NFM-P main, NFM-P database, NFM-P Auxiliary	>15000	80	TCP	None	any	PKI server
Auxiliary database	>15000					
WS-NOC	>49192					

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
Network element	any	162	UDP	None	mediation	SNMP traps
Network element	n/a	n/a	ICMP	None	mediation	ICMP traffic between NSP and NEs.
browser/OSS clients	any	443	TCP	Dynamic Encryption provided by TLS	client	HTTPS communications for NSP applications, REST API, session management
Simulation Tool	>32768	443	TCP	Dynamic Encryption provided by TLS	internal	authentication, authorization, REST API
redundant NSP	>32768	443	TCP	Dynamic Encryption provided by TLS	internal	redundancy communications (DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
NFM-P main, NFM-P Auxiliary	>15000	443	TCP	Dynamic Encryption provided by TLS	internal	authentication, authorization, REST API
WS-NOC	>49192	443	TCP	Dynamic Encryption provided by TLS	client	authentication, authorization, REST API
Identity Provider(s)	any	443	TCP	Dynamic Encryption provided by TLS	client	User login services for Identity Provider applications
Network element	any	2055	UDP	None	mediation	Flow Collector: netflowV5
NFM-P main, NFM-P Auxiliary	>15000	2281	TCP	Dynamic Encryption provided by TLS	internal	Secure Zookeeper communications
WS-NOC	>49192					

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	4152	TCP	Dynamic Encryption provided by TLS	internal	ASM module (DR only)
Network element	any	4739	UDP	None	mediation	Flow Collector: ipfixV10IANA
Network element	any	4740	UDP	None	mediation	Flow Collector: ipfixV10NokiaAA
Network element	any	4741	UDP	None	mediation	Flow Collector: ipfixV10NokiaSYS
Network element	any	4742	UDP	None	mediation	Flow Collector: ipfixV10NokiaBBNAT
remote DR NSP cluster	>32768	5000, 5001	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (DR only)
remote DR NSP cluster	>32768	5002	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (HA/DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	5100, 5101	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (DR only)
remote DR NSP cluster	>32768	5102	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (HA/DR only)
remote DR NSP cluster	>32768	5200, 5201	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (DR only)
remote DR NSP cluster	>32768	5202	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (HA/DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	6000, 6001	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (DR only)
remote DR NSP cluster	>32768	6002	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (HA/DR only)
remote DR NSP cluster	>32768	6100, 6101	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (DR only)
remote DR NSP cluster	>32768	6102	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (HA/DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	6200, 6201	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (DR only)
remote DR NSP cluster	>32768	6202	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (HA/DR only)
remote DR NSP cluster	>32768	6379	TCP	Dynamic Encryption provided by TLS	internal	nspos-redis/Devicestore data replication
remote DR NSP cluster	>32768	6380	TCP	Dynamic Encryption provided by TLS	internal	devicestore-fileserver/ Initiate data replication

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	6432	TCP	Dynamic Encryption provided by TLS	internal	Postgres database
NFM-P main	>15000					
WS-NOC	>49192					
remote DR NSP cluster	>32768	7000, 7001	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (DR only)
remote DR NSP cluster	>32768	7002	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (HA/DR only)
remote DR NSP cluster	>32768	7100, 7101	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	7102	TCP	Dynamic Encryption provided by TLS	internal	nsp-tomcat neo4j (HA/DR only)
remote DR NSP cluster	>32768	7200, 7201	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (DR only)
remote DR NSP cluster	>32768	7202	TCP	Dynamic Encryption provided by TLS	internal	nrcx-tomcat neo4j (HA/DR only)
remote DR NSP cluster	>32768	7687	TCP	Dynamic Encryption provided by TLS	internal	nspos-neo4j (DR only)

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
remote DR NSP cluster	>32768	7690	TCP	Dynamic Encryption provided by TLS	internal	nrcx-neo4j (DR only)
remote DR NSP cluster	>32768	8001	TCP	Dynamic Encryption provided by TLS	internal	Role Manager (DR only)
remote DR NSP cluster	>32768	8482	TCP	Dynamic Encryption provided by TLS	internal	System Admin App (DR only)
OSS clients	any	8545	TCP	Dynamic Encryption provided by TLS	client	MDM applications

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
NSP deployer node	any	8548	TCP	Dynamic Encryption provided by TLS	internal	adaptor installation
OSS clients	any	8548	TCP	Dynamic Encryption provided by TLS	client	mdmTomcat
OSS clients	any	8565	TCP	Dynamic Encryption	client	file service SFTP
remote DR NSP cluster	>32768	8566	TCP	Dynamic Encryption provided by TLS	internal	File synchronization with redundant NSP

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
NE	any	8567	TCP	Dynamic Encryption provided by TLS	mediation	File transfer with Nokia NEs.
NFM-P main	>15000	8575	TCP	Dynamic Encryption provided by TLS	internal	system token for components external to NSP
remote DR NSP cluster	>32768	8663	TCP	Dynamic Encryption provided by TLS	internal	CAM data synchronization (DR only)
OSS clients	any	9192	TCP	Dynamic Encryption provided by TLS	client	Kafka

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
NFM-P main, NFM-P Auxiliary	>15000	9192	TCP	Dynamic Encryption provided by TLS	client/internal	Kafka Applies to NSP deployments where client/internal communications are on same network interface.
WS-NOC	>49192					
OSS clients	any	9193, 9194	TCP	Dynamic Encryption provided by TLS	client	Kafka - enhanced NSP deployments only
NFM-P main, NFM-P Auxiliary	>15000	9193, 9194	TCP	Dynamic Encryption provided by TLS	client/internal	Kafka - enhanced NSP deployments only Applies to NSP deployments where client/internal communications are on same network interface.
WS-NOC	>49192					
NFM-P main, database	>15000	9200	TCP	Dynamic Encryption provided by TLS	internal	Opensearch log collection

Table 6-2 NSP Kubernetes virtual machine communications (continued)

Source component(s)	Source Port	NSP Destination Port	Transport Protocol	Encryption	Network Interface	Description/Purpose
NFM-P main, NFM-P Auxiliary	>15000	9292	TCP	Dynamic Encryption provided by TLS	internal	Kafka Applies to NSP deployments where client/internal communications are on separate network interfaces.
WS-NOC	>49192					
NFM-P main, NFM-P Auxiliary	>15000	9293, 9294	TCP	Dynamic Encryption provided by TLS	internal	Kafka - enhanced NSP only Applies to NSP deployments where client/internal communications are on separate network interfaces.
WS-NOC	>49192					

Table 6-3 Network Element Communications

Source component	Source port	NE Destination Port	Transport Protocol	Description
NSP kubernetes VM	n/a	n/a	ICMP	Network connectivity test, ICMP type 8 (request) and type 0 (reply)
System administration server	any	22	TCP	Administrator SSH access, SFTP
NSP kubernetes VM	>32768			
NSP kubernetes VM	>32768	161	UDP	SNMP mediation
NSP kubernetes VM	>32768	830	TCP	NETCONF mediation
NSP kubernetes VM	>32768	57400	TCP	gRPC

Table 6-3 Network Element Communications (continued)

Source component	Source port	NE Destination Port	Transport Protocol	Description
NSP kubernetes VM	>32768	21	TCP	telnet, FTP access - use only where required

Table 6-4 VSR-NRC Communications

Source component	Source port	VSR-NRC Destination Port	Transport Protocol	Description
NSP kubernetes VM	>32768	4199	TCP	Network topology information, service management
Primary VSR-NRC	4199	4199	TCP	One possible inter-VSR-NRC link
Standby VSR-NRC	4199	4199	TCP	Another possible inter-VSR-NRC link

Refer to the *Security Best Practices and Hardening Guide* for detailed information on secure communications with VSR-NRC.

Refer to section 6.10 of this guide for a complete list of firewall rules for NFM-P and associated components.

Table 6-5 NFM-P Main Server Communications

Source component	Source port	NFM-P Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	7879	TCP	internal	CPROTO port
NSP kubernetes VM	>32768	8087	TCP	client	web applications communications
NSP kubernetes VM	>32768	8089	TCP	client	web applications communications
NSP kubernetes VM	>32768	8443	TCP	client	XML API

Table 6-5 NFM-P Main Server Communications (continued)

Source component	Source port	NFM-P Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	9100	TCP	internal	node exporter

NSP communicates with NFM-P Database Server and NFM-P Auxiliary Server for collecting metrics.

Table 6-6 NFM-P Database Server Communications

Source component	Source port	NFM-P Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	9100	TCP	internal	node-exporter

Table 6-7 NFM-P Auxiliary Server Communications

Source component	Source port	NFM-P Aux Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	9100	TCP	internal	node exporter

Table 6-8 Auxiliary Database Server Communications

Source Component	Source Port	AuxDB Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	5433	TCP	internal	
NSP kubernetes VM	>32768	7299	TCP	internal	
NSP kubernetes VM	>32768	9100	TCP	internal	node exporter

Refer to WS-NOC documentation for a complete list of WS-NOC application communications.

Table 6-9 WS-NOC Communications

Source Component	Source Port	WS-NOC Destination Port	Transport Protocol	Network Interface	Description
NSP kubernetes VM	>32768	443	TCP	client	
NSP kubernetes VM	>32768	8443	TCP	client	GUI
NSP kubernetes VM	>32768	8543	TCP	client	WS-RC REST API

Table 6-10 Syslog Server Communications

Source Component	Source Port	Destination Component	Destination Port	Transport Protocol	Description
NSP kubernetes VM	>32768	Syslog server	514	TCP	syslog notifications

Table 6-11 Mail Server Communications

Source Component	Source Port	Destination Component	Destination Port	Transport Protocol	Description
NSP kubernetes VM	>32768	Mail Server	25	TCP	SMTP mail server (unsecure)
NSP kubernetes VM	>32768	Mail Server	465	TCP	SMTPTS mail server (secure)
NSP kubernetes VM	>32768	Mail Server	587	TCP	STARTTLS mail server (secure)

Table 6-12 Remote Authentication Servers and Identity Provider Communications

Source Component	Source Port	Destination Component	Destination Port	Transport Protocol	Description
NSP kubernetes VM	n/a	RADIUS server	n/a	ICMP	Test connectivity of NSP and RADIUS server
NSP kubernetes VM	n/a	TACACS server	n/a	ICMP	Test connectivity of NSP and TACACS server
NSP kubernetes VM	n/a	Identity Provider	n/a	ICMP	Test connectivity of NSP and Identity Providers
NSP kubernetes VM	>32768	LDAP server	389	TCP	LDAP (unsecure)
NSP kubernetes VM	>32768	LDAP server	636	TCP	LDAP (secure)
NSP kubernetes VM	>32768	RADIUS server	1812	UDP	RADIUS
NSP kubernetes VM	>32768	TACACS server	49	TCP	TACACS
NSP kubernetes VM	>32768	Identity Provider	443	TCP	NSP communications with Identity Provider. Destination port may vary depending on Identity Provider configuration.

i **Note:** Remote authentication servers and Identity Provider applications may have security applications which block certain traffic types. For NSP to test connectivity to these remote server applications, ICMP traffic will need to be permitted as indicated

Table 6-13 Splunk Server Communications

Source Component	Source Port	Destination Component	Destination Port	Transport Protocol	Description
NSP kubernetes VM	>32768	Splunk Server	8088 (see Note)	TCP	NSP application logs to Splunk

i **Note:** Destination port determined by Splunk server configuration.

6.6 NSP Kubernetes Platform Communications

6.6.1 Overview

The tables provided in this section identify the listening ports on a deployer node and on worker nodes for an NSP cluster deployment. (Note: worker nodes of an NSP cluster includes MDM only and PCE only nodes.) These ports must be accessible between the deployer and worker nodes within an NSP deployment. The SSH ports on all servers must be accessible by a system administrator for installation and maintenance functions.

i **Note:** The 'node exporter' referred to in this section is used by Prometheus to pull performance metrics from a VM.

Table 6-14 Ports used by deployer node

Default port(s)	Type	Application
22	TCP	SSH
443	TCP	HTTPS
6443	TCP	kubernetes API server
8443	TCP	helm repo, container registry
9100	TCP	node exporter
10250	TCP	kubelet metrics
30000-32767	TCP	kube proxy

Table 6-15 Ports used by worker nodes

Default Port(s)	Type	Node(s)	Application
22	TCP	all	sshd
53	TCP	all	node-cache

Table 6-15 Ports used by worker nodes (continued)

Default Port(s)	Type	Node(s)	Application
179	TCP	all	bird
443	TCP	any two nodes in enhanced profile, any one node for all other profiles	nsp-piraeus-operator
2049	TCP	one of the nodes that has linstor-controller	NFS server service (only when NSP provides NFS storage)
2375	TCP	all	containerd
2379	TCP	all	etcd
2380	TCP	all	etcd
3370	TCP	last 3 nodes of NSP cluster	linstor-controller
4789	UDP	all	VXLAN encapsulation
6443	TCP	all	kubernetes API server
7472	TCP	all	metalLB (metrics)
7946	TCP & UDP	all	metalLB (load balancer communications)
8081	TCP	all	nginx
9100	TCP	all	node exporter
9253	TCP	all	node cache
9254	TCP	all	node cache
9353	TCP	all	node-cache
10250	TCP	all	kubelet metrics
10251	TCP	all	kube-scheduler
10256	TCP	all	kube-proxy
10257	TCP	all	kube controller
10259	TCP	all	kube scheduler
30000-32767	TCP	all	kube-proxy

6.7 Securing the NFM-P

6.7.1 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 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 Installation and Upgrade Guide*.

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

- Install a clean operating system environment with the minimum required packages documented in the *NSP Installation and Upgrade Guide*.
- Install the latest recommended patch cluster for RHEL. For customers using the Nokia provided NSP RHEL OS image, only the NSP RHEL OS update can be used for applying OS patches. For customer sourced and manually deployed RHEL OS instances, the patches must be obtained from Red Hat.
- Harden the RHEL operating system installation based upon recommendations and best practices in the *NSP Security Hardening Guide*. Reference the *NSP System Architecture Guide* for the Recommendations and Compliance statements. The *NSP Security Hardening Guide* best practices are already implemented on the NSP RHEL OS 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 [6.7.5 “Deploying NFM-P with firewalls” \(p. 137\)](#). 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 [6.10 “Additional NSP component firewall and NAT rules” \(p. 146\)](#) for more details
- Installation of NFM-P with a secure configuration described in [6.7.3 “Installing the NFM-P components” \(p. 136\)](#)
- Network element connection configuration as described in [6.7.4 “NFM-P network element communication” \(p. 136\)](#)
- 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.

6.7.2 Operating system installation for NFM-P stations

For customer sourced and manually deployed RHEL OS instances, Nokia supports customers applying RHEL patches provided by Red Hat which include security fixes as well as functional fixes. Nokia also supports customers applying Windows patches provided by Microsoft. If a patch is found to be incompatible with the NSP, the patch may need to be removed until a solution to the incompatibility is provided by Red Hat or Nokia. See the NSP Release Notice for up-to-date information about the recommended RHEL maintenance update and patch levels.

For customers using the Nokia provided NSP RHEL OS images, only the NSP RHEL OS update provided by Nokia can be applied.

NFM-P is supported on RHEL installed with the list of required RHEL Packages documented in the *NSP Installation and Upgrade Guide*. SELinux is supported in both permissive and enforcing mode for most NSP components. Auxiliary databases 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.

6.7.3 Installing the NFM-P components

Nokia recommends performing 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 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 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

6.7.4 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

6.7.5 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 6.10 “Additional NSP component firewall and NAT rules” (p. 146) 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 4.13 “Network element specific requirements” (p. 71) there is a requirement for the NFM-P GUI client to communicate directly with the network element using specialized configuration tools.

Figure 6-2 Firewalls and NFM-P standalone deployments





Figure 6-3 Firewalls and NFM-P redundant deployments



22667

6.8 Additional NSP component port information

6.8.1 Default ports

The following table describes the listening ports on the various additional NSP components.

Table 6-16 Port information

Default port	Type	Encryption	Description
NFM-P server and NFM-P statistics auxiliary			

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
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 6.9 "FTP" (p. 146))
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.
80	TCP	None See port 443 for secure communications.	HTTP This port redirects 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 NSP GUI.
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 station to adjust the NFM-P auxiliary behavior. (Example: adjusting log levels, shutting down the auxiliary server, etc)
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 (for example: database changes or alarm notifications, etc)

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
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.
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
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
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.
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.
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.
8097	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	REST port used for internal communication for DR features (DR alignment, dashboard)

Table 6-16 Port information (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 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.
8483	TCP	None	JBoss RMI port for WebServices 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
9100	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 node-exporter metrics.
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
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.

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
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.
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 stations 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.
NSP 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

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
4803	TCP	None	Spread Client connections Inter-node communication only.
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	Dynamic Encryption (if secure=true) Encryption provided by TLS. Strong ciphers are supported using various AES ciphers provided by TLS.	JDBC Client communication port (NFM-P server, statistics auxiliary, NSP Cluster)
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	TCP	Dynamic Encryption (if secure=true) Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS.	RMI NSP auxiliary database proxy port.
7300–7309	TCP	None	RMI NSP auxiliary database proxy ports. Not used if secure=true.
9100	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 node-exporter metrics.
50000	TCP	None	Rsync Inter-node and inter-cluster communication
32768-60999	TCP	None	Vertica - Zygot Inter-node communication only

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
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 6.9 "FTP" (p. 146) 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.
80	TCP	None	HTTP This port is required for the NFM-P client to communicate with the network element Web GUIs. See 4.13 "Network element specific requirements" (p. 71) 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.
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, 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			

Table 6-16 Port information (continued)

Default port	Type	Encryption	Description
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.
9003	TCP	Dynamic Encryption Encryption provided by TLS. Strong ciphers are supported using various CBC and AES ciphers provided by TLS	Database file transfer Port This port is used by the NFM-P database stations in a redundant station configuration. This port allows database transfers between the primary and standby databases. For example: when the standby database is reinstated, or when the standby database is installed for the first time.
9100	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 node-exporter metrics.
NFM-P client and client delegate server			
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.

6.9 FTP

6.9.1 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 station 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.

6.10 Additional NSP component firewall and NAT rules

6.10.1 Overview

Firewall rules are applied to the incoming network interface traffic. As a rule, firewall rules are not applied to the outgoing network interface traffic.

For 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 component. All others should disable the RHEL firewall.

The installation of each component will include the filter rules lists to be applied for successful communication between different NSP components, XML API clients, and network elements. The table below defines the location

Table 6-17 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 auxiliary	IPv4/IPv6	/opt/nsp/nfmp/auxserver/nms/sample/firewall/

Table 6-17 Sample firewalld filter rules lists file locations (continued)

Component	Protocol	File location
NSP 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 server	IPv4/IPv6	<base client install dir>/nms/sample/firewall/

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

See 7.4 “NFM-P Network Address Translation” (p. 163) for supported NAT configurations.

6.10.2 NFM-P server firewall and NAT rules

When there is a firewall at the NFM-P server, the following initiating connection details must be considered. Network Interface number is in reference to Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 160)

Table 6-18 Firewall rules for traffic connecting to the NFM-P server

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P server (public address)	Any	21	TCP	3	Connection to NFM-P server private address (if NAT in use)
XML API client	Any	21	TCP	3	If FTP is required
NSP cluster	Any	21	TCP	1	If FTP is used
NFM-P server	Any	22	TCP	1	From the redundant NFM-P server
XML API client	Any	22	TCP	3	If SCP / SFTP is required
NSP cluster	Any	22	TCP	1	If SCP / SFTP is used
9500 MPR / Wavence	Any	22	TCP	2 / 4	NE backups
NFM-P GUI client	Any	443	TCP	3	HTTPS
Managed Network	Any	162	UDP	2 / 4	SNMP trap initiated from the NE
Managed Network	Any	--	ICMP	2 / 4	Ping policy
1830 SMS HSM Server	5552	758	TCP	2 / 4	nlogin

Table 6-18 Firewall rules for traffic connecting to the NFM-P server (continued)

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P server (public address)	>1023	>1023	TCP	3	Connection to NFM-P server private address (if NAT in use)
NFM-P GUI client / XML API client	Any	1097	TCP	3	JMS
NFM-P auxiliary server	Any	1097	TCP	1	JMS
NFM-P server	Any	1099	TCP	1	From the redundant NFM-P server
NFM-P GUI client / XML API client	Any	1099	TCP	3	JNDI
NFM-P auxiliary server	Any	1099	TCP	1	JNDI
NSP cluster	Any	1099	TCP	1	JNDI
NFM-P GUI client / XML API client	Any	4447	TCP	3	JMS
NFM-P auxiliary server	Any	4447	TCP	1	JMS
NFM-P server	Any	7879	TCP	1	From the redundant NFM-P server
NSP cluster	Any	7879	TCP	1	CPROTO
XML API client	Any	8080	TCP	3	HTTP
NSP cluster	Any	8080	TCP	1	HTTP
NFM-P GUI client	Any	8085	TCP	3	HTTP
NFM-P server	Any	8087	TCP	1	From the redundant NFM-P server
NFM-P GUI client	Any	8087	TCP	3	HTTP(S)
NFM-P GUI client	Any	8089	TCP	3	HTTP(S)
NSP	Any	8097	TCP	1	From NSP
XML API client	Any	8443	TCP	3	HTTPS
NSP cluster	Any	8443	TCP	1	HTTPS

Table 6-18 Firewall rules for traffic connecting to the NFM-P server (continued)

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P GUI client	Any	8444	TCP	3	HTTPS
NFM-P server	Any	9010	TCP	1	From the redundant NFM-P server
NSP Cluster	Any	9100	TCP	1	Node-exporter
Web client	Any	9443	TCP	3	Swagger interface for HSM
NSP Cluster	Any	9990	TCP	1	Main Server node-exporter metrics
NSP Cluster	Any	10090	TCP	1	JMS Server node-exporter metrics
NFM-P server	Any	10290	TCP	1	From the redundant NFM-P server
NFM-P server	Any	11800	TCP	1	From the redundant NFM-P server
NFM-P server	Any	12010	TCP	1	From the redundant NFM-P server
NFM-P server	Any	12300-12307	TCP	1	From the redundant NFM-P server
NFM-P auxiliary server	Any	12300-12307	TCP	1	--
NFM-P auxiliary server	Any	12800	TCP	1	--

i **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).

6.10.3 NFM-P database firewall

When there is a firewall at the NFM-P database, the following initiating connection details must be considered. Network Interface number is in reference to [Figure 7-3, "Distributed NFM-P server/database deployment with multiple network interfaces" \(p. 160\)](#)/

Table 6-19 Firewall rules for traffic connecting to the NFM-P database

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P database	Any	22	TCP	1	
NFM-P server	Any	1523	TCP	1	
NFM-P database	Any	1523	TCP	1	From the redundant NFM-P database
NFM-P auxiliary server	Any	1523	TCP	1	

Table 6-19 Firewall rules for traffic connecting to the NFM-P database (continued)

Source	Source port	Destination port	Protocol	Network Interface	Notes
NSP server	Any	1523	TCP	1	
NFM-P server	Any	9002	TCP	1	
NFM-P database	9002	9002	TCP	1	From the redundant NFM-P database
NFM-P auxiliary server	Any	9002	TCP	1	
NFM-P server	Any	9003	TCP	1	
NFM-P database	9003	9003	TCP	1	From redundant NFM-P database
NFM-P auxiliary server	Any	9003	TCP	1	

6.10.4 NFM-P auxiliary server firewall and NAT rules

When there is a firewall at the NFM-P auxiliary server, the following initiating connection details must be considered. Network Interface number is in reference to [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160).

Table 6-20 Firewall rules for traffic connecting to the NFM-P auxiliary server

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P auxiliary server public address	Any	21	TCP	3	Connect to NFM-P auxiliary server private address (for NAT) If FTP is required
XML-API Client	Any	21	TCP	3	If FTP is required
XML-API Client	Any	22	TCP	3	If SFTP is required
NFM-P auxiliary server	Any	22	TCP	1	From redundant NFM-P auxiliary server
NFM-P auxiliary server public address	>1023	>1023	TCP	3	Connect to NFM-P auxiliary server private address (for NAT) If FTP is required
NFM-P auxiliary server	Any	1095	TCP	1	From redundant NFM-P auxiliary server

Table 6-20 Firewall rules for traffic connecting to the NFM-P auxiliary server (continued)

Source	Source port	Destination port	Protocol	Network Interface	Notes
NSP Cluster	Any	9100	TCP	1	Node-exporter
NSP Cluster	Any	10190	TCP	1	Auxiliary Server node-exporter metrics
NFM-P server	Any	12300 - 12307	TCP	1	--
NFM-P auxiliary server	Any	12300 - 12307	TCP	1	From redundant NFM-P auxiliary server
NFM-P server	Any	12800	TCP	1	--
NFM-P auxiliary server	Any	12800	TCP	1	From redundant NFM-P auxiliary server

i **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).

6.10.5 NSP auxiliary database firewall rules

When there is a firewall at the NSP auxiliary database, the following initiating connection details must be considered. Network Interface number is in reference to [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160).

Since the inter-node communication should traverse a private LAN, it is not recommended to implement a firewall on this interface.

Table 6-21 Firewall rules for traffic connecting to the NSP auxiliary database

Source	Source port	Destination port	Protocol	Network Interface	Notes
NSP auxiliary database	Any	22	TCP	1	SFTP between clusters
NFM-P server	Any	5433	TCP	1	JDBC
NFM-P statistics auxiliary	Any	5433	TCP	1	JDBC
NSP cluster	Any	5433	TCP	1	JDBC
NFM-P server	Any	7299	TCP	1	RMI secure = true
NFM-P server	Any	7299 - 7309	TCP	1	RMI secure - false
NSP auxiliary database	Any	50000	TCP	1	Rsync between clusters

6.10.6 NFM-P client and client delegate firewall rules

When there is a firewall at the client or client delegate, the following initiating connection details must be considered. Network Interface number is in reference to [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160).

Table 6-22 Firewall rules for traffic connecting to the NFM-P client and client delegate

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P server	--	--	ICMP	3	Ping Client delegate only
Managed Network	Any	20	TCP	2 / 4	Active FTP 9500 MPR / Wavence (NEtO)
Managed Network	Any	21	TCP	2 / 4	9500 MPR / Wavence (NEtO)
Managed Network	Any	22	TCP	2 / 4	9500 MPR / Wavence (NEtO)
Managed Network	Any	162	UDP	2 / 4	9500 MPR / Wavence (NEtO)
Managed Network	>1023	>1023	TCP	2 / 4	Passive FTP 9500 MPR / Wavence (NEtO)
Managed Network	5010	5010	UDP	2 / 4	9500 MPR / Wavence (NEtO)

6.10.7 Managed network firewall rules

When there is a firewall at the managed network, the following initiating connection details must be considered. Network Interface number is in reference to [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160).

Table 6-23 Firewall rules for traffic connecting to the managed network

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P server	Any	21	TCP	2 / 4	FTP
NFM-P auxiliary server	Any	21	TCP	2 / 4	FTP NFM-P statistics auxiliary
NFM-P client (NEtO)	Any	21	TCP	2 / 4	FTP 9500 MPR / Wavence Management
NFM-P server	Any	22	TCP	2 / 4	SSH
NFM-P auxiliary server	Any	22	TCP	2 / 4	SSH NFM-P statistics auxiliary

Table 6-23 Firewall rules for traffic connecting to the managed network (continued)

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P client (NEtO)	Any	22	TCP	2 / 4	SFTP 9500 MPR / Wavence Management (MSS-8/4/1)
NFM-P server	Any	23	TCP	2 / 4	Telnet
NFM-P auxiliary server	Any	23	TCP	2 / 4	Telnet NFM-P statistics auxiliary
NFM-P client (NEtO)	Any	23	TCP	2 / 4	Telnet 9500 MPR / Wavence Management
NFM-P client	Any	80	TCP	2 / 4	HTTP (GNE / Omni)
NFM-P client (NEtO)	Any	80	TCP	2 / 4	HTTP 9500 MPR / Wavence (MSS-8/4/1 and 9400 AWY)
NFM-P server	Any	161	UDP	2 / 4	SNMP
NFM-P auxiliary server	Any	161	UDP	2 / 4	SNMP NFM-P statistics auxiliary
NFM-P client (NEtO)	Any	161	UDP	2 / 4	SNMP 9500 MPR / Wavence Management
NFM-P client	Any	443	TCP	2 / 4	HTTPS (GNE / Omni)
NFM-P server	>1023	>1023	TCP	2 / 4	Passive FTP transfer
NFM-P auxiliary server	>1023	>1023	TCP	2 / 4	Passive FTP transfer NFM-P statistics auxiliary
NFM-P client (NEtO)	>1023	>1023	TCP	2 / 4	Passive FTP transfer 9500 MPR / Wavence Management
NFM-P server	Any	1491	TCP	2 / 4	SNMP Streaming
NFM-P server	Any	5001	TCP	2 / 4	CPAA / vCPAA
NFM-P client (NEtO)	5010	5010	UDP	2 / 4	SNMP 9500 MPR / Wavence (MSS-8/4/1)
NFM-P client (NEtO)	Any	11500	UDP	2 / 4	Equipment View (GUI) 9500 MPR / Wavence (MSS-1C / MPR-e)

6.10.8 syslog firewall rules

When there is a firewall at the syslog server, the following connection details must be considered.

Table 6-24 Firewall rules for traffic connecting to the syslog server

Source	Source port	Destination port	Protocol	Network Interface	Notes
NFM-P server	Any	As configured	TCP	N/A	Port is configurable

7 NSP deployment with multiple network interfaces and IP addresses

7.1 Support for multiple network interfaces

7.1.1 Introduction

The NSP and its associated components communicate with different entities that typically exist in different network spaces. Isolating different types of traffic to different networks provides better security and helps manage traffic volume on different networks.

NSP supports configuring different network interfaces to handle the following types of traffic in a multi-homed system.

- A client network interface can be used for connecting users to NSP GUI and to connect external OSS systems to NSP.
- An internal network interface can be used to handle traffic between NSP systems that does not need to be accessed by external systems or with managed network elements. Internal traffic includes, but is not limited to, resync of network topology information, security communications, application registration and data synchronization between redundant components.
- A mediation network interface can be used to communicate with network elements (provisioning, NE database backups, monitoring, operations, etc).

7.1.2 Component support and limitations

An NSP cluster can be configured with network interfaces for client traffic, for internal network management traffic, and for managed network traffic. In a multi-node NSP cluster, each node must have the same number of interfaces. Each node of the NSP cluster must have mediation network connectivity to all managed NE devices as MDM pods may be rescheduled on different nodes in the NSP cluster.

The following table describes the supported NSP deployments with multiple network interfaces.

Table 7-1 Supported multi interface NSP deployments

NSP Cluster Node Network Connectivity	NSP traffic management
single network	All NSP traffic uses a single network interface.
two networks, option 1	Client and Internal traffic share one network interface. Mediation traffic uses second network interface.
two networks, option 2	Client traffic uses one network interface. Internal and Mediation traffic share one network interface.

Table 7-1 Supported multi interface NSP deployments (continued)

NSP Cluster Node Network Connectivity	NSP traffic management
three networks	Client traffic uses one interface. Internal traffic uses one interface. Mediation traffic uses one interface.

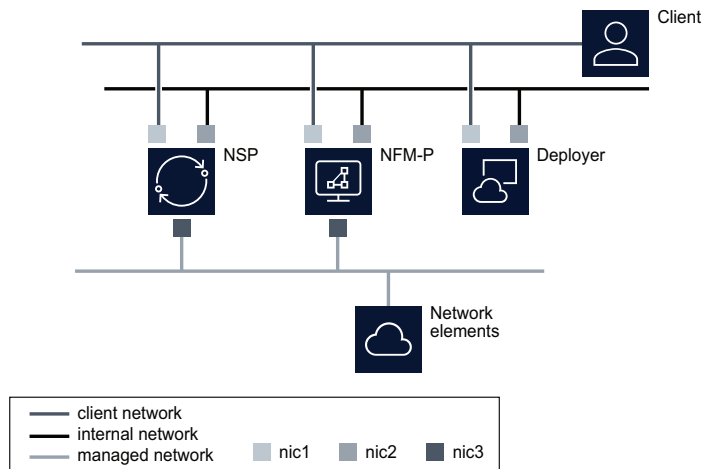
Note: An NSP deployment of client and mediation traffic on one network interface, and internal traffic on a second interface is not supported.

A deployer node must have connectivity to all NSP cluster nodes on the internal network. The deployer node must also have access to the client network.

All servers in an NSP deployment must support the same network configuration for client and internal networks. For example, an NSP cluster deployed with network interfaces for client and internal networks cannot be deployed with a NFMP server that is configured for one interface supporting client and internal communications.

The following figure shows a multi network interface deployment of NSP cluster, deployer host and NFM-P.

Figure 7-1 Multi-interface NSP deployment



36809

An NSP cluster must communicate with a VSR-NRC using the BOF interface.

The NFM-P supports integrated deployment with NSP cluster using network interfaces for client, internal and managed network traffic.

An integrated deployment of multi-interface NSP cluster with an older release NFM-P is supported but a workaround procedure is required on the NFM-P system. The workaround will enable the separation of client and internal network traffic between NFM-P and NSP cluster. See the *NSP Installation and Upgrade Guide* for details on the integration procedure for older release NFM-P.

When installing NSP components on stations with multiple network interfaces, each interface must reside on a separate subnet, with the exception of interfaces that are to be used in IP Bonding.

There is no requirement on the NSP cluster to use the first network interface (eg. eth0, bge0) to communicate with client applications.

Additional network interfaces can be configured on the NSP cluster, at the customer's discretion, for other operations such as archiving database backups or activity logs.

When an NSP cluster is deployed with WS-NOC , the separation of client and internal network traffic is not supported. NSP and WS-NOC must use a single network for client and internal communications.

When using custom TLS certificates in a multi-network configuration, the NSP server certificate requires the IP address or hostname or FQDN of the client network interface (or virtual IP) and the IP address or hostname or FQDN of the internal network interface (or virtual IP) in the certificate SAN field.

For ingress configurations and topology diagrams, please refer to Appendix D in the *NSP Installation and Upgrade Guide*.

7.1.3 Multi-interface support in IPv4 and IPv6 networks

The NSP cluster can use IPv4 or IPv6 addressing on the client, internal and mediation network interfaces. In addition to the limitations and restrictions documented in section 4.2.3, the following conditions apply:

- The NSP cluster can only use IPv4 or IPv6 communications on the client network interface and on the internal network interface (ie. NSP cluster client traffic can not use IPv4 and IPv6 simultaneously). The system network interfaces can be dual stack, but NSP communications on those interfaces can only use IPv4 or IPv6.
- The NSP cluster mediation interface supports IPv4 only, IPv6 only and IPv4 and IPv6 simultaneously for NEs. When NSP is configured with IPv4 and IPv6 mediation on a dual stack network interface, the NSP must have a dedicated mediation interface not shared with client and internal network communications.
- In an NSP deployment with separate network interfaces for client and internal communications, the client and internal networks must both use IPv4 or IPv6 addressing. Example, client communications on IPv4 and internal communications on IPv6 is not supported.

7.1.4 Multi-interface NSP deployment and traffic management policies

Customers can use traffic management applications to protect NSP components but should be applied with care to ensure that NSP applications are not negatively impacted. NSP communications are defined in Section 6.5 of this guide.

The following table summarizes the traffic management policies for an NSP cluster deployment by each network or network interface.

Table 7-2 NSP cluster communications by network interface

Network description	Permitted communications
Client network	Client communications as defined in Table 6-2, “NSP Kubernetes virtual machine communications” (p. 115) Kafka communications on ports 9192, 9193, 9194
Internal network	All communications with NFMP and with redundant datacenter All communications between cluster nodes and deployer node Kafka communications on ports 9292, 9293, 9294
Mediation network	Mediation communications as defined in Table 6-2, “NSP Kubernetes virtual machine communications” (p. 115)

The NSP cluster can communicate with some external components on any network interface, including

- remote authentication servers (LDAP, RADIUS, TACACS)
- VSR-NRC
- syslog server
- email server

Each node in an NSP cluster must allow the same traffic on each network interface.

i **Note:** NSP requires ICMP connectivity to the configured remote authentication servers to verify network connectivity.

7.2 NSP Network Address Translation

7.2.1 Overview

NSP supports the use of Network Address Translation (NAT) between the following components:

- NSP cluster and clients (web application users, REST API clients)
- NSP cluster and network elements
- NSP cluster and other components in the NSP deployment (eg. NFM-P)

NSP does not support the use of NAT between nodes within an NSP cluster deployment, including the deployer host.

The cluster nodes in an NSP deployment must be able to route to the NAT client VIP address of the NSP cluster.

7.3 NFM-P multihoming

7.3.1 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 the entities may 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)).

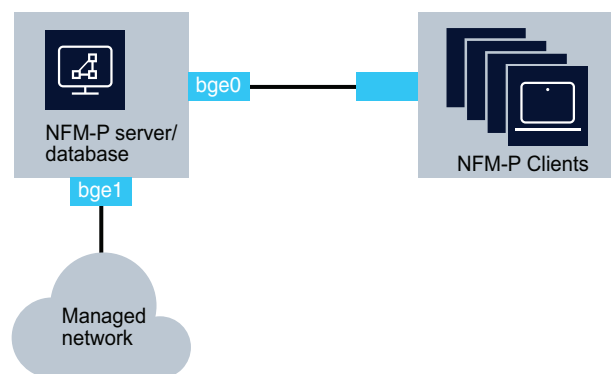
It is a security requirement that all IP communications from an NFM-P auxiliary collector to the NFM-P main server use only one IP address. This IP Address must be the same IP address as the auxiliary collector IP address configured when installing the main server. Any other IP communications originating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

When installing NFM-P components on stations 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 7-2, “Collocated NFM-P server/database deployment with multiple network interfaces”](#) (p. 159) 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 station (for example ce0, bge0) to communicate with the NFM-P GUI clients.

Figure 7-2 Collocated NFM-P server/database deployment with multiple network interfaces



22666

Figure 7-3 Distributed NFM-P server/database deployment with multiple network interfaces

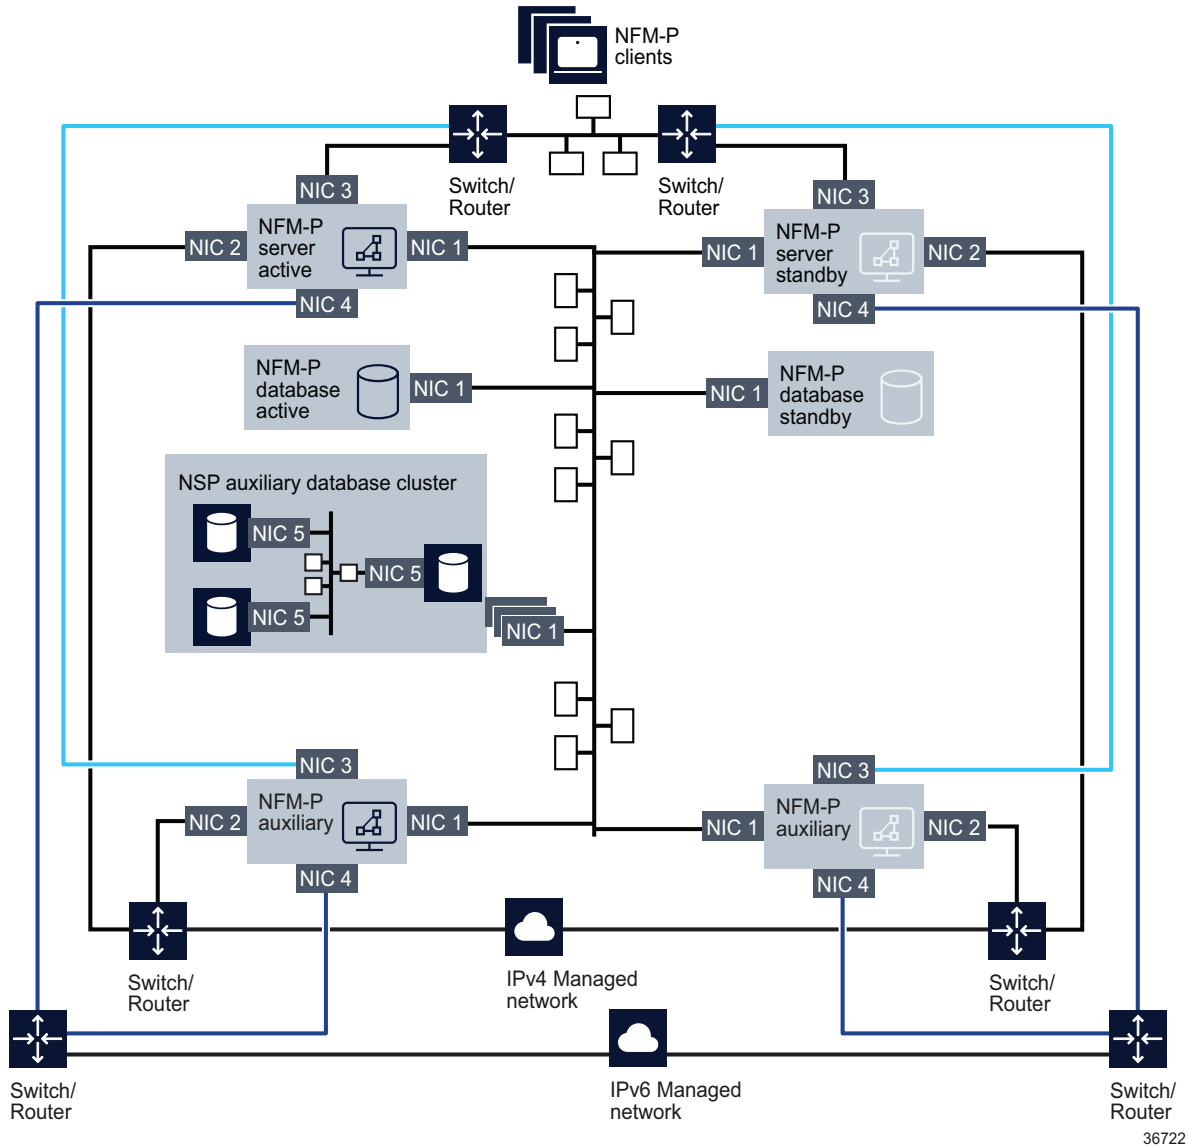


Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” (p. 160) illustrates a distributed, redundant NFM-P deployment where the NFM-P components are configured to actively use more than one network interface.

Due to limitations with the inter-process and inter-station communication mechanisms, a specific network topology and the use of hostnames is required (see 7.5 “Use of hostnames for the NFM-P client” (p. 165)). Contact your Nokia account representative to obtain further details.

7.3.2 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)) 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)). 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 - for example, 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, NSP auxiliary database, and NFM-P auxiliary collectors as well as any redundant NFM-P components should they be present (NIC 1 on [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)). This network interface contains the IP address that the NFM-P database, NSP 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 station, at the customer’s discretion, to perform maintenance operations such as station backups.
- IPv4 and IPv6 network elements can be managed from the same interface or from separate interfaces. (NIC2 and/or NIC4 on [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)).

7.3.3 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)) 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#)

(p. 160)). 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, NSP auxiliary database cluster as well as any redundant NFM-P components should they be present (NIC 1 on [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)). This network interface contains the IP address that the NFM-P server, NFM-P database, NSP 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 station, at the customer’s discretion, to perform maintenance operations such as station backups.
- IPv4 and IPv6 network elements can be managed from the same interface or from separate interfaces. (NIC2 and/or NIC4 on [Figure 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)).

7.3.4 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)) 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces”](#) (p. 160)) 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 station, at the customer’s discretion, to perform maintenance operations such as station backups.

7.3.5 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” \(p. 160\)](#)) 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” \(p. 160\)](#)). 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” \(p. 160\)](#)) 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 7-3, “Distributed NFM-P server/database deployment with multiple network interfaces” \(p. 160\)](#)). 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 station, at the customer’s discretion, to perform maintenance operations such as station backups.

7.4 NFM-P Network Address Translation

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

Figure 7-4 NFM-P server deployments with NAT between the server and the managed network


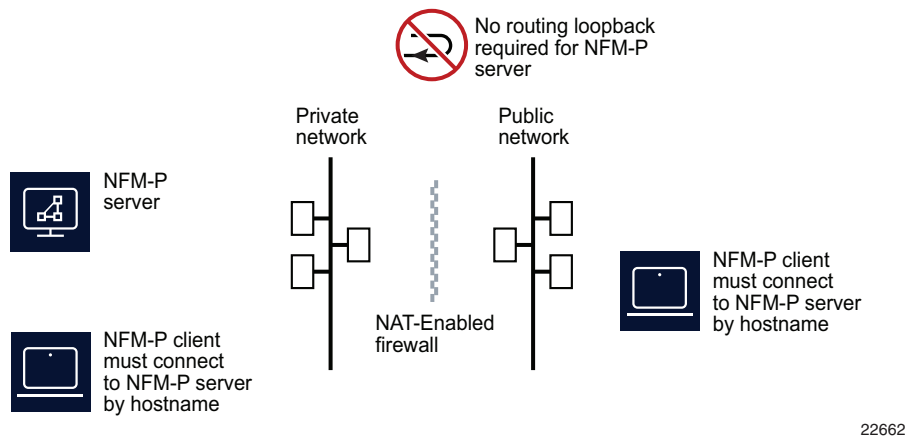


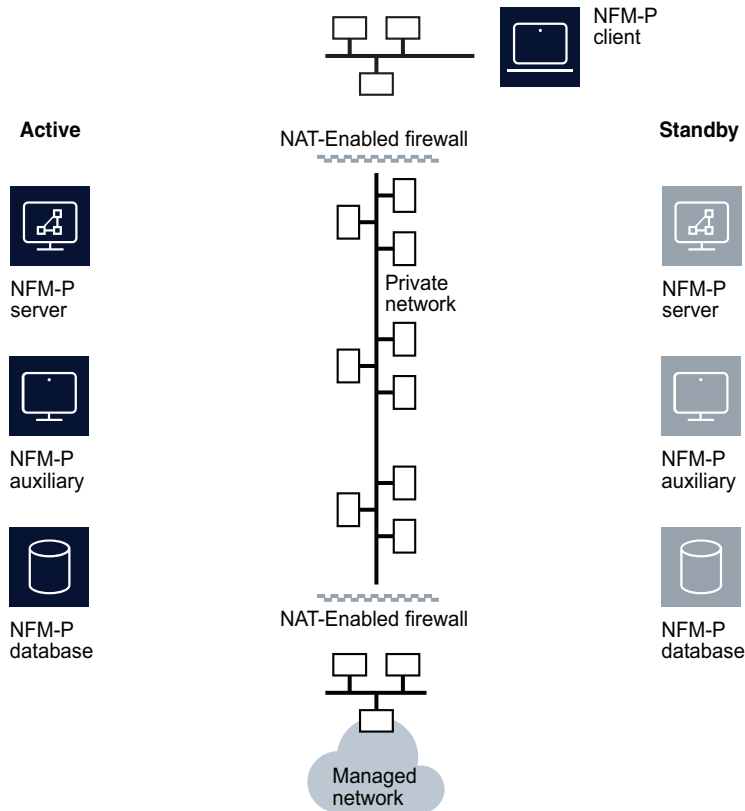
Figure 7-5 NFM-P server deployment using NAT with name resolution based communication



In Figure 7-5, “NFM-P server deployment using NAT with name resolution based communication” (p. 164), 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 7-6 NFM-P deployment with NAT



22661

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.

7.5 Use of hostnames for the NFM-P client

7.5.1 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.

8 Appendix A

8.1 Storage-layer I/O performance tests

8.1.1 Introduction

Use the commands in this section to determine if the storage-layer performance meets the NSP cluster deployment requirements. Both tests (storage latency and storage throughput) must be executed and meet or exceed the provided requirements.

8.1.2 Storage latency

In this example, the /test directory is on the same disk where etcd runs.

Enter the following as the root user to run the test:

```
# fio --rw=write --ioengine=sync --fdatasync=1 --directory=/var/lib/test
--size=22m --bs=2300 --name=mytest ↵
```

The command produces output like the following:

```
Starting 1 process
mytest: Laying out IO file (1 file / 22MiB)
Jobs: 1 (f=1)
mytest: (groupid=0, jobs=1): err= 0: pid=40944: Mon Jun 15 10:23:23 2020
  write: IOPS=7574, BW=16.6MiB/s (17.4MB/s) (21.0MiB/1324msec)
    clat (usec): min=4, max=261, avg= 9.50, stdev= 4.11
      lat (usec): min=4, max=262, avg= 9.67, stdev= 4.12
    clat percentiles (nsec):
      | 1.00th=[ 5536],  5.00th=[ 5728], 10.00th=[ 5920], 20.00th=[
6176],
      | 30.00th=[ 7584], 40.00th=[ 8896], 50.00th=[ 9408], 60.00th=[
9792],
      | 70.00th=[10432], 80.00th=[11584], 90.00th=[12864], 95.00th=
[14528],
      | 99.00th=[20352], 99.50th=[23168], 99.90th=[28800], 99.95th=
[42752],
      | 99.99th=[60672]
    bw ( KiB/s): min=16868, max=17258, per=100.00%, avg=17063.00,
stdev=275.77, samples=2
    iops        : min= 7510, max= 7684, avg=7597.00, stdev=123.04,
samples=2
```

```
lat (usec) : 10=64.21%, 20=34.68%, 50=1.08%, 100=0.02%, 500=0.01%
```

In the second block of output, which is shown below, the 99.00th percentile must be less than 10 ms. In this test output, the 99.00th percentile is less than 1 ms.

```
fsync/fdatasync/sync_file_range:
  sync (usec): min=39, max=1174, avg=120.71, stdev=63.89
  sync percentiles (usec):
    | 1.00th=[ 42], 5.00th=[ 45], 10.00th=[ 46], 20.00th=[
48],
    | 30.00th=[ 52], 40.00th=[ 71], 50.00th=[ 153], 60.00th=[
159],
    | 70.00th=[ 167], 80.00th=[ 178], 90.00th=[ 192], 95.00th=[
206],
    | 99.00th=[ 229], 99.50th=[ 239], 99.90th=[ 355], 99.95th=[
416],
    | 99.99th=[ 445]
  cpu : usr=2.95%, sys=29.93%, ctx=15663, majf=0, minf=35
  IO depths : 1=200.0%, 2=0.0%, 4=0.0%, 8=0.0%, 16=0.0%, 32=0.0%,
>=64=0.0%
  submit : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%,
>=64=0.0%
  complete : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%,
>=64=0.0%
  issued rwts: total=0,10029,0,0 short=10029,0,0,0 dropped=0,0,0,0
  latency : target=0, window=0, percentile=100.00%, depth=1
```

8.1.3 Storage throughput

To run this test, first change to the directory where the test is to be performed. The test will create a local file. The output from the command contains read and write IOPS values which must be evaluated against the minimum requirements provided in [Table 2-7, “Minimum NSP cluster IOPS requirements - local storage” \(p. 31\)](#)

Enter the following as the root user to run the test:

```
# fio --randrepeat=1 --ioengine=libaio --direct=1 --gtod_reduce=1
--name=test --filename=random_read_write.fio --bs=4k --iodepth=64
--size=4G --readwrite=randrw --rwmixread=50 ↵
```

The command produces output like the following:

```
test: (g=0): rw=randrw, bs=(R) 4096B-4096B, (W) 4096B-4096B, (T)
4096B-4096B, ioengine=libaio, iodepth=64
fio-3.7
```

```
Starting 1 process
test: Laying out IO file (1 file / 4096MiB)
Jobs: 1 (f=1): [m(1)][100.0%][r=22.1MiB/s,w=22.2MiB/s][r=5645,w=5674
IOPS][eta 00m:00s]
test: (groupid=0, jobs=1): err= 0: pid=32439: Mon Sep 21 10:25:11 2020
read: IOPS=6301, BW=24.6MiB/s (25.8MB/s) (2049MiB/83252msec)
    bw ( KiB/s): min=13824, max=39088, per=99.57%, avg=25098.60,
stdev=5316.27, samples=166
    iops        : min= 3456, max= 9772, avg=6274.49, stdev=1329.11,
samples=166
write: IOPS=6293, BW=24.6MiB/s (25.8MB/s) (2047MiB/83252msec)
    bw ( KiB/s): min=13464, max=40024, per=99.56%, avg=25062.73,
stdev=5334.65, samples=166
    iops        : min= 3366, max=10006, avg=6265.57, stdev=1333.67,
samples=166
    cpu         : usr=5.13%, sys=18.63%, ctx=202387, majf=0, minf=26
    IO depths   : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%,
>=64=100.0%
    submit     : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%,
>=64=0.0%
    complete   : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%,
>=64=0.0%
    issued rwts: total=524625,523951,0,0 short=0,0,0,0 dropped=0,0,0,0
    latency    : target=0, window=0, percentile=100.00%, depth=64
Run status group 0 (all jobs):
    READ: bw=24.6MiB/s (25.8MB/s), 24.6MiB/s-24.6MiB/s (25.8MB/s-25.
8MB/s), io=2049MiB (2149MB), run=83252-83252msec
    WRITE: bw=24.6MiB/s (25.8MB/s), 24.6MiB/s-24.6MiB/s (25.8MB/s-25.
8MB/s), io=2047MiB (2146MB), run=83252-83252msec
Disk stats (read/write):
    vda: ios=523989/526042, merge=0/2218, ticks=3346204/1622070,
in_queue=4658999, util=96.06%
```

