NOKIA

NSP
Network Services Platform

Release 22.3

System Administrator Guide

3HE-18167-AAAA-TQZZA

Issue 2

April 2023

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.

Contents

About this document

Purpose

The *NSP System Administrator Guide* is intended for operators who have NSP system administrator privileges and need to understand or perform Network Services Platform system management and maintenance. The guide describes how to perform operations for system and component configuration, security, application access, and database management.

Scope

The scope of this document is limited to NSP system administration. Readers of the guide are advised to familiarize themselves with the different aspects of the administration process. Each chapter or section describes a specific area of interest or administrative function.

Safety information

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

Document support

Customer documentation and product support URLs:

- Documentation Center
- Technical support

How to comment

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

1 NSP administration overview

1.1 Introduction

1.1.1 Guide description

The *NSP System Administrator Guide* describes how to perform various NSP management operations as requirements arise, or as directed by technical support.

The guide is written for an NSP operator who has the NSP administrator role assigned to their NSP user group. For information about NSP role-based user management, see Chapter 2, “NSP security administration”.

1.1.2 NSP administrator responsibilities

An NSP system administrator can manage all NSP functional areas, and is primarily responsible for the following:

- system security, such as TLS configuration and user management, as described in Chapter 2, “NSP security administration”
- application setup, as described in Chapter 3, “NSP application administration”; application usage is described in the online application help
- starting, stopping, and configuring system components, as described in Chapter 4, “NSP global system administration”
- database management, such as restoration after a failure, as described in Chapter 6, “NSP database administration”

Note: It is strongly recommended that you perform an administrative procedure in this guide only under the guidance of technical support.

2 NSP security administration

2.1 NSP security administration overview

2.1.1 Introduction

This chapter describes the following NSP security areas of interest:

- NSP TLS implementation, and operations such as certificate renewal; see 2.2 “NSP TLS administration” (p. 11)
- SELinux administration; see 2.4 “SELinux implementation and management” (p. 17)
- NSP user password management; see 2.8 “User password policies” (p. 26)
- user-activity and application log forwarding; see 2.9 “User activity log forwarding” (p. 27)
- occasional NSP system security operations; see 2.10 “General NSP security administration procedures” (p. 28)

User access control, session management, and activity logging are documented in the *NSP User Manager Application Help*.

2.2 NSP TLS administration

2.2.1 NSP TLS administration overview

The NSP uses TLS to secure the following:

- Kubernetes infrastructure; see 2.2.2 “Kubernetes infrastructure TLS” (p. 12)
- internal NSP subsystems and services; see 2.2.3 “NSP internal TLS” (p. 12)
- external interfaces between NSP components and for client access; see 2.2.4 “NSP external TLS” (p. 12)

The NSP includes a Public Key Infrastructure, or PKI server, to distribute TLS certificates. A PKI server can generate internal and external certificates using private root CA certificates. See the *NSP Deployment and Installation Guide* for more information about NSP TLS deployment using a PKI server.

Support for deprecated TLS versions

By default, the NSP uses TLS 1.2. External systems such as OSS clients may use older versions, which are deprecated. You can enable the older TLS versions in the NSP for compatibility with such systems.

Note: You must enable support for the deprecated TLS versions in a shared-mode NSP system that includes a Release 20 NFM-P or NFM-T, if the **secure** parameter in the **nspos** section of the NSP configuration file is set to true.

The following parameters enable or disable the support for the deprecated TLS versions:

- NSP cluster—tlsv1ProtocolsEnabled, in the nsp-config.yml file

- IP resource control and cross-domain resource control—tlsv1_protocols_enabled, in the config.yml file

2.2.2 Kubernetes infrastructure TLS

The TLS certificates that secure the NSP Kubernetes infrastructure expire one year from installation. See 2.2.5 “Managing NSP TLS certificates” (p. 12) for information about Kubernetes infrastructure certificate expiry notifications and certificate renewal.

2.2.3 NSP internal TLS

The internal TLS certificate secures the internal NSP processes. For maximum security, an NSP PKI server uses an internally generated private root CA to create the certificate. Consequently, no certificate from any external CA is trusted for access to system processes.

A PKI server generates an internal certificate automatically during initialization. During an NSP system installation or upgrade, the NSP PKI server must be running in order for each component to request and receive an internal certificate, as described in each NSP and NFM-P installation and upgrade procedure.

When you add or replace an NSP system element such as an NSP Flow Collector or an NFM-P component, the PKI server provides an internal certificate during initialization, as described in the component installation procedure.

Note: To reduce complexity, each upgrade procedure instructs you to start the PKI server, regardless of the upgrade conditions.

2.2.4 NSP external TLS

The external TLS certificate secures the NSP interfaces used by clients and external systems. The certificate can be signed by an external CA, or by the private root CA of an NSP PKI server.

2.2.5 Managing NSP TLS certificates

NSP internal or external TLS certificate replacement may be required when:

- a certificate nears or reaches expiry
- a component is added to the NSP system
- an NSP component is replaced
- an NSP component address changes

2.3 “To replace the internal or external NSP TLS certificate” (p. 13) describes how to replace the internal TLS certificate, the external certificate, or both, in an NSP system.

Kubernetes infrastructure TLS certificate renewal

The NSP monitors the Kubernetes infrastructure TLS certificates for expiry, and automatically renews the certificates; no administrative action is required.

Internal certificate replacement

To replace the internal certificate used in an NSP system, you must start the PKI server, enable the NSP to regenerate internal certificates, and then run the installation script.

External certificate replacement

The external certificate replacement method depends on the TLS deployment method:

- Manual—The replacement process is the same as the manual TLS deployment process described in the *NSP Deployment and Installation Guide*.
- Automated—The following options are available for generating private root-CA-signed certificates.
  - Provide a set of TLS key and certificate files to the PKI server for signing certificate requests from NSP components during deployment or configuration.
  - Start the PKI server without providing a TLS file set. The PKI server prompts the operator for certificate parameters, signs the certificate using the embedded private root CA, and then generates a TLS file set.

Note: Some system conversion or migration operations may include additional TLS configuration requirements; see the *NSP Deployment and Installation Guide* for more information.

TLS certificate expiry notifications

The NSP checks the expiry date of each TLS certificate during initialization, and every 24 hours thereafter. After an NSP TLS certificate expires, the NSP cluster continues to operate, but functions that depend on secure communication are unavailable.

When a certificate expires or approaches expiry, the NSP raises one of the following server or internal certificate alarms:

- Warning, if the certificate is to expire within 30 days of the current time
- Critical, if the certificate is to expire within 7 days of the current time
- Critical, if the certificate is expired

Note: The NSP raises one alarm per certificate.

Note: The alarms for internal or external NSP certificate expiry do not clear automatically.

Note: An expiry alarm for an NSP Kubernetes infrastructure certificate clears automatically 60 minutes after the certificate renewal.

Note: The Days Remaining value in an expiry alarm is based on the number of complete 24-hour periods until the certificate expiry time. If fewer than 24 hours remain until expiry, the Days Remaining value is zero; however, the NSP does not raise an alarm about the certificate expiry until the next periodic check, 24 hours later.

2.3 To replace the internal or external NSP TLS certificate

2.3.1 Purpose

Perform this procedure to replace the PKI-server-generated TLS certificates, or custom CA TLS certificates, or both, in a NSP system.

Note: You must perform the procedure on each NSP cluster in a DR deployment.

Note: You require root user privileges on each NSP cluster VM in each data center.

Note: The install.sh utility requires SSH access to a target station. To enable SSH access, you must do one of the following.

- Configure the required SSH keys on the stations.
- If each remote station has the same root user password, add the --ask-pass argument to the install.sh command; for example:

  ./install.sh --ask-pass --target *remote_station*

Note: *release-ID* in a file path has the following format:

R.r.p-rel.*version*

where

R.r.p is the NSP release, in the form *MAJOR.minor.patch*

version is a numeric value

2.3.2 Steps

1

Log in as the root user on the NSP Configurator VM.

2

Open a console window.

3

Configure the NSP to preserve the existing deployment.

1. Open the following file using a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml
2. Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

   ```
   deleteOnUndeploy:false
   ```
3. Save and close the file.
4. Enter the following:

   ```
   # /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --undeploy ↵
   ```

4

If you are changing the deployment, such as adding or removing a component, or changing a component address, update the NSP configuration.

1. Edit the following file as required using a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml
2. Save and close the file.

5

If you are updating the PKI-server-generated TLS certificates, copy the new CA certificate files to the following directory

/opt/nsp/NSP-CN-*release-ID*/tls/ca

6

If you are updating the custom-CA-signed TLS certificates, modify the NSP configuration to include the new certificate information.

1. Open the following file using a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml
2. Edit the parameters in the **tls** section as required.
3. Save and close the file.

7

Enter the following:

```
# nsp-config.bash -config --deploy ↵
```

8

If the NSP deployment does not include IP resource control or cross-domain resource control, go to Step 12.

9

If you need to update the PKI-generated TLS certificates for IP resource control or cross-domain resource control, perform the following steps.

1. Enter the following:

   ```
   # cd /opt/nsp/NSP-CN-release-ID/tools/pki ↵
   ```
2. Enter the following:

   ```
   # cp ../../tls/ca/ca* . ↵
   ```
3. Enter the following:

   ```
   # ./pki-server & ↵
   ```

   The PKI server starts.
4. Log in as the root user on the station that has the resource-control installation software.
5. Open the following file using a plain-text editor such as vi:

   NSP_installer_directory/config/config.yml

   where *NSP_installer_directory* is the directory that contains the installation files
6. Configure the following parameters in the **tls** section as shown below:

   ```
   pki_server: PKI_server_address
   regenerate_certs: true
   ```

   where *PKI_server_address* is the IP address of the NSP configurator VM

7. Save and close the file.

10

If you need to update the custom-CA-signed TLS certificates for IP resource control or cross-domain resource control, perform the following steps.

1. Log in as the root user on the station that has the resource-control installation software.
2. Open the following file using a plain-text editor such as vi:

   NSP_installer_directory/config/config.yml

   where *NSP_installer_directory* is the directory that contains the installation files
3. Configure the following parameters in the **tls** section using the new values:

   ```
   custom_key_alias
   custom_keystore_path
   custom_truststore_path
   ```
4. Save and close the file.

11

Enter the following to update the configuration of the IP resource control and cross-domain resource control servers:

Note: If you do not include the --extra-vars option and required parameters, the installer prompts you for the TLS keystore and truststore passwords.

Note: The --target parameter is optional, if not specified, all IP resource control and cross-domain resource control servers are updated. If you include the parameter, only the servers with the specified IP addresses are updated.

```
# ./install.sh --ask-pass --target target_list --extra-vars
"variable_1=value_1,variable_2=value_2" ↲
```

where

target_list is a comma-separated list of server IP addresses

variable_1=value_1 and *variable_2=value_2* are the following:

- tls_custom_keystore_password=*TLS_keystore_password*
- tls_custom_truststore_password=*TLS_truststore_password*

Note: A special character in a password value must be escaped using a backslash [\] character.

You are prompted for the root password of each IP resource control or cross-domain resource control server, and the configuration is updated.

12

Configure each other NSP component to obtain the updated TLS configuration.

For information about configuring TLS for components such as NSP Flow Collectors, Flow Collector Controllers, or analytics servers, see the *NSP Deployment and Installation Guide*.

For information about configuring TLS for other components and products such as the NFM-P or NFM-T, see the specific component or product documentation.

13

If the PKI server is running, enter Ctrl+C in the NSP Configurator VM console window to stop the PKI server.

Note: You must not stop the PKI server until each NSP component has obtained the updated certificates from the PKI server.

14

Close the open console windows.

END OF STEPS

2.4 SELinux implementation and management

2.4.1 Introduction

For greater system security, you can enable RHEL SELinux on NSP components. SELinux logs user operations in Application Visibility and Control, or AVC messages that are stored in local logs. SELInux has two modes, permissive and enforcing; the support for each is described in 2.4.2 “SELinux support scope” (p. 17).

See the RHEL documentation for comprehensive SELinux configuration and implementation information.

SELinux permissive mode

No SELinux policy is enforced in permissive mode, and no operations are denied. However, SELinux does log AVC messages while in permissive mode. AVC messages may be of use for troubleshooting, debugging, and SELinux policy improvements. An AVC message is logged each time a violation occurs.

SELinux enforcing mode

In enforcing mode, SELinux enforces the policies specified in the NSP SELinux configuration, and logs AVC messages as required.

2.4.2 SELinux support scope

You can enable SELinux in enforcing mode only on the following:

- NFM-P main servers
- NFM-P main databases
- NFM-P auxiliary servers, excluding PCMD and call-trace

SELinux is supported only in permissive mode on the following:

- NSP cluster VMs
- NSP Flow Collectors, Flow Collector Controllers

- NSP analytics servers
- NFM-P auxiliary databases

2.4.3 SELinux troubleshooting

In the event that an NFM-P system in SELinux enforcing mode has functional issues and an AVC is present, a change to permissive mode may resolve the issue. If enabling permissive mode resolves the issue, and the AVC is in the NSP domain, it is strongly recommended that you raise a support ticket to report the AVC.

2.5 To enable SELinux in permissive mode

2.5.1 Purpose

CAUTION

Service Disruption

Enabling SELinux in a standalone or redundant NFM-P system creates a network management outage. A standalone system requires a full shutdown and restart; a redundant system requires one or more server activity switches that each may cause a brief service interruption.

Perform the procedure only during a scheduled maintenance period of sufficient duration with the guidance of technical support.

Perform this procedure to enable SELinux in permissive mode on the components of an NFM-P system.

Note: You must enable permissive mode on the components before you can enable enforcing mode on the components.

Note: You require the following user privileges:

- on each main and auxiliary server station — root, nsp
- on each main database station — root

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

2.5.2 Steps

Check for required OS packages

1

Before you can enable SELinux on a station, you must ensure that the required RHEL OS packages are installed.

Perform the following steps on each main server, main database, and auxiliary server station.

1. Log in to the station as the root user.
2. Open a console window.
3. Enter the following:

   ```
   # /opt/nsp/nfmp/config/selinux/tools/bin/selinuxenable.sh -c ↵
   ```

   A status message is displayed.
4. If the message indicates that one or more required SELinux packages are not installed, enter the following:

   ```
   # yum -y install policycoreutils setools-console libselinux-devel
   setroubleshoot-server selinux-policy-devel selinux-policy-doc ↵
   ```

   The packages are installed.

Close client sessions

2

Close the open GUI and XML API client sessions, as required.

1. Open a GUI client using an account with security management privileges, such as admin.
2. Choose Administration→Security→NFM-P User Security from the main menu. The NFM-P User Security - Security Management (Edit) form opens.
3. Click on the Sessions tab.
4. Click Search. The form lists the open GUI and XML API client sessions.
5. Identify the GUI session that you are using based on the value in the Client IP column.
6. Select all sessions except for the following:
   - the session that you are using
   - the sessions required to monitor the network during a redundant system upgrade
7. Click Close Session.
8. Click Yes to confirm the action.
9. Click Search to refresh the list and verify that only the required sessions are open.
10. Close the NFM-P User Security - Security Management (Edit) form.
11. Close your GUI client.
12. Sign out of the NSP Launchpad, if you are signed in.

3

If the NFM-P system is standalone:

1. Perform Step 5 to Step 10 on the main server, man database, and statistics-collection auxiliary servers.
2. Go to Step 13.

4

If the NFM-P system is redundant:

1. Perform Step 5 to Step 12 on the standby server complex.

   After this step, the initial standby server complex is the new primary complex.
2. Perform Step 5 to Step 10 on the initial primary server complex, which is the new standby server complex.
3. If you want to restore the initial primary and standby roles of the server complexes, go to Step 11. Otherwise, go to Step 13.

Stop system components

5

Stop the main server.

1. Log in to the main server station as the nsp user.
2. Open a console window.
3. Enter the following:

   ```
   bash$ cd /opt/nsp/nfmp/server/nms/bin ↵
   ```
4. Enter the following:

   ```
   bash$ ./nmsserver.bash stop ↵
   ```
5. Enter the following:

   ```
   bash$ ./nmsserver.bash appserver_status ↵
   ```

   The server status is displayed; the server is fully stopped if the status is the following:

   ```
   Application Server is stopped
   ```

   If the server is not fully stopped, wait five minutes and then repeat this step. Do not perform the next step until the server is fully stopped.
6. Enter the following to switch to the root user:

   ```
   bash$ su - ↵
   ```
7. Enter the following:

   ```
   # systemctl stop nspos-nspd.service ↵
   ```

6

Stop the Oracle proxy and database services.

1. Log in to the database station as the root user.
2. Open a console window.
3. Enter the following to stop the Oracle proxy:

   ```
   # systemctl stop nfmp-oracle-proxy.service ↵
   ```
4. Enter the following to stop the main database:

   ```
   # systemctl stop nfmp-main-db.service ↵
   ```

7

If the system includes one or more statistics-collection auxiliary servers, stop each such auxiliary server.

1. Log in to the auxiliary server station as the nsp user.
2. Open a console window.
3. Enter the following:

   bash$ **/opt/nsp/nfmp/auxserver/nms/bin/auxnmsserver.bash auxstop** ↵

   The auxiliary server stops.

Enable SELinux on components

8

Enter the following as the root user on each main server, main database, and statistics-collection auxiliary server station:

/opt/nsp/nfmp/config/selinux/tools/bin/selinuxenable.sh -p ↵

Apply SELinux labels and reboot

9

Perform the following steps as the root user on each main server, main database, and statistics-collection auxiliary server station.

1. Enter the following:

   # **/opt/nsp/nfmp/config/selinux/installer/bin/nsp-selinux-config.bash** ↵

2. Enter the following:

   # **systemctl reboot** ↵

   The stations reboots.

   After the reboot, the SELinux labels are applied to both the RHEL OS and NFM-P processes, and SELinux is running in targeted permissive mode.

Verify system startup

10

After each station is rebooted, verify that the main server, main database, and auxiliary servers are operational.

Note: If any command in a substep indicates that the component is not yet operational, wait one minute and then re-issue the command.

1. Enter the following as the root user on the main database station:

   # **systemctl status nfmp-main-db.service** ↵

   If the command output includes the following, the database is operational:

```
Active: active (running) since time
```

2. Enter the following as the root user on the main database station:

   ```
   # systemctl status nfmp-oracle-proxy.service ↵
   ```

   If the command output includes the following, the database proxy is operational:

   ```
   Active: active (running) since time
   ```

3. Enter the following as the nsp user on the main server station:

   ```
   bash$ ./nmsserver.bash appserver_status ↵
   ```

   If the command output includes the following, the main server is operational:

   ```
   Application Server process is running.  See nms_status for more
   detail.
   ```

4. On each statistics-collection auxiliary server station, enter the following as the nsp user:

   ```
   bash$ ./auxnmsserver.bash auxappserver_status ↵
   ```

   If the command output includes the following, the auxiliary server is operational:

   ```
   Auxiliary Server process is running.  See auxnms_status for more
   detail.
   ```

Switch redundancy roles

11

If automatic database realignment is not enabled, perform a database switchover.

1. As the nsp user on the main server station, enter the following:

   ```
   bash$ /opt/nsp/nfmp/server/nms/bin/switchoverdb.bash -u username -p
   password ↵
   ```

   where *username* and *password* are the login credentials of an NFM-P user with the required privilege level and scope of command

   The script displays the following confirmation message:

   ```
   The standby database will become the new primary database, and the
   old primary will become the new standby. Do you want to proceed?
   (YES/no) :
   ```

2. Enter the following to initiate the switchover:

   ```
   YES ↵
   ```

   The NFM-P server initiates a database switchover. Progress is indicated by a rolling display of dots in the console window. The database switchover is complete when the CLI prompt reappears.

12

Enter the following to perform a server activity switch:

```
bash$ /opt/nsp/nfmp/server/nms/bin/nmsserver.bash force_restart ↵
```

The server activity switch begins. The standby main server restarts as the primary main server, and the primary restarts as the standby.

13

Close the open console windows.

END OF STEPS

2.6 To enable SELinux enforcing mode

2.6.1 Purpose

CAUTION

Potential Security Risk

Enabling SELinux enforcing mode when any AVCs remain unresolved may pose a security risk.

Before you attempt to enable enforcing mode, you must resolve any AVCs associated with the nsp_domain_t domain that are raised during a soak period in permissive mode.

It is strongly recommended that the system run in permissive mode for at least seven days with no nsp_domain_t AVCs on any NFM-P main server, main database, or auxiliary server.

Perform this procedure to enable SELinux enforcing mode in an NFM-P system.

Note: You must perform the procedure on each component that supports SELinux enforcing mode, as listed in 2.4.2 “SELinux support scope” (p. 17).

Note: You must enable permissive mode on each component, as described in 2.5 “To enable SELinux in permissive mode” (p. 18), before you can enable enforcing mode on the components.

Note: You do not need to stop any NFM-P processes in order to switch from permissive to enforcing mode.

Note: You require root user privileges on each station.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

2.6.2 Steps

1

Log in to the component station as the root user.

2

Open a console window.

3

Enter the following:

```
# cd /opt/nsp/nfmp/config/selinux/tools/bin ↵
```

4

Enter the following to list the AVCs:

```
# ./setroubleshoot.bash collect-avcs ↵
```

The AVCs are listed.

5

If any NSP-domain AVCs are listed, perform the following steps to resolve the AVCs.

1. Enter the following:

   ```
   # ./setroubleshoot.bash resolve-nsp-avcs policy_module ↵
   ```

   where *policy_module* is a policy module file name other than nsp_domain

   A list of commands is displayed.

2. Enter each listed command.

The AVCs are resolved.

6

Enter the following:

```
# ./selinuxenable.sh -e ↵
```

7

Enter the following:

```
# getenforce ↵
```

The SELinux mode is displayed.

8

View the command output to verify that SELinux is enabled in enforcing mode.

9

Close the console window.

END OF STEPS

2.7 To troubleshoot SELinux

2.7.1 Purpose

Perform this procedure if SELinux enforcing mode is enabled and you suspect that SELinux is affecting NFM-P operation.

Note: The procedure applies only to the components that support SELinux enforcing mode, as listed in 2.4.2 “SELinux support scope” (p. 17).

Note: You must perform the procedure on each station that has SELinux enforcing mode enabled.

Note: You require root user privileges on each station.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

2.7.2 Steps

1

Log in as the root user on the standalone or primary NFM-P main server station.

2

Open a console window.

3

Enter the following:

```
# cd /opt/nsp/nfmp/config/selinux/tools/bin ↵
```

4

Switch to SELinux permissive mode.

Note: The NFM-P main server can remain running during the switch from enforcing to permissive mode.

1. Enter the following:

   ```
   # ./selinuxenable.sh -p ↵
   ```

2. Enter the following to verify that SELinux is in permissive mode:

   ```
   # getenforce ↵
   ```

   The SELinux mode is displayed.

3. View the command output to verify that SELinux is enabled in permissive mode.

5

To list all system and NSP-domain AVCs, enter the following:

```
# ./setroubleshoot.bash collect-avcs ↵
```

The AVCs are listed.

6

To resolve the NSP-domain AVCs, perform the following steps.

1. Enter the following:

   ```
   # ./setroubleshoot.bash resolve-nsp-avcs policy_module ↵
   ```

   where *policy_module* is a policy module file name other than nsp_domain

A list of commands is displayed.

2. Enter each listed command.

The AVCs are resolved.

7

To resolve the AVCs in all domains, enter the following:

1. # **./setroubleshoot.bash resolve-all-avcs *policy_module*** ↵

   where *policy_module* is a policy module file name other than nsp_domain

   A list of commands is displayed.

2. Enter each listed command.

The AVCs are resolved.

8

Close the console window.

END OF STEPS

2.8 User password policies

2.8.1 Introduction

When an operator attempts to sign in to the NSP Launchpad and a password change is required, the new password must conform to the password policy of the authenticating agent, as described in the following table.

Application	Requirement
NFM-P	When an NFM-P-authenticated user is prompted to change their password during an NSP login attempt, the new password must conform to the NFM-P password requirements. See the *NSP NFM-P Administrator Guide* for the NFM-P password requirements and expiration policy.
NFM-T	When an NFM-T-authenticated user is prompted to change their password during an NSP login attempt, the password must conform to the NFM-T password requirements, which are described in the Common Functions section of the *NFM-T Administration Guide*.
LDAP, RADIUS and TACACS+	A password-change policy is not applied during an NSP user login attempt. If a password change is required, the user must contact the system administrator for information about the LDAP, RADIUS, or TACACS+ password requirements.

2.9 User activity log forwarding

2.9.1 User activity log forwarding overview

If the forwarding of NSP user activity logs to a remote server is enabled, each NSP user action is forwarded to a remote syslog server specified in the NSP configuration during system deployment.

2.9.2 User activity syslog record format

Each generated remote syslog message for user activity has the following fields:

- timestamp
- hostname of syslog producer
- program name
- User Activity Log entry

User Activity Log syslog record example

The following is an example of an NFM-P User Activity Log record forwarded to a remote syslog server:

Note: The record is displayed as three separate sections for illustration purposes; an actual record is contiguous.

```
May 18 09:56:36 nsp-1a3 activitylogs: {"app":"User Manager","clientHost":
"203.0.100.5","reqMethod":"POST","addlParams":"{}","actionParams":
[{"val":"

{\"retentionPeriod\":32,\"activityLogsMaxSize\":1000000,
\"activityLogsWarningThreshold\":95,\"activityLogsCriticalThreshold\":
100,\"activityLogsWarningPurgePercent\":5,
\"activityLogsCriticalPurgePercent\":10}

","key":"jsonRequest"}],"respCodePhrase":"OK","timeStamp":"2020/05/27
10:47:14 821 +0000","affObjs":"{}","uid":
"a0d3b09f66acb238d9f95ab1155d075e","host":"198.51.100.16","action":"set",
"time":"1590576434821","user":"admin","reqURL":"https://198.51.100.
16/activitylogs-api/rest/api/v1/activityLogs/settings/set","respCode":
"200"}
```

The fields in the example have the following values; the actionParams section, which is the second section in the example, indicates that the action involved setting user-activity log parameters:

- timestamp—May 18 09:56:36
- hostname of syslog entry producer—nsp-1a3
- program name—activitylogs
- User Activity Log entry—remainder that begins with "app":"User Manager"; is in JSON format, and includes the following:
  - app—application name from which action performed
  - clientHost—remote hostname or IP address that invokes action
  - reqMethod—type of action performed

- actionParams—array; contains parameters passed to action
- addlParams—array; contains parameters or other such values not in other fields
- respCodePhrase—human-readable action response code
- timeStamp—time at which action completed
- affObjs—array of affected-object attributes, for example, FDN and ID
- uid—record ID
- host—IP address of server on which action performed
- action—name of action performed
- user—username under which action performed
- reqURL—HTTP URL of the executed HTTP request
- respCode—action response code, in integer format

2.10 General NSP security administration procedures

2.10.1 Introduction

The following procedures describe NSP system security actions that may occasionally be required.

2.11 To change the nsp system user password

2.11.1 Purpose

Perform this procedure to change the password of the RHEL nsp user on a cross-domain resource control or IP resource control server.

2.11.2 Steps

1

Log in as the root user on a cross-domain resource control or IP resource control server.

2

Open a console window.

3

Enter the following:

```
# passwd nsp ↵
```

The following prompt is displayed:

```
New Password:
```

4

Enter the new password and press ↵.

The following prompt is displayed:

```
Confirm New Password:
```

5

Enter the new password again and press ↵. The password is changed.

6

Record the new password and store it in a secure location.

7

Close the console window.

END OF STEPS

2.12 To add an NSP analytics server to the allowlist for OSS report requests

2.12.1 Purpose

When an OSS embeds reports from the Analytics application in its own OSS web application, requests to retrieve these reports may be identified as cross-origin requests. Such requests are blocked by the NSP CORS policy, as only the hosts in the allowlist are accepted as NSP clients. Perform this procedure to disable the blocking of report retrieval by an OSS.

2.12.2 Steps

1

Log in as the nsp user on the active IP resource control server.

2

Open a console window.

3

Enter the following to acquire an access token from the REST API Gateway:

```
bash$ curl --insecure -X POST https:
//server/rest-gateway/rest/api/v1/auth/token -H 'authorization: Basic
credentials' -H 'cache-control: no-cache' -H 'content-type:
application/x-www-form-urlencoded' -d grant_type=client_credentials ↵
```

where

server is the IP address or hostname of the server that hosts the active nspOS instance

credentials is the base64-encoded credentials, expressed as *username*:*password*

The REST API Gateway returns an access token.

4

Enter the following once for each NSP analytics server to add the server as an allowlist target:

Note: You must add each NSP analytics server in the deployment to the allowlist.

```
bash$ curl -kv https:
//server/session-manager/api/v1/whitelist/allowedHosts -H
'Content-Type: application/json' -H "Authorization: Bearer
access_token" --data '{"host":"allowlist_target"}' -X POST ↵
```

where

server is the IP address or hostname of the server that hosts the active nspOS instance

access_token is the access token returned in Step 3

allowlist_target is the hostname or IP address of the NSP analytics server

The NSP analytics server address is added to the allowlist.

5 Close the console window.

END OF STEPS

3 NSP application administration

3.1 NSP application administration overview

3.1.1 Introduction

3.2 “NSP application access” (p. 31) describes NSP application access requirements and best practices. The chapter also includes procedures for configuring global application settings and application access.

3.9 “Syslog record format for NSP application logging” (p. 40) describes the syslog record format for NSP application log entries that are forwarded to a syslog server.

3.1.2 Identifying alarm sources in shared-mode deployments

In a shared-mode NSP deployment that includes a system with a similar type of component, for example, a database, a similar alarm may be raised by each system in the event of a fault.

Before you take action to correct such a fault, it is vitally important that you identify the source system that raises the alarm. The Source Type field of the alarm names either the NSP or NFM-P as the source of the alarm.

For example, in a shared-mode NSP and NFM-P deployment, an alarm is raised against the standby database. The Fault Manager Source Type field of the alarm reads “NFM-P”, so an operator determines that the standby NFM-P main database is at fault, rather than the standby NSP PostgreSQL database.

Note: In an NFM-P main database alarm, the Site ID and Site Name fields identify the NFM-P main server that raises the alarm. By contrast, in an NSP database alarm, the Site ID and Site Name fields identify the PostgreSQL database instance.

Regardless of the Source Type, the Additional Text field of a system alarm displays the IP address of the faulted component.

See the NSP Fault Management application help for more information about NSP alarms.

3.2 NSP application access

3.2.1 Application activation

The NSP applications that are required in a typical management network are activated by default and available from the NSP Launchpad.

Applications that are required for more specialized functions are deactivated by default in a new NSP system.

Note: The NSP does not load a deactivated application, or display the application icon on the NSP Launchpad.

However, for a few minutes immediately after an NSP system installation, the icon of an application that is deactivated by default may be displayed on the NSP Launchpad.

Note: An NSP system upgrade preserves the current activation setting of an application.

Note: The JNLP GUI-client installation method is deprecated, and is to be removed in a future release. You can enable or disable the JNLP installation method, as described in 3.3 “How do I configure global NSP application settings?” (p. 34).

The following applications are deactivated by default in a new NSP deployment:

- Subscriber Manager
- Wireless NE Views
- Wireless Supervision

An administrator can reactivate a deactivated application, and also control the loading of other applications. See 3.6 “How do I activate or deactivate NSP applications?” (p. 37) for information.

3.2.2 Browser access to redundant NSP clusters

If you open a browser to the primary NSP cluster URL in a DR deployment, the primary server NSP sign-in page opens.

If you open a browser to the standby NSP URL, the browser is redirected to the primary NSP URL if the standby server is operational; otherwise, the browser shows the standby URL as unreachable.

Single-address DR NSP system access

To reduce the number of IP addresses that an NSP operator requires for access to the servers in a DR NSP deployment, you can use a reverse-proxy server to set one IP address for NSP access, regardless of which NSP cluster is active.

See 4.4 “To enable single-address DR NSP system access” (p. 47) for proxy-server configuration information.

3.2.3 OSS access

Applications that use REST APIs publish a set of URLs for managed application resources or web services. Each domain application documents the URLs that are available to users. The API URLs are accessible through a browser to authorized users, including OSS applications, which can use the URLs for cross-launching.

See the Network Developer Portal for information about OSS access to the NSP using a REST API.

3.2.4 Application help and user documentation

You can open the NSP Help Center from each NSP application user interface by clicking on the ? icon. The Help Center provides application-specific help, as well as access to other NSP documentation.

3.2.5 Application-server connection loss

NSP application sessions that are terminated by a server connection loss may require up to two minutes to reset after the server connection is restored. In the interim, the application GUI may seem to function, but executing a GUI command results in a Server Not Found browser error. The condition persists until an automated system function clears the former application session.

3.2.6 Best practices for application access

Some HTTP errors or stalled user sessions can be avoided by adhering to the following best practices:

- All NSP applications are supported on the latest version of Google Chrome. Although other browser types are supported, Chrome is the preferred browser. For information about additional supported browsers for NSP applications, see the *NSP Planning Guide*.
- It is recommended to use the NSP Launchpad for access to NSP applications, as user-created links to individual applications may be broken by a server activity switch or software upgrade.
- Enable cookies in your browser.
- Sign in to the NSP Launchpad before opening additional NSP applications in other tabs.
- Before signing in as a different user, close all other NSP tabs and sign out of the last tab.
- If multiple NSP applications are open in one browser, close all other NSP tabs before signing out of the last NSP tab; do not just close the browser.
- Avoid pausing a polling application for more than ten minutes.
- In the event of an NSP server activity switch or shutdown, close all browser tabs; you can sign in again when the server returns to service.

3.2.7 Keyboard-based navigation

You can use the keyboard to navigate and interact with most NSP applications. Keyboard navigation allows you to highlight and select interactive elements of the application using keystrokes instead of a pointing device.

The following table lists the accessibility options.

Keystroke	Action
Tab	Advance to next element
Shift + Tab	Return to previous element
Alt + down arrow Option/Alt + down arrow in Apple/OSX	Open pop-up or drop-down menu
Shift + F10 Shift + Fn + F10 in Apple/OSX	Open contextual menu
Ctrl + c Command + c in Apple/OSX	Copy

Keystroke	Action
Ctrl + v Command + v in Apple/OSX	Paste
Enter	Open folder or expandable object such as tile Invoke action on button or menu item
F8 Fn + F8 in Apple/OSX	Move over larger elements or to next page
F5 Shift + Fn + F5 in Apple/OSX	Refresh
Shift + F1 Shift + Fn + F1 in Apple/OSX	Open tool tip
Esc	Close tool tip or menu
Arrow	After tile selected using Tab key, navigate across tiles in matrix such as Fault Management Top Unhealthy NEs view or Service Supervision matrix view Up and down arrows for navigation through items in open contextual or pop-up menu Up and down arrows for navigation between table rows Left and right arrows for navigation across table column headers
Shift + right or left arrow	Reorder data-table columns in selected header

3.3 How do I configure global NSP application settings?

3.3.1 Purpose

Use this procedure to specify the default operating parameters of NSP applications.

3.3.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **System Settings**.

4

If required, configure the Global parameter for applications, Polling time (seconds) and Language.

5

Set or modify the Security statement text, if required.

6

If required, disable JNLP single-user GUI client and client delegate server installation by deselecting the Enable JNLP installation for NFM-P client parameter.

Note: The JNLP installation method is deprecated, and is to be removed in a future release. Installation using the binary client installer, as described in the *NSP NFM-P Deployment and Installation Guide*, is recommended.

7

Enable or disable severity-level background colors for alarm display, if required, using the Row color with severity parameter.

8

If required, configure the Time Zone parameter, which affects alarm messages in Assurance applications.

9

To specify a tile server for map drill-down operations in NSP applications, configure the Map settings parameters:

- Background Map Layer URL—link to a map available under an open license, in the following format:
  https://*tile_server*/*path*/*file*.png
- Background Map Layer Attribution—optional free-form text field for crediting an open license provider for legal purposes

10

Click **SAVE**.

End of steps

3.4 How do I configure alarm-severity colors?

3.4.1 Purpose

Use this procedure to specify the display colors for alarm levels.

3.4.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **Alarm Colors**.

4

Under Select alarm type, select a severity level, and then use one of the following methods to assign a background color to the severity level:

- Click on a color tile in the palette.
- Enter a hexadecimal color code in the text field beside the palette.

Repeat this step to set custom colors for other alarm severities, as required.

5

Click **SAVE**.

End of steps

3.5 How do I configure linked URLs?

3.5.1 Purpose

Use this procedure to link up to 20 external URLs that application users can open in a new browser tab from the More menu on the NSP Launchpad.

3.5.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **Linked URLs**.

4

Click **+Add**. The Add Linked URLs form opens.

5

Configure the Name and URL parameters.

6

Click **Add**. The form closes.

7

Click **SAVE**.

8

To edit a linked URL, click ⋮ **(More), Edit**.

9

To remove a linked URL, click ⋮ **(More), Delete**.

END OF STEPS

3.6 How do I activate or deactivate NSP applications?

3.6.1 Purpose

Use this procedure to specify which NSP applications are available to operators from the NSP Launchpad.

Note: Deactivating and reactivating an NSP application may cause the NFM-P web server to restart unexpectedly. To avoid this behavior, you must restart the NFM-P web server between application deactivation and reactivation.

CAUTION

Service Disruption

This procedure involves a restart of each NSP server.

It is strongly recommended that you perform this procedure only during a scheduled maintenance period.

3.6.2 Steps

1

Sign in to the NSP Launchpad as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **Application Deployment Control**.

4

Select the application category, and then select **INSTALL** or **UNINSTALL** to activate or deactivate applications in the category. Installing an application changes the Inactive status to Active. Uninstalling changes the Active status to Inactive.

Note: If you are reactivating an application, there may be a brief delay before the Launchpad displays the application icon.

5

To deactivate an application, click **UNINSTALL** and click **DISABLE** to confirm. The application is deactivated

End of steps

3.7 How do I configure event logging?

3.7.1 Purpose

Use this procedure to configure the recording of assurance events, or to purge all event records from the database.

Note: Events can be retained for up to 30 days.

3.7.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **Event Logging Policy**.

4

Configure the Enable event logging parameter.

5

To specify how long event records are retained, configure the Retention Time parameter.

6

If required, click **DELETE STORED EVENTS**.

7

Click **SAVE**.

END OF STEPS

3.8 How do I configure an e-mail server for alarm notifications?

3.8.1 Purpose

Use this procedure to configure connection information to an e-mail server. This e-mail feature may be used by NSP to contact NSP users or send alarm notifications as configured in an alarm policy.

3.8.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose **User**, **Settings** from the NSP banner bar.

3

Click **E-mail Notification**.

4

Configure the E-mail Server Settings parameters.

Note: The E-mail server address parameter value must be an IPv4 address or a hostname. You cannot specify a literal IPv6 address; instead, you must use a hostname that is resolvable by DNS.

5

Click **SAVE**.

END OF STEPS

3.9 Syslog record format for NSP application logging

3.9.1 Introduction

The NSP can be configured to forward NSP application log entries, including NFM-P server log entries, to a remote syslog server, as described in the *NSP Deployment and Installation Guide*. The following topics describe the syslog record formats of NSP application and NFM-P server log entries.

NSP application log entries

Different NSP applications log different types and amounts of information. Consequently, an NSP application log entry does not have a fixed length or number of fields, so the following topic describes only the syslog record format, and not each type of application entry.

NFM-P server log entries

When NFM-P server log forwarding is enabled, the standalone or primary NFM-P main server forwards each entry written to the local EmsServer.log file to the specified syslog server.

3.9.2 NSP application syslog record format

Each syslog record for an NSP application log entry has the following fields:

- timestamp
- hostname of syslog message producer
- program name, which is appslogs
- application log entry

Note: NSP application log entries do not have a common format, as the number and type of fields in an entry is application-specific.

The following is an example of a syslog record that contains an NSP application log entry:

```
Nov  9 11:10:28 nsp-1a3 appslogs: {app-specific_log_entry}
```

The fields in the example record have the following values:

- timestamp—Nov 9 11:10:28
- hostname of syslog entry producer—nsp-1a3
- program name—appslogs
- application log entry—*app-specific_log_entry*, which is a comma-separated list of colon-separated attribute-value pairs that contain the log-entry message and other information, for example:

  ```
  "attribute1":"value","attribute2":"value","attribute3":"value"
  ```

3.9.3 NFM-P server syslog record format

Each NFM-P server log entry has the following fields:

- timestamp
- hostname of syslog message producer

- program name, which is nfmpserverlogs
- server log entry

The following is an example of a syslog record that contains an NFM-P server log entry:

```
Nov 7 05:40:15 nfmp-dc_1 nfmpserverlogs:{EMS_server_log_entry}
```

The fields in the example record have the following values:

- timestamp—Nov 7 05:40:15
- hostname of syslog entry producer—nfmp-dc_1
- program name—nfmpserverlogs
- server log entry—*EMS_server_log_entry*, which is a comma-separated list of colon-separated attribute-value pairs that contain the log-entry message and other information, for example:

  ```
  "attribute1":"value","attribute2":"value","attribute3":"value"
  ```

4 NSP global system administration

4.1 Overview

4.1.1 Purpose

This chapter describes NSP system-level administration and management actions.

4.1.2 Contents

NSP system configuration

4.2 Introduction

4.2.1 Description

The following procedures describe global system administration operations that affect NSP user experience or access to NSP applications.

4.3 To add or remove installation options

4.3.1 Purpose

Perform this procedure to add an installation option to an existing NSP system, or to remove an installation option.

CAUTION

Service disruption

Performing the procedure may require a restart of each NSP cluster, which is service-affecting.

You must perform the procedure only during a scheduled maintenance period.

Note: You must perform the procedure on each NSP cluster.

Note: In a DR deployment, you must perform the steps first on the standby NSP cluster.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

4.3.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

If you intend to remove any installation options, you must configure the NSP to preserve the existing deployment.

1. Open the following file using a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

2. Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

```
    deleteOnUndeploy:false
```

3. Save and close the file.
4. Enter the following:

   # /opt/nsp/NSP-CN-*release-ID*/bin/nsp-config.bash --undeploy ↵

3

Open the following file using a plain-text editor such as vi:

/opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

4

Locate the **installationOptions** section, which resembles the following:

```
  installationOptions:
    - name: "NSP Platform - Base Services"
      id: platform-baseServices
    - name: "NSP Platform - Logging and Monitoring"
      id: platform-loggingMonitoring
#    - name: "Other Installation Option"
#      id: otherInstallationOption
```

5

To add an installation option, uncomment the installation option name and id lines by removing the leading # character from each line.

Note: You must preserve the leading spaces in each line.

6

To remove an installation option, convert the installation option name and id lines to comments by inserting a # character at the beginning of each line.

Note: You must preserve the spaces that follow the # character.

7

Save and close the nsp-config.yml file.

8

Enter the following to put the changes into effect:

 # /opt/nsp/NSP-CN-*release-ID*/bin/nsp-config.bash --config --deploy ↵

The installation options are added or removed, as required.

END OF STEPS

4.4 To enable single-address DR NSP system access

4.4.1 Purpose

Use this procedure to reduce the number of IP addresses a user requires for access to the NSP clusters in a DR NSP deployment.

The procedure describes implementing a reverse proxy that presents only one IP address for system access. The reverse proxy maps the IP address to the appropriate NSP cluster.

Note: The procedure describes using the mod_proxy Apache HTTP module. Using a different proxy agent or mod_proxy configuration is supported but not described. Also, mod_proxy installation is not described. Reverse proxy implementation is specific to a network; the network administrator must determine which implementation is best suited to the management network.

4.4.2 Steps

1

Log in as the root user on the station that is to host the reverse proxy.

2

Open a console window.

3

Open the httpd.conf file in the mod_proxy installation directory using a plain-text editor such as vi.

4

Edit the file to include the following:

```
<VirtualHost *:*>
    <Proxy nspOS://dr>
        BalancerMember http://NSP1
        BalancerMember http://NSP2
    </Proxy>
    ProxyPreserveHost Off
    ProxyPass / nspOS://dr/
    ProxyPassReverse / nspOS://dr/
</VirtualHost>
```

where

NSP1 and *NSP2* are the advertised addresses of the NSP clusters

5

Close the console window.

END OF STEPS

4.5 To disable NSP websocket event notifications

4.5.1 Purpose

Websocket-based events are used by some NSP applications. The following steps describe how to disable websocket event notifications, if required.

Note: The websocket connection used by the NSP may not function if a browser or any client is behind a proxy. Websocket communication through an entity between the websocket client and server, for example, a proxy server, firewall, or load balancer, is dependent on the entity configuration.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

4.5.2 Steps

1

Log in as the nsp user on the IP resource control server.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/configure/config ↵
```

4

Open the wsc-security.conf file using a plain-text editor such as vi.

5

Modify the following section to read:

```
websocket {
enableEvents=false
}
```

6

Enter the following to restart the web server:

Note: If the NSP deployment is redundant, you must perform the step on each IP resource control server.

```
# systemctl restart nsp-tomcat ↵
```

The web server restarts, and websocket event notifications are disabled.

7

Close the console window.

END OF STEPS

4.6 To install custom Mistral actions for Workflow Manager

4.6.1 Purpose

Perform this procedure to install custom Mistral actions created by Nokia Professional Services.

Note: Installation of self-developed actions is currently restricted to lab use only.

Before you start, the zip file containing the custom actions must be saved in a directory accessible to the root user. The zip file must follow Python directory packaging rules. The subdirectories in the zip file must be called `actions` and `expressions`.

Note: The Mistral pods must be restarted for the custom actions to be applied in Workflow Manager. The restart will affect Workflow Manager operation.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

4.6.2 Steps

1

Log in as the root user on the station where the NSP container was installed.

2

Execute one of the following:

- To copy the custom actions to the NSP without restarting the pods:

  ```
  # /opt/nsp/NSP-CN-release-ID/tools/wfm/bin/custom-actions.bash
  --install path ↵
  ```

- To copy the custom actions to the NSP and restart the pods immediately:

```
# /opt/nsp/NSP-CN-release-ID/tools/wfm/bin/custom-actions.bash
--install path --restart ↵
```

where *path* is the path to the zip file containing the custom actions you need to install.

3

To restart the Mistral pods, execute the script with the **--restart** option.

END OF STEPS

4.7 To configure a generic mediator for Intent Manager or Workflow Manager

4.7.1 Purpose

Use this procedure to install or uninstall a generic mediator to connect to an external controller. A generic mediator will handle authentication to the controller, allowing a call to be made to the controller from Intent Manager or Workflow Manager without the need to provide credentials.

You can configure as many generic mediators as required, with different external controller parameters. Docker images and helm charts are provided in the Nokia Git repository, along with the charts and images for other mediators.

To provide generic mediator parameters to the NSP you must create a values.yaml file; see the Intent Based Management Framework tutorial on the Network Developer Portal for information.

Naming

Each generic mediator must have a unique name.

Names must be unique in the following areas:

- the name of the instance in helm

  You can allow helm to autogenerate a unique name, however, using the -n option when installing the helm chart will allow you to use a meaningful user identifiable name.

- the name of the pod and Kubernetes configuration

  This is specified in the values.yaml file using the mediator_name value. When the name is given, the Kubernetes structures will be given a name like `nsp-mdt-<name>-mediator`. For example, to name the mediator SF1, enter the mediator_name: "SF1" in the values.yaml file. This will produce a pod and Kubernetes structures prefixed with nsp-mdt-SF1-mediator. If multiple words are in the name, they must be separated with a dash character (-).

The helm instance name and the mediator_name do not have to be the same, however, using the same name may make alignment simpler.

Certs files

Certs files may need to be copied into the pod for requests and authentication to work properly. This is done by a combination of the `copy_certs` and `certsFileName` properties and the `--set-file` flag on the helm command.

When copy_certs is set to true, the NSP will attempt to copy a certs file into the pod in the `/opt/nsp/os/ssl/certs/custom` directory. Since the name of the certs file might be important and the mediator itself will not be aware of this, the file name to give this file is specified in the `certsFileName` value.

For example, if you have these values in the values.yaml file:

- `copy_certs: true`
- `certsFileName: "ca.pem"`

A file named ca.pem is created in the `/opt/nsp/os/ssl/certs/custom` directory.

If you specify `copy_certs: true` in the values.yaml file but do not add the `--set-file` flag to the helm command, the pod will not initialize.

4.7.2 Steps

Perform helm installation

1

Execute the helm installation based on the following example:

```
# helm install helm chart file -n nsp-mdt-name-mediator -f values file
--set-file externalControllerConfig.externalControllerAuth.certsFile=
certs file ↵
```

where

helm chart file is the path to the helm chart

name is the name for the mediator

values file is the path to the values.yaml file

certs file is the path to the certs file

If you do not want to copy certs files to the pod, the **`--set-file`** flag is not required.

Perform helm uninstallation

2

To uninstall a generic mediator, you must delete them using helm. Execute the following:

```
helm delete --purge mediator-instance-name ↵
```

where

mediator instance name is the mediator instance, for example, nsp-mdt-generic-mediator

END OF STEPS

4.8 To configure a Workflow Manager trigger framework

4.8.1 Purpose

Use this procedure to install or uninstall a WFM framework. The WFM Trigger framework is designed to allow you to trigger workflows in the Workflow Manager application from an external system such as Git or HTTPS, without the need for Workflow Manager to monitor the external system directly.

The framework provided by NSP installs a single microservice that can handle at least one user -provided plugin installed into that microservice. The framework handles connections to the internal Kubernetes environment and to the Workflow Manager application, while the plugin connects to the external system. The plugin must include a specified Python file to allow it to push to a Kafka topic in the NSP. From there, a Kafka Trigger can be created in Workflow Manager to trigger a workflow.

You can configure as many microservices as required, with different parameters. Docker images and helm charts are provided in the Nokia Git repository.

This procedure provides general information. See the WFM Trigger Framework tutorial on the Network Developer Portal for more details.

Docker

You will need to use a Docker image generated from the base Docker image provided by NSP, that includes the code written to handle the external system. The name of the base image can be found either by browsing the NSP Docker repository or by looking at the provided defaulted values.yaml file.

Helm charts

A default values.yaml file is provided, and new values.yaml files must be derived from this one to get all the appropriate NSP-specific information that the pod needs to run. You need to change the following values:

- `pod_name`: Each pod must be given a unique pod name to avoid clashes over resources and so to install properly.
- `topic_name`: This specifies the topic name that messages will be sent on (and thus that the Kafka Trigger must listen on). The topic will always be wfm.trigger.<*topic_name*>. If this is not given, the `pod_name` will be used as the topic name.
- `pluginConfigFileName`: The file name of any config file that the user happens to be using. If this is not set, it defaults to `plugin.conf`.
- `pluginConfigFileInstallPath`: The install path of the config file. If this is not set, it defaults to `/opt/nsp/configure/config`.
- `pluginPythonClientFile`: The name and path of the client file that's importing the framework_server_client. If this is not set, it defaults to `/opt/nsp/launch_python_app.py`.
- `pluginCertsFileName`: The file name for any certs that are being used. If this is not set, it defaults to `plugincert.pem`.
- `pluginCertsInstallPath`: The install path for any certs that are being used. If this is not set, it defaults to `/opt/nsp`.

4.8.2 Steps

Perform helm installation

1

Execute the helm installation based on the following example:

helm install *helm name helm path* -f *values file* --set-file pluginCertsFile= *certs file* --set-file pluginConfigFile=*plugin file* ↵

where

helm name is the identifier of the helm chart

helm path is path to the helm chart

values file is the path to the values.yaml file

certs file is the path to and filename of the certs file

plugin file is the path to and filename of the plugin file

If you do not want to copy certs files to the pod, the **--set-file pluginCertsFile** flag is not required.

Perform helm uninstallation

2

To uninstall a framework, you must delete them using helm. Execute the following:

helm delete --purge *helm name* ↵

where

helm name is the framework helm chart identifier.

END OF STEPS

NSP cluster administration

4.9 Introduction

4.9.1 Description

The following procedures describe the following basic Kubernetes administration operations for the NSP:

- starting and stopping clusters
- identifying the master node in a cluster
- displaying the DR status
- performing DR role switches

4.9.2 NSP cluster startup and shutdown

Low-level routine maintenance such as applying a RHEL OS patch to the hosts in an NSP Kubernetes cluster may require that you stop and start Kubernetes in the cluster.

Stopping Kubernetes in an NSP cluster stops the NSP software in the cluster, and creates a network management outage in a standalone deployment.

In a DR deployment, you can avoid a network management outage by stopping and starting the Kubernetes clusters in sequence, as specified in 4.10 “Workflow to stop and start both DR NSP Kubernetes clusters” (p. 54).

4.10 Workflow to stop and start both DR NSP Kubernetes clusters

4.10.1 Description

CAUTION

System Degradation

If the primary NSP cluster and associated primary components outside the cluster are not in the same data center, the maintenance shutdown described in this workflow may cause an undesired DR switchover of one or more components, which can be service-affecting.

Before you attempt to use this workflow, you must ensure that the primary NSP cluster and associated primary components outside the cluster are in the same data center. If not, perform the appropriate procedure to ensure that all component roles in each data center are aligned.

The following is the sequence of high-level actions required to stop and start the active and standby NSP Kubernetes clusters in a graceful manner for maintenance purposes.

See the following procedures for information about stopping and starting a cluster:

- 4.12 “To stop an NSP Kubernetes cluster” (p. 60)
- 4.11 “To start an NSP Kubernetes cluster” (p. 59)

4.10.2 Stages

Perform orderly shutdown of standby components outside NSP cluster

1

If the deployment includes the IPRC or CDRC, stop the standby IPRC and CDRC, as described in 5.4 “To start or stop IP resource control or cross-domain resource control” (p. 78).

2

If the NSP deployment includes NSP analytics servers, stop each analytics server in the standby data center, as described in 5.14 “To start or stop an NSP analytics server” (p. 94).

3

If the NSP deployment includes NSP Flow Collectors:

1. Stop each Flow Collector in the standby data center, as described in 5.19 “To start or stop an NSP Flow Collector” (p. 101).
2. Stop each Flow Collector Controller in the standby data center, as described in 5.31 “To start or stop an NSP Flow Collector Controller” (p. 112).

4

If the NSP deployment includes an auxiliary database, stop the auxiliary database cluster in the standby data center, as described in the *NSP NFM-P Administrator Guide*.

5

If the NSP deployment includes the NFM-P:

1. Stop each Reserved NFM-P auxiliary server of the standby main server, as described in the *NSP NFM-P Administrator Guide*.
2. Stop each Preferred NFM-P auxiliary server of the standby main server, as described in the *NSP NFM-P Administrator Guide*.
3. Stop the standby main server, as described in the *NSP NFM-P Administrator Guide*.
4. Stop the standby main database, as described in the *NSP NFM-P Administrator Guide*.

Perform standby NSP cluster maintenance

6

Stop the standby NSP cluster, as described in 4.12 “To stop an NSP Kubernetes cluster” (p. 60).

7

Perform the required maintenance on the standby cluster.

8

Start the standby cluster, as described in 4.11 “To start an NSP Kubernetes cluster” (p. 59).

Switch NSP cluster roles

9

Perform an activity switch to change the standby cluster role to active, as described in 4.16 “To switch the NSP cluster roles in a DR deployment” (p. 64).

The standby cluster assumes the active role.

Perform orderly startup of standby components outside cluster

10

If the NSP deployment includes the NFM-P:

1. Start the former standby main database, as described in the *NSP NFM-P Administrator Guide*.
2. Start the former standby main server, as described in the *NSP NFM-P Administrator Guide*.
3. Start each Preferred NFM-P auxiliary server of the former standby main server, as described in the *NSP NFM-P Administrator Guide*.
4. Start each Reserved NFM-P auxiliary server of the former standby main server, as described in the *NSP NFM-P Administrator Guide*.

11

If the NSP deployment includes an auxiliary database, start the auxiliary database cluster in the former standby data center, as described in the *NSP NFM-P Administrator Guide*.

12

If the NSP deployment includes NSP Flow Collectors:

1. Start each Flow Collector Controller in the former standby data center, as described in 5.31 “To start or stop an NSP Flow Collector Controller” (p. 112).
2. Start each Flow Collector in the standby data center, as described in 5.19 “To start or stop an NSP Flow Collector” (p. 101).

13

If the NSP deployment includes NSP analytics servers, start each analytics server in the former standby data center, as described in 5.14 “To start or stop an NSP analytics server” (p. 94).

14

If the deployment includes the IPRC or CDRC, start the former standby IPRC and CDRC, as described in 5.4 “To start or stop IP resource control or cross-domain resource control” (p. 78).

Perform orderly shutdown of former primary components outside NSP cluster

15

If the deployment includes the IPRC or CDRC, stop the former primary IPRC and CDRC, as described in 5.4 “To start or stop IP resource control or cross-domain resource control” (p. 78).

16

If the NSP deployment includes the NFM-P:

1. Stop each Preferred NFM-P auxiliary server of the former primary main server, as described in the *NSP NFM-P Administrator Guide*.
2. Stop each Reserved NFM-P auxiliary server of the former primary main server, as described in the *NSP NFM-P Administrator Guide*.
3. Stop the former primary main server, as described in the *NSP NFM-P Administrator Guide*.
4. Stop the former primary main database, as described in the *NSP NFM-P Administrator Guide*.

17

If the NSP deployment includes NSP Flow Collectors:

1. Start each Flow Collector Controller in the standby data center, as described in 5.31 “To start or stop an NSP Flow Collector Controller” (p. 112).
2. Start each Flow Collector in the standby data center, as described in 5.19 “To start or stop an NSP Flow Collector” (p. 101).

18

If the NSP deployment includes NSP analytics servers, stop each analytics server in the standby data center, as described in 5.14 “To start or stop an NSP analytics server” (p. 94).

Perform former primary NSP cluster maintenance

19

Stop the former active cluster, as described in 4.12 “To stop an NSP Kubernetes cluster” (p. 60).

20

Perform the required maintenance on the former active cluster.

21

Start the former active cluster, as described in 4.11 “To start an NSP Kubernetes cluster” (p. 59).

Perform orderly startup of former primary components outside cluster

22

If the NSP deployment includes the NFM-P:

1. Start the former primary main database, as described in the *NSP NFM-P Administrator Guide*.
2. Start the former primary main server, as described in the *NSP NFM-P Administrator Guide*.
3. Start each Preferred NFM-P auxiliary server of the former primary main server, as described in the *NSP NFM-P Administrator Guide*.
4. Start each Reserved NFM-P auxiliary server of the former primary main server, as described in the *NSP NFM-P Administrator Guide*.

23

If the NSP deployment includes an auxiliary database, start the auxiliary database cluster in the former primary data center, as described in the *NSP NFM-P Administrator Guide*.

24

If the NSP deployment includes NSP Flow Collectors:

1. Start each Flow Collector Controller in the former primary data center, as described in 5.31 “To start or stop an NSP Flow Collector Controller” (p. 112).
2. Start each Flow Collector in the primary data center, as described in 5.19 “To start or stop an NSP Flow Collector” (p. 101).

25

If the NSP deployment includes NSP analytics servers, start each analytics server in the former primary data center, as described in 5.14 “To start or stop an NSP analytics server” (p. 94).

26

If the deployment includes the IPRC or CDRC, start the former primary IPRC and CDRC, as described in 5.4 “To start or stop IP resource control or cross-domain resource control” (p. 78).

Restore initial primary/standby NSP cluster roles

27

If required, perform an activity switch to restore the initial active and standby roles, as described in 4.16 “To switch the NSP cluster roles in a DR deployment” (p. 64).

4.11 To start an NSP Kubernetes cluster

4.11.1 Purpose

The following steps describe how to start the Kubernetes software in an NSP cluster.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

4.11.2 Steps

1

Log in as the root user on the NSP configurator VM in the cluster.

2

Open a console window.

3

Enter the following:

```
# nsp-config.bash --deploy ↵
```

The NSP cluster starts.

4

Enter the following periodically to display the Kubernetes cluster status:

```
# kubectl get pods ↵
```

The NSP is operational when the status of each pod is Running.

5

When the NSP cluster is operational, close the console window.

END OF STEPS

4.12 To stop an NSP Kubernetes cluster

4.12.1 Purpose

CAUTION

Network Management Disruption or Outage

Performing the procedure in a standalone deployment completely stops the NSP and creates a network management outage that persists until you start the cluster. In a DR deployment, stopping an NSP cluster may initiate a server activity switch that may temporarily affect network management.

Perform the procedure only during a scheduled maintenance period and under the guidance of technical support.

The following steps describe how to stop the Kubernetes software in an NSP cluster, for example, when the NSP hosts in the cluster require maintenance, or for cluster decommissioning.

Note: If you are stopping the NSP clusters in a DR deployment, ensure that you perform the procedure at the appropriate stage of 4.10 “Workflow to stop and start both DR NSP Kubernetes clusters” (p. 54).

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

4.12.2 Steps

1

Log in as the root user on the NSP configurator VM in the cluster.

2

Open a console window.

3

Open the following file using a plain-text editor such as vi:
/opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

4

Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

```
deleteOnUndeploy:false
```

5

Save and close the file.

6

Enter the following:

```
# nsp-config.bash --undeploy ↵
```

The NSP cluster stops.

7

Enter the following periodically to display the Kubernetes cluster status:

```
# kubectl get pods ↵
```

The NSP cluster is stopped when only the following output is displayed:

```
NAME                    READY   STATUS    RESTARTS   AGE
nsp-backup-storage-0    1/1     Running   0          age
```

8

When the NSP cluster is stopped, close the console window.

END OF STEPS

4.13 To identify the master node in an HA NSP cluster

4.13.1 Purpose

The following steps describe how to list the HA NSP cluster VMs and identify which has the master role.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

4.13.2 Steps

1

Log in as the root user on an NSP cluster VM.

2

Enter the following:

```
# kubectl get nodes -o wide ↵
```

A list of VMs like the following is displayed.

```
NAME    STATUS   ROLES    AGE   VERSION   INTERNAL-IP
EXTERNAL-IP
```

```
node1   Ready    master   nd   version   int_IP   ext_IP
node2   Ready    <none>   nd   version   int_IP   ext_IP
node3   Ready    <none>   nd   version   int_IP   ext_IP
```

3

Log out of the VM.

END OF STEPS

4.14 To display the NSP cluster status

4.14.1 Purpose

The following steps describe how to view the status of standalone or redundant NSP clusters.

Note: You require root user privileges on each NSP cluster VM in each data center.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

4.14.2 Steps

1

Log in as the root user on any NSP cluster VM.

2

For a standalone deployment, enter the following:

```
# kubectl exec -it $(kubectl get pods -l app=nsp-role-manager -o
jsonpath='{.items[0].metadata.name}') --
/opt/nsp/os/rolemgr/bin/rmgrctl status ↵
```

Cluster status output like the following is displayed:

```
Site:   cluster_name
Status: active
Since:  timestamp
```

3

For a DR deployment, enter the following:

```
# kubectl exec -it $(kubectl get pods -l app=nsp-role-manager -o
jsonpath='{.items[0].metadata.name}') --
/opt/nsp/os/rolemgr/bin/rmgrctl statusAll ↵
```

Cluster status output like the following is displayed:

```
Site:   active_cluster_name
Status: active
```

```
Since:  timestamp
Site:   standby_cluster_name
Status: standby
Since:  timestamp
```

4

Close the open console window.

END OF STEPS

4.15 To restart a Kubernetes pod

4.15.1 Purpose

Perform this procedure to restart a Kubernetes pod in an NSP cluster; for example, as directed by technical support during a maintenance operation.

4.15.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Enter the following to list the pod that you need to restart:

```
# kubectl get pods | grep string ↵
```

where *string* is part of the pod name

The pod instance names that include *string* are listed.

4

Enter the following:

```
# kubectl delete pod pod_name ↵
```

where *pod_name* is the name of the pod to restart

The pod is deleted and recreated.

5

Close the console window.

END OF STEPS

4.16 To switch the NSP cluster roles in a DR deployment

4.16.1 Purpose

CAUTION

Service disruption

Performing this procedure causes a temporary loss of network visibility, which may be service-affecting.

You must perform the procedure only with the assistance of technical support during a scheduled maintenance period.

The following steps describe how to switch the active and standby roles of the redundant NSP clusters in a DR deployment.

Note: You require root user privileges on each NSP cluster VM in each data center.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

4.16.2 Steps

1

Log in as the root user on a station in the active NSP cluster.

2

Open a console window.

3

Enter the following:

```
# kubectl exec -it $(kubectl get pods -l app=nsp-role-manager -o
jsonpath='{.items[0].metadata.name}') --
/opt/nsp/os/rolemgr/bin/rmgrctl toStandby ↵
```

The active cluster assumes the standby role.

4

Log in as the root user on a station in the standby NSP cluster.

5

Enter the following:

```
# kubectl exec -it $(kubectl get pods -l app=nsp-role-manager -o
jsonpath='{.items[0].metadata.name}') --
/opt/nsp/os/rolemgr/bin/rmgrctl toActive ↵
```

The standby cluster assumes the active role.

6

Enter the following periodically to display the status of the clusters as the roles change:

```
# kubectl exec -it $(kubectl get pods -l app=nsp-role-manager -o
jsonpath='{.items[0].metadata.name}') --
/opt/nsp/os/rolemgr/bin/rmgrctl statusAll ↵
```

Output like the following is displayed when the role changes are complete:

```
Site:   active_cluster_name
Status: active
Since:  timestamp
Site:   standby_cluster_name
Status: standby
Since:  timestamp
```

7

When the role changes are complete, close the open console windows.

END OF STEPS

NSP cluster lifecycle management

4.17 Introduction

4.17.1 Description

The following procedures describe NSP Kubernetes cluster lifecycle management operations that may occasionally be required, such as:

- moving a pod to a different cluster node
- adding, removing, and replacing cluster nodes

4.18 To move a Kubernetes pod to a different node

4.18.1 Purpose

The following steps describe how to move a non-pinned Kubernetes pod to a different node in an NSP cluster. This action may be required, for example, when you need to allocate additional node capacity to a pod that is pinned to a specific node.

Note: You can only move pods that are not pinned to a specific node using labels.

4.18.2 Steps

CAUTION

System Degradation

The procedure includes operations that fundamentally reconfigure the NSP system.

You must contact Nokia support for guidance before you attempt to perform the procedure.

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Enter the following:

```
# kubectl get pods -o wide ↵
```

The pods are listed.

4

Enter the following to list all nodes:

```
# kubectl get nodes ↵
```

5

Enter the following:

kubectl describe nodes *node* ↵

where *node* is the name of the node that has the pod to move

The node resources are listed.

6

Record the Memory Requests and CPU Requests values.

7

List the nodes to ensure that another node has sufficient capacity for the pod to move:

kubectl describe nodes ↵

The nodes are listed.

8

View the Allocated resources value for each node to ensure that sufficient capacity exists according to the Requests value.

9

Enter the following to cordon the node that the pod is currently running on in order to prevent new pods from being scheduled on the node:

kubectl cordon *node* ↵

where *node* is the name of the node to cordon

Note: Existing pods on the node continue to run and are unaffected.

10

Restart the pod you intend to move.

1. Enter the following:

   # **kubectl get deployments** ↵

   The pods are listed.

2. Enter the following to stop the pod:

   # **kubectl scale deployment *pod* --replicas=0** ↵

   where *pod* is the pod name

3. Enter the following to start the pod:

   # **kubectl scale deployment pod --replicas=1** ↵

The pod is scheduled on a different node.

11

Verify that the pod is successfully moved to another node; enter the following:

```
# kubectl get pods -o wide ↵
```

The pods are listed.

12

Uncordon the node; enter the following:

```
# kubectl uncordon node ↵
```

where *node* is the node name

The node enters the Initializing state, and then moves to the "Running" state.

Note: If the pod remains in the Pending state, no other node may have sufficient resources to host the pod. Use the following command, where *pod* is the pod name, to obtain information about why the pod remains in the Pending state:
`kubectl describe pod`

13

Close the console window.

END OF STEPS

4.19 To add an NSP cluster node

4.19.1 Purpose

The following steps describe how to add nodes to an NSP cluster.

The procedure creates a new worker node in the cluster that can run pods that do not require local storage, such as pods that are not pinned to a specific node.

4.19.2 Steps

CAUTION

System Degradation

The procedure includes operations that fundamentally reconfigure the NSP system.

You must contact Nokia support for guidance before you attempt to perform the procedure.

1

Log in as the root user on the NSP deployer host.

2

Open a console window.

3

Open the following file using a plain-text editor such as vi:

/opt/nsp/kubespray/inventory/nsp-deployer-default/hosts.yml

4

Add a subsection for the new node to the hosts section of the file, for example.

```
nsp-37-2:
  ansible_host: 203.0.113.37
  ip: 203.0.113.37
  access_ip: 203.0.113.37
```

The new node needs to be added to two sections, under the all.hosts section and under the children.kube-node.hosts section.

5

Add the new node name to the children section, kube-node subsection of the file, for example.

```
kube-node:
  hosts:
    nsp-36-1:
    nsp-37-2:
```

6

Save and close the file.

7

Enter the following to deploy the containerization environment:

 # nspdeployerctl install -i *deployer_host_IP* -l *profile* -e *extra_var1* -e *extra_var2* -e *extra_varN* ↵

Note: The action in this step configures the required Kubernetes infrastructure on the new node, but does not affect any running Kubernetes component.

where

deployer_host_IP is the NSP deployer host IP address

profile is the label profile used for the NSP cluster

extra_var1 to *extra_varN* are one or more of the following attribute-value pairs:

- access_ip_is_vip=true
  required if the access address of each NSP cluster member is the IP address of a virtual network interface rather than a physical interface
- enable_dual_stack_networks=true
  required if the cluster VMs support both IPv4 and IPv6 addressing

8

Log in as the root user on the NSP configurator VM.

9

Enter the following to verify that the new node is added to the cluster:

```
# kubectl get nodes ↵
```

An action such as the following causes pod deployment on the new node:

- moving a pod to the node, as described in 4.18 “To move a Kubernetes pod to a different node” (p. 66)
- enabling additional NSP features that do not require local storage, as described in the NSP Deployment and Installation Guide

10

Close the open console windows.

END OF STEPS

4.20 To remove an NSP cluster node

4.20.1 Purpose

The following steps describe how to remove a node from an NSP cluster.

Note: You can use the procedure to remove only a node added using procedure 4.19 “To add an NSP cluster node” (p. 68).

4.20.2 Steps

CAUTION

System Degradation

The procedure includes operations that fundamentally reconfigure the NSP system.

You must contact Nokia support for guidance before you attempt to perform the procedure.

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Enter the following:

```
# kubectl get nodes ↵
```

The NSP cluster nodes are listed.

4

Record the name of the node that you intend to remove.

5

Enter the following to stop all pods that are running on the node:

 # kubectl drain *node* --ignore-daemonsets --delete-local-data ↵

where *node* is the name of the node that you intend to remove

6

Enter the following:

 # kubectl delete node *name* ↵

where *name* is the name of the node to remove

The node is removed from the cluster.

7

Enter the following:

 # kubectl get nodes ↵

The NSP cluster nodes are listed.

8

Verify that the node is removed from the cluster.

9

Log in as the root user on the NSP deployer host.

10

Open the following file using a plain-text editor such as vi:

/opt/nsp/kubespray/inventory/nsp-deployer-default/hosts.yml

11

Remove each reference to the removed node to ensure that the node cannot be added back into the cluster during a redeployment operation.

12

Close the console window.

END OF STEPS

4.21 To restore the NSP deployer host of an NSP cluster

4.21.1 Purpose

The following steps describe how to restore the NSP deployer host VM in an NSP cluster, for example, if the VM fails and must be recreated.

Note: In order to perform the procedure, you require a backup of the NSP deployer host configuration, which is performed during the NSP system installation.

4.21.2 Steps

1

Create a new deployer host VM to manage the NSP cluster. The new deployer host can have the same IP address as the original NSP deployer host, or a new address.

You must use the VM image that created the original NSP deployer host that you wish to replace.

See “To install the NSP” in the *NSP Deployment and Installation Guide* for information about creating the NSP deployer host VM.

2

Log in as the root user on the NSP deployer host.

3

Open a console window.

4

Restore the backup of the /opt/nsp/kubespray/inventory/nsp-deployer-default/hosts.yml file from the original deployer host.

5

You must generate an SSH key for password-free NSP deployer host access to each NSP cluster VM.

Enter the following:

```
# ssh-keygen -N "" -f ~/.ssh/id_rsa -t rsa ↵
```

6

Enter the following for each NSP cluster VM to distribute the SSH key to the VM.

```
# ssh-copy-id -i ~/.ssh/id_rsa.pub root@address ↵
```

where *address* is the NSP cluster VM IP address

7

If you have applied any patches to the existing NSP cluster, apply the same patches to the deployer host VM. This ensures that the new deployer host has any docker images or helm charts in the patch.

8

If the NSP deployer host IP address is not the same as the previous NSP deployer host address, perform the following steps on each NSP cluster VM.

1. Open the /etc/hosts file with a plain-text editor such as vi.
2. Update the following parameters in the file using the new IP address.
   - `repo.nspos.nokia.local`
   - `charts.nspos.nokia.local`
   - `registry.nspos.nokia.local`
3. Save and close the file.

9

Close the console window.

END OF STEPS

4.22 To replace a failed NSP cluster node

4.22.1 Purpose

The following describes how to replace a failed NSP cluster node.

4.22.2 Steps

1

Contact Nokia support for assistance. The procedure varies, depending on which node has failed, and the state of the failed node.

END OF STEPS

5 NSP component administration

5.1 Overview

5.1.1 Purpose

This chapter describes the administration and management actions for individual NSP components.

5.1.2 Contents

IP resource control / cross-domain resource control administration

5.2 Introduction

5.2.1 Description

The following procedures describe IP resource control and cross-domain resource control management functions such as:

- displaying the status
- starting and stopping
- applying a licence

5.3 To display the status of IP resource control or cross-domain resource control

5.3.1 Purpose

The following steps describe how to view the operational status of an IP resource control or cross-domain resource control instance.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

5.3.2 Steps

1

Log in as the root user on the IP resource control or cross-domain resource control server.

2

Open a console window.

3

Perform one of the following.

a. To display the local instance status, enter the following:

```
# nspdctl status ↵
```

b. To display the status of the peer instance in a DR deployment, enter the following:

```
# nspdctl --host server status ↵
```

where *server* is the instance IP address or hostname

The status is displayed.

The instance is operational if the State value is “running” and the required services are shown as “active”.

Note: The nsp-sdn-replication service is shown as active only in a DR deployment.

4

Close the console window.

END OF STEPS

5.4 To start or stop IP resource control or cross-domain resource control

5.4.1 Purpose

The following steps describe how to start or stop an IP resource control or cross-domain resource control instance.

CAUTION

Service disruption

Stopping IP resource control or cross-domain resource control may create a network-management outage; also, starting the functions out of sequence in a DR deployment may initiate a server activity switch that is disruptive to network management.

Perform the procedure only under the guidance of technical support during a scheduled maintenance period.

Note: A leading # character in a command line represents the root user prompt, and is not to be included in a typed command.

5.4.2 Steps

1

Log in as the root user on the IP resource control or cross-domain resource control server.

2

Open a console window.

3

To start the server, enter the following:

```
# nspdctl --host server start ↵
```

where *server* is the server IP address or hostname

The server starts.

4

To stop the server, enter the following:

```
# nspdctl --host server stop ↵
```

where *server* is the server IP address or hostname

The server stops.

5

Close the console window.

END OF STEPS

5.5 To apply an NSP license to IP resource control

5.5.1 Purpose

CAUTION

Service disruption

Performing the procedure may cause a temporary IP resource control outage, which is service-affecting.

If the new license includes a change to the system specifications, an IP resource control restart is required; a simple license renewal with no configuration change does not require a restart, so is not service-affecting.

You must perform the procedure only under the guidance of technical support during a scheduled maintenance period.

Perform the procedure to apply a new or updated NSP license to the IP resource control in a data center.

Note: You must perform the steps in each data center of a DR deployment, and first in the standby data center.

Note: In a 1+1 DR NSP deployment, restarting the nsp-tomcat service in Step 10 may trigger an IP resource control switchover to the standby.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

5.5.2 Steps

Import license

1

Log in as the nsp user on the station used for the initial IP resource control deployment.

2

A change to the license specifications, for example, the addition of DR or another function, may require an equivalent update to the IP resource control configuration; if you are unsure whether your license specifications affect the IP resource control configuration, contact technical support.

Note: No configuration update is required for the renewal of an existing license that has unchanged specifications.

If the license includes one or more changed specifications, update the IP resource control configuration.

1. Open the following file using a plain-text editor such as vi:

   NSP_installer_directory/config/config.yml

   where *NSP_installer_directory* is the directory contains the NSP software bundle
2. As required, modify the parameters to match the updated license specifications.
3. Save and close the file.

3

Copy the new license file to the *NSP_installer_directory*/license directory.

4

If you are renewing an existing license with unchanged specifications, no further action is required; go to Step 11.

Restart IP resource control

5

Wait two minutes for IP resource control to import the new license.

6

Open a console window.

7

Enter the following:

```
bash$ cd NSP_installer_directory/bin ↵
```

8

Enter the following to apply the license:

```
bash$ ./install.sh ↵
```

The license is distributed to each local IP resource control server.

9

Enter the following to switch to the root user:

```
bash$ su ↵
```

10

Enter the following to restart the web server and activate the license:

```
# systemctl restart nsp-tomcat ↵
```

The web server restarts, and the license is applied.

11

Close the console window.

END OF STEPS

5.6 To enable additional IP resource control functions

5.6.1 Purpose

Use this procedure to enable IP resource control functions that are disabled by default.

Note: You must perform the procedure on each IP resource control server in the NSP system.

Note: The following RHEL CLI prompt in a command line denotes the nsp user, and is not to be included in a typed command:

- bash$

5.6.2 Steps

Note: You must edit a file in the procedure using only a plain-text editor such as vi.

1

Log in as the nsp user on the IP resource control server.

2

Enter the following to stop the server:

```
bash$ nspdctl --host server stop ↵
```

where *server* is the IP resource control server IP address

3

To enable BGP-LS topology learning; edit the /opt/nsp/configure/config/arm-system.conf file to read as follows:

```
nrcp {
  bgpLs
   {
      isTopoSourceBgpLS=true
```

4

If the NSP deployment includes redundant VSR-NRCs, edit the /opt/nsp/configure/config/sros-vms.conf file to configure one virtual ID for the redundant VMs:

```
sros-vms {
  enabled=false
  vms =[
       {
         .
         .
         .
         v_id=virtual_ID
```

where *virtual_ID* is a positive integer

5

To enable PCEP for PCC- and PCE-initiated LSP creation; edit the /opt/nsp/configure/config/sros-vms.conf file to read as follows:

```
sros-vms {
  enabled=false
  vms =[
       {
         .
         .
         .
         pcep=true
```

6

To enable OpenFlow for flow steering; edit the /opt/nsp/configure/config/sros-vms.conf file to read as follows:

```
sros-vms {
  enabled=false
  vms =[
       {
         .
         .
         .
         openflow=true
```

7

Enter the following to start the IP resource control server:

```
bash$ nspdctl --host server start ↵
```

where *server* is the server IP address

The server starts.

8

Close the console window.

END OF STEPS

MDM Administration

5.7 Introduction

5.7.1 Description

The following workflow and procedures describe MDM administration operations.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

5.8 Workflow to commission a device for model-driven management

5.8.1 Description

The following sequence of high-level actions describes how to prepare a device for MDM management.

Note: In order for the NSP to manage an NE using MDM, the NE must not currently be managed by an NFM-P system in the NSP deployment.

5.8.2 Stages

1

If the NFM-P currently manages the NE, unmanage the NE from the NFM-P, as described in the “Device discovery” chapter of the *NSP NFM-P User Guide*.

2

Configure the following on the device:

- device identification—NE name used for NSP filtering, configuration and monitoring
- management interface protocol configuration—authentication and communication parameters for device management interface

See the device and adaptor documentation for information.

3

Use the NSP Device Administrator application to discover the device and to verify the device management.

See the Device Administrator application help for information about MDM device discovery.

5.9 To restart an MDM server

5.9.1 Purpose

Perform the following steps to restart an MDM server in an NSP cluster.

5.9.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Enter the following to display the MDM server instances:

```
# kubectl get pods | grep mdm-server ↵
```

The MDM server instances and pod numbers are listed.

4

Enter the following:

```
# kubectl delete pod mdm-server-n ↵
```

where *n* is the MDM server pod number

The MDM server restarts.

5

Repeat Step 4 to restart an additional MDM server, as required.

6

Close the console window.

END OF STEPS

5.10 To install or upgrade MDM adaptors

5.10.1 Purpose

Perform this procedure to install or upgrade MDM adaptors to enable MDM NE discovery and management.

NSP adaptors are available for evaluation in a lab environment, and as a starting point for development. The adaptors are validated only against specific network devices and configurations; the readme file in an adaptor package lists the supported devices and configurations.

The following must be true before you attempt to perform the procedure:

- The NSP system includes MDM, and is installed.
- The NSP clusters, MDM servers, mdm-tomcat service, and file service are initialized and operational.

Note: Because your requirements or configuration may differ from the tested configurations, it is strongly recommended not to use the adaptors in a production network.

Contact Nokia to obtain an adaptor suite zip file, and for assistance with adaptor customization.

Note: It is not necessary to uninstall any adaptors before installing a newer adaptor package.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

5.10.2 Steps

1

Log in as the root user on the NSP configurator VM in the standalone or primary NSP cluster.

Note: In a DR deployment, you perform the installation or upgrade only on the primary NSP cluster, which then replicates the adaptor configuration to the standby cluster.

2

Transfer the MDM adaptor suite zip file to an empty temporary directory.

3

Enter the following to verify that the REST token for the NSP cluster is valid:

```
# /opt/nsp/NSP-CN-release-ID/tools/mdm/bin/adaptor-suite.bash --user
admin --pass password --list ↵
```

Messages like the following are displayed if the token is valid.

```
Using log file: log_file
INFO: timestamp -> Listing MDM adaptors suites ...
INFO: timestamp -> Token requested...
INFO: timestamp -> Token acquired
INFO: timestamp -> Listing adaptor suites...
Installed adaptor suites:
```

list_of_adaptors

INFO: *timestamp* -> Done.

4

Navigate to the directory that contains the adaptor suite zip file.

5

Enter the following:

Note: Multiple files and wildcards are supported. For example, **install sros*** installs all adaptor suites that have filenames beginning with sros.

/opt/nsp/NSP-CN-*release-ID*/tools/mdm/bin/adaptor-suite.bash --install *zip_file* ↵

where

zip_file is the name of the adaptor zip file

The script prompts you for administrator credentials.

6

Enter the NSP admin user credentials.

The adaptors are transferred to each MDM instance.

Note: It may take 30 minutes or more for all adaptors to load.

After the adaptors load, you can use the NSP Device Administrator application to discover compatible devices. See the application online help for information about device discovery.

7

If the NSP system is a DR deployment, verify that the primary and standby file-service-app pods are communicating.

1. Log in as the root user on the NSP configurator VM in the standby NSP cluster.
2. Enter the following:

   # **kubectl exec -it nsp-file-service-app-0 -- ls /opt/nsp/containers/nspvolume/fileservice/nokia/nsp/mdm/features/suite/** ↵

   Messages like the following are displayed, and the MDM adaptor zip files are listed.

   Defaulting container name to nsp-file-service-app

   Use 'kubectl describe pod/nsp-file-service-app-0 -n default' to see all of the containers in this pod.

   If the adaptor zip files are listed, the primary and standby pods are communicating; the adaptors are installed on the standby cluster upon DR activation of the standby cluster.

END OF STEPS

5.11 To enable TLS for MDM telemetry and gNMI on_change support

5.11.1 Purpose

To enable TLS communication between MDM and managed NEs after an NSP system deployment, you must deploy a self-signed TLS certificate to each MDM-managed device that supports gRPC TLS, and import the certificate to each MDM truststore.

The following steps describe how to secure the following NSP communication with NEs by importing an NE TLS certificate:

- MDM telemetry
- gNMI on_change notifications

Note: A gRPC certificate is separate from a certificate used for secure communication within the NSP system.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

5.11.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Transfer the NE certificate files to the /opt/nsp/NSP-CN-*release-ID*/tls/telemetry directory.

Note: You must not modify or delete any existing file in the directory.

3

Enter the following to delete the nsp-tls Kubernetes secret:

```
# kubectl delete secret nsp-tls ↵
```

4

Enter the following:

```
# /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --config ↵
```

5

Enter the following:

```
# /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --deploy ↵
```

The certificate file is made available to each MDM instance, but not yet imported.

Import certificate to MDM servers

6

Perform one of the following to import the TLS certificate to the TLS truststore on each MDM server.

a. Manually import the certificate; perform the following steps for each MDM server to activate the gRPC certificate file.

Note: A manual import is not service-affecting, and is the recommended option.

1. Enter the following to copy the certificate file to the MDM server:

   # **kubectl cp /opt/nsp/NSP-CN-*release-ID*/tls/telemetry/*gRPC_certificate_file* mdm-server-*n*:/opt/nsp/os/ssl/certs/telemetry/** ↵

   where

   n is the mdm-server pod number

   gRPC_certificate_file is the gRPC certificate file

2. Enter the following:

   # **kubectl exec -it mdm-server-*n* -- /opt/nsp/os/jre/bin/keytool -alias *gRPC_alias* -file /opt/nsp/os/ssl/certs/telemetry/*gRPC_certificate_file* -import -keystore /opt/nsp/os/ssl/nsp.truststore -storepass *password*** ↵

   where

   n is the mdm-server pod number

   gRPC_alias is the TLS keystore alias for the unique gRPC certificate

   gRPC_certificate_file is the gRPC certificate file

   password is the TLS keystore password

   You are prompted to import the certificate.

3. Enter yes ↵.

   The MDM server imports the certificate to the local TLS truststore.

b. Restart the MDM server pod; perform 5.9 “To restart an MDM server” (p. 85) for each MDM server pod.

Note: Restarting an MDM server pod is service-affecting, and must be performed only during a scheduled maintenance period.

7

Close the open console windows.

End of steps

5.12 To manage MDM mappings and model definitions

5.12.1 Purpose

The following steps describe how to manage the YANG model definitions, NE model definitions, and JSON device mappings in an NSP cluster.

Note: The NSP system must be fully operational when you perform the procedure.

Note: If you are adding device mappings or model definitions:

- The new device -mapping or model-definition files must be in a directory accessible to the NSP configurator.
- If a zipped collection of files is being installed, all files in the collection must be of the same type, for example, telemetry device mappings.

Note: If you are adding an NE model, the required files are in a compressed archive file; contact Nokia for access to the file.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

5.12.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Enter the following:

```
# cd /opt/nsp/NSP-CN-release-ID/tools/mdm/bin ↵
```

YANG models

3

To list the YANG module sets, or the YANG models in a module set, enter the following:

```
# ./yang-files.bash --list --modulesetname name --save directory
--user user --pass password ↵
```

Note: The --modulesetname option is optional; if included, the model names in the specified module set are listed; otherwise, the module set names are listed.

where

name is a module set name defined in the NSP

directory is an existing local directory in which to save the list . The --model option is optional and allows you to specify model names to save.

user and *password* are the credentials of an NSP administrative user

4

To add one or more YANG model definitions, enter the following:

Note: Because of the heavy processing load associated with importing model definitions, it is strongly recommended that you do not attempt to install all YANG files in a bundle at once.
If you need to install a large number of YANG files, the recommended method is to install only the immediately required files, in small batches, and install any additional YANG files later, as required.
If the following message is displayed after you try to install YANG files, you may need to install fewer files at a time; you can use the command option in the message to verify that the current operation installed all specified files:
`WARN: Timeout uploading yang files. Please verify with --list --modulesetname name`

```
# ./yang-files.bash --add path --modulesetname name --user user --pass password ↵
```

where

path is the path to a YANG definition file, or to a directory that contains definition files

name is the name of a module set in the NSP

user and *password* are the credentials of an NSP administrative user

5

To remove one or more YANG model definitions, enter the following:

```
# ./yang-files.bash --remove definition_1 definition_2 ...definition_n --modulesetname name --user user --pass password ↵
```

where

definition_1 definition_2 ...definition_n is a list of model definitions to remove; a definition has the following format: modelname,namespace,revision—for example, old-definition,urn:nokia.com:nsp:telemetry:mymodel,2021-03-14

name is the name of a module set in the NSP

user and *password* are the credentials of an NSP administrative user

NE models

6

To add one or more NE model definitions, enter the following:

```
# ./ne-model.bash --install path filename --user user --pass password ↵
```

where

path is the path to the NE model zip file

filename is the filename of the NE model zip file

user and *password* are the credentials of an NSP administrative user

JSON device mappings

7

To list the current JSON device mappings, enter the following:

Note: The --ne option applies only to resync mappings, and is optional.

Note: The --class option applies only to nfmp-resync mappings, and is optional.

Note: If the --ne and --class options are omitted, all mappings of the specified type are listed.

Note: The --save option saves the list to a file, and is valid only for telemetry mappings.

```
# ./json-files.bash --list --mapping mapping_type --ne device_def
--class class_path --save path --user user --pass password ↵
```

where

mapping_type is one of the following mapping types: telemetry, resync, or nfmp-resync

device_def is a device type definition, for example, SR-7750,21.10.R1

class_path is the NFM-P class path of the device type, in XPath format

path is an existing local directory in which to save the mapping list

user and *password* are the credentials of an NSP administrative user

8

To add one or more JSON mapping files, enter the following:

```
# ./json-files.bash --add path --user user --pass password ↵
```

where

path is the path to a JSON mapping file, or to a directory that contains mapping files

user and *password* are the credentials of an NSP administrative user

9

Enter the following to remove one or more JSON device mappings:

Note: The --filename option is required only for telemetry mappings.

Note: The --ne option applies only to resync mappings, and is optional.

Note: The --class option applies only to nfmp-resync mappings, and is optional.

Note: If the --ne and --class options are omitted, all mappings of the specified type are removed.

```
# ./json-files.bash --remove --mapping mapping_type --filename
mapping_1 mapping_2 ...mapping_n --ne device_type --class class_path
--user user --pass password ↵
```

where

mapping_type is one of the following mapping types: telemetry, resync, or nfmp-resync

mapping_1 mapping_2 ...mapping_n is a list of telemetry mapping files to remove

device_type is a device type definition, for example, SR-7750,21.10.R1; a value has the following format: neType,neVersion—for example, SR-7750,21.10.R1

class_path is the NFM-P class path of the device type, in XPath format

user and *password* are the credentials of an NSP administrative user

10

Close the console window.

END OF STEPS

NSP analytics server administration

5.13 Introduction

5.13.1 Description

The following procedures describe NSP analytics server administration operations.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

5.14 To start or stop an NSP analytics server

5.14.1 Purpose

The following steps describe how to start or stop the NSP analytics server software on a station.

5.14.2 Steps

1

Log in as the nsp user on the analytics server station.

2

Open a console window.

3

To start the NSP analytics server, enter the following:

```
bash$ /opt/nsp/analytics/bin/AnalyticsAdmin.sh start ↵
```

The following is displayed:

```
Starting Analytics Application
```

When the analytics server is started, the following is displayed.

```
Analytics Application successfully started!
```

4

To stop the NSP analytics server, enter the following:

```
bash$ /opt/nsp/analytics/bin/AnalyticsAdmin.sh stop ↵
```

The following is displayed:

```
Stopping Analytics Application
```

When the analytics server is stopped, the following is displayed:

```
Analytics Application is not running
```

5

Close the console window.

End of steps

5.15 To manage images on an analytics server

5.15.1 Purpose

The following steps describe how to upload logo images from an analytics server to the Images folder in the NSP Analytics application repository, or to update or remove existing images. You can use logo images for Analytics report branding.

Before you begin, the images must be saved to the analytics server in one of the following formats:

- JPEG
- JPG
- GIF
- PNG
- SVG
- BMP

An image name or filename can include only the following characters:

- alphanumerics
- underscore (_)
- period (.)

Note: You can also manage images from the NSP Analytics application; see the application online help for information.

5.15.2 Steps

1

Log in as the nsp user on the NSP analytics server station.

2

Open a console window.

3

Create a text file that contains the information for each image that you want to deploy to the Analytics application; add one line for each image, in the following format:

image_name|*path*/*image_filename*

where

image_name is the name to assign to the Resource ID that a user must specify when adding the image to a report

path is the absolute path of the image file

image_filename is the name of the image file, and is the name that the Analytics application applies to the image in the Repository folder

4

To remove an image from the application Repository folder, add the following line to the text file:

```
image_filename|delete
```

5

Enter the following to deploy the images:

```
bash$ /opt/nsp/analytics/bin/AnalyticsAdmin.sh deployImage text_file ↵
```

where *text_file* is the absolute path of the text file created in Step 3

The analytics server deploys the images and displays progress messages.

6

When the image deployment is complete, close the console window.

END OF STEPS

5.16 To enable and manage analytics server logging

5.16.1 Purpose

The following steps describe how to enable, configure, or disable the logging of Analytics application events on an NSP analytics server, for example, when troubleshooting an application problem.

By default, an analytics server logs only error events.

CAUTION

System disruption

Performing the procedure restarts the analytics server.

Also, the logging is verbose; the created log files may consume excessive disk space if logging is enabled for an extended period.

Perform the procedure only if required, and only for the period required to collect the log entries of interest. Contact technical support for assistance or more information.

Note: The following RHEL CLI prompt in a command line denotes the nsp user, and is not to be included in a typed command:

- bash$

Note: If the analytics servers are redundant, you must perform the procedure on each analytics server to ensure that all log events are collected, for example, in the event that the Analytics application begins using a different analytics server, or if analytics load balancing is enabled.

5.16.2 Steps

1

Log in as the nsp user on the NSP analytics server station.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/analytics/bin ↵
```

4

Enter the following:

```
bash$ ./AnalyticsAdmin.sh enableLog object ↵
```

where *object* is one of the following:

- ADHOC—generates logs during ad hoc report design
- SQL—logs the SQL commands that the analytics server generates during report execution
- ALL—enables all available log objects

The following message and prompt are displayed.

```
This Action requires Analytics Server restart.
Please type 'YES' to continue.
```

5

Enter YES.

The following is displayed as the analytics server restarts and the logging begins.

```
Stopping Analytics Application
date time Starting Analytics Application
Waiting for Analytics Server to come up
date time Analytics Server is UP and Running
Analytics Server successfully started!
```

The log entries are stored in the following file:

- /opt/nsp/analytics/log/analytics.server.log

6

To change the logging level, perform the following steps:

Note: You must use the resetLog option to disable any logging level that is enabled. For example, if ALL logging is enabled, and you want only SQL logging, you must disable ALL logging, and then enable SQL logging; using the enableLog option does not disable any previously enabled logging level.

1. Reset the logging level, as described in Step 7.
2. Go to Step 4.

7

To reset the logging function to the default of logging only error events, perform the following steps.

1. Enter the following:

   ```
   bash$ AnalyticsAdmin.sh resetLog ↵
   ```

   The following message and prompt are displayed.

   ```
   This Action requires Analytics Server restart.
   Please type 'YES' to continue.
   ```

2. Enter YES.

   The following is displayed as the analytics server restarts and the logging is reset to the default level.

   ```
   Stopping Analytics Application
   date time Starting Analytics Application
   Waiting for Analytics Server to come up
   date time Analytics Server is UP and Running
   Analytics Server successfully started!
   ```

8

Close the console window.

END OF STEPS

5.17 To collect analytics-server log files

5.17.1 Purpose

Use this procedure to collect the relevant log files for troubleshooting an NSP analytics server if requested by technical support.

5.17.2 Steps

1

Log in as the root user on the NSP analytics server station.

2

Open a console window.

3

Enter the following:

```
# cd /opt/nsp/analytics/bin ↵
```

4

Enter the following:

Note: You cannot specify /tmp, or any directory below /tmp, as the output directory.

```
# ./getDebugFilesAnalytics.bash output_dir days ↵
```

where

output_dir is a local directory that is to contain the output files

days is the optional number of days for which to collect log files; if not specified, all logs are collected

Messages like the following are displayed as the logs are collected:

```
Please wait, capturing workstation information files. This may take a
few minutes...
Done capturing workstation information files.
Please wait, capturing analytics server debug files. This may take
several minutes...
Done capturing analytics server debug files.
Please wait, capturing nsp os log files. This may take several
minutes...
Done capturing nsp os log files.
----------------------------------------
Please ftp the output_dir/filespec.tar files to the Nokia ftp server
ftp to IP_address, login as anonymous
Put the files in /pub/<ER_NAME>/incoming
Contact your Nokia support representative for assistance
----------------------------------------
```

5

Transfer the files as directed in the script output.

6

Close the console window.

END OF STEPS

NSP Flow Collector administration

5.18 Introduction

5.18.1 Description

The following procedures describe NSP Flow Collector administration operations.

Note: Some operations require a familiarity with flow statistics and the associated NSP collection, transfer, and storage mechanisms. See the “Flow statistics collection” chapter in the *NSP NFM-P Statistics Management Guide* for information about the flow statistics collection and processing options.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

5.19 To start or stop an NSP Flow Collector

5.19.1 Purpose

The following steps describe how to start or stop an NSP Flow Collector that is on a dedicated station, or collocated with an NSP Flow Collector Controller.

CAUTION

System degradation

On a station that hosts a collocated NSP Flow Collector and NSP Flow Collector Controller, starting or stopping the Flow Collector also starts or stops the Flow Collector Controller, and affects all Flow Collectors associated with the Controller.

Before you stop an NSP Flow Collector that is collocated with a Flow Collector Controller, ensure that you understand the implications of the action.

5.19.2 Steps

1

Log in as the nsp user on the NSP Flow Collector station.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/flow ↲
```

4

To stop the NSP Flow Collector, perform one of the following:

a. If the NSP Flow Collector is collocated with an NSP Flow Collector Controller, enter the following:

```
bash$ ./fcc/bin/flowCollectorController.bash stop ↲
```

The NSP Flow Collector and NSP Flow Collector Controller stop.

b. If the NSP Flow Collector is on a dedicated station, enter the following:

```
bash$ ./fc/bin/flowCollector.bash stop ↲
```

The NSP Flow Collector stops.

5

To start the NSP Flow Collector, perform one of the following:

a. If the NSP Flow Collector is collocated with an NSP Flow Collector Controller, enter the following:

```
bash$ ./fcc/bin/flowCollectorController.bash start ↲
```

The NSP Flow Collector and NSP Flow Collector Controller start.

b. If the NSP Flow Collector is on a dedicated station, enter the following:

```
bash$ ./fc/bin/flowCollector.bash start ↲
```

The NSP Flow Collector starts.

6

Close the console window.

END OF STEPS

5.20 To display the NSP Flow Collector status or release level

5.20.1 Purpose

The following steps describe how to view the operational status or software release of an NSP Flow Collector that is on a dedicated station, or collocated with an NSP Flow Collector Controller.

5.20.2 Steps

1

Log in as the nsp user on the NSP Flow Collector station.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/flow ↵
```

4

To display the NSP Flow Collector status, perform one of the following:

a. If the NSP Flow Collector is collocated with an NSP Flow Collector Controller, enter the following:

```
bash$ ./fcc/bin/flowCollectorController.bash status ↵
```

The NSP Flow Collector Controller and Flow Collector status information is displayed.

b. If the NSP Flow Collector is on a dedicated station, enter the following:

```
bash$ ./fc/bin/flowCollector.bash status ↵
```

The NSP Flow Collector status is displayed.

If the NSP Flow Collector is running, the line that begins with flow-collector includes Started.

If the NSP Flow Collector is not running, the following is displayed:

```
nspos-karaf.service is not running
```

5

To display the NSP Flow Collector release level, perform one of the following.

a. If the NSP Flow Collector is collocated with an NSP Flow Collector Controller, enter the following:

```
bash$ ./fcc/bin/flowCollectorController.bash version ↵
```

The NSP Flow Collector Controller and Flow Collector release level is displayed.

b. If the NSP Flow Collector is on a dedicated station, enter the following:

```
bash$ ./fc/bin/flowCollector.bash version ↵
```

The NSP Flow Collector release level is displayed.

6

Close the console window.

END OF STEPS

5.21 Workflow to configure flow statistics collection

5.21.1 Stages

1

Open the NSP Flow Collector web UI; see 5.22 “To open the NSP Flow Collector web UI” (p. 104) .

2

Specify the NEs from which to collect flow statistics; see 5.23 “To specify the NEs and MDAs for flow statistics collection” (p. 105).

3

Specify whether the statistics in a domain are to be stored for NSP Analytics use, or forwarded to a file server; see 5.24 “To configure the AA flow data persistence” (p. 105).

4

Configure the aggregation parameters; see 5.25 “To configure flow statistics aggregation” (p. 106).

5

If required, configure the transfer of collected flow statistics to a remote server; see 5.26 “To configure the transfer of result files” (p. 107).

6

If required, specify the transfer options for statistics files in CSV format; see 5.27 “To configure CSV file compression and renaming” (p. 108).

7

Configure a special-study policy, if required; see 5.28 “To configure an AA Cflowd special-study policy” (p. 109).

8

Configure AA application or protocol filters as required; see 5.29 “To configure an AA application or protocol filter” (p. 110).

5.22 To open the NSP Flow Collector web UI

5.22.1 Steps

1

Use a browser to open the following URL:

https://*server*:8443/fc/admin

where *server* is the NSP Flow Collector IP address or hostname

2

If a login form opens, enter the required user credentials and click OK. The NSP Flow Collector page opens.

END OF STEPS

5.23 To specify the NEs and MDAs for flow statistics collection

5.23.1 Steps

1

Open the NSP Flow Collector web UI. The Collection Policy page is displayed.

2

Add NEs and IS-AA MDAs, as required.

1. Click Add. A new table row is displayed.
2. Configure the following parameters:
   - System ID

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

     You can specify multiple MDAs on one NE by adding one table row for each MDA and using the same System ID in each row.
   - Source IPFIX Address

     The Source IPFIX Address value is the NE address specified in the discovery rule for the NE. This address will be used to send IPFIX traffic to the NSP Flow Collector.
   - Description
3. If the NSP Flow Collector is to collect system Cflowd statistics, configure the following parameters:
   - Port
   - Flow Protocol

Note: For system Cflowd statistics, the combination of System ID, Source IPFIX Address, and Port must be unique.

3

To delete an NE, select the Delete on save check box beside the NE.

4

Click Save Configuration. The configuration is saved.

END OF STEPS

5.24 To configure the AA flow data persistence

5.24.1 Purpose

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

5.24.2 Steps

1

Open the NSP Flow Collector web UI. The Collection Policy page is displayed.

2

Click on the Results Persistence tab. The AA Cflowd statistics domains are listed.

3

Select or deselect the persistence options for each statistics domain, as required.

- Select IPDR to enable data encoding in XDR format.
- Select Aux DB to enable data storage for use by the NSP Analytics application.
- Select CSV to enable data encoding in CSV format.

4

Click Save Configuration. The persistence settings are applied.

5

Save your changes and close the forms.

End of steps

5.25 To configure flow statistics aggregation

5.25.1 Steps

1

Open the NSP Flow Collector web UI. The Collection Policy page is displayed.

2

Click on the Aggregation Policy tab.

3

Perform one of the following:

a. If the NSP Flow Collector is to collect system Cflowd statistics, select the required aggregation types from the tabs in the lower panel.

b. If the NSP Flow Collector is to collect AA statistics, select one or more statistics classes in the Subscriber Collection panel to enable aggregation for the classes.

4

Use the Interval drop-down menus in the Aggregation Intervals panel to specify the aggregation interval for each statistic type, as required.

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

Note: For BB NAT statistics, you must set the collection interval no higher than the following, based on the expected flow rate:

- 350 000 flows/s—1 minute
- 50 000 flows/s—5 minutes
- 25 000 flows/s—15 minutes

5

The Interval Closing Timeout parameter specifies a latency value that is applied at the end of a collection interval to ensure that any queued statistics are written to the current file. Typically, the default value of one second is adequate; configure the parameter only at the request of technical support.

6

Specify the aggregations for each statistic type, as required.

1. Click on the tab in the lower panel that corresponds to the statistic type.
2. Select or deselect aggregations, as required.

7

Click Save Configuration. The configuration is saved.

End of steps

5.26 To configure the transfer of result files

5.26.1 Purpose

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

- NAT
- IPDR
- PGW-EDR
- CSV

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

Note: SFTP transfers are considerably slower than FTP transfers.

5.26.2 Steps

1

Click on one of the following tabs:

- NAT Transfer
- IPDR Transfer
- PGW-EDR Transfer
- CSV Transfer

2

Configure the parameters:

- Enable Transfer—whether file transfers are enabled
- Transfer Protocol—FTP or SFTP
- IP Address / Host name—file server address
- Port—file server port
- Location—file server directory that is to contain the files
- User—FTP or SFTP username
- Password—FTP or SFTP password

3

Click Save Configuration. The configuration is saved.

End of steps

5.27 To configure CSV file compression and renaming

5.27.1 Purpose

Perform this procedure to enable the following for CSV-formatted statistics output files:

- timestamp inclusion in file name
- gzip compression

5.27.2 Steps

1

Log in to the NSP Flow Collector station as the nsp user.

2

Open the following file using a plain-text editor such as vi:

/opt/nsp/flow/fc/cfg/CfdStorage.properties

3

Locate the csv file section.

4

To enable gzip file compression of the output files:

1. Locate the section that begins with the following:

   ```
   # compression option for CSV files
   ```

2. Set the following parameter to true, as shown below:

   ```
   com.alu.bsx.cfd.collector.storage.disk.csv.gz=true
   ```

5

To enable the addition of a timestamp as a filename prefix:

Note: The prefix format is the following:
*YYYYMMDDhhmmss_*
where *YYYYMMDDhhmmss* is the collection interval start time

1. Locate the section that begins with the following:

   ```
   # add time prefix on transfer option for CSV files
   ```

2. Set the following parameter to true, as shown below:

   ```
   com.alu.bsx.cfd.collector.storage.disk.csv.tpot=true
   ```

6

Save and close the CfdStorage.properties file.

END OF STEPS

5.28 To configure an AA Cflowd special-study policy

5.28.1 Steps

1

Open the NSP Flow Collector web UI. The Collection Policy page is displayed.

2

Click on the Special Study Policy tab.

3

Click Add. A new table row is displayed.

4

Configure the Filter Type parameter.

5

If the Filter Type parameter is set to something other than All Traffic per Subscriber, configure the Application / Application Group Name parameter.

6

Configure the Subscriber Type parameter.

7

Configure the Subscriber ID parameter.

Note: In the mobile domain, the subscriber ID must be prefixed by IMSI, MSISDN, or IMEI. For example, IMSI 88123398891xxxx.

8

If the Subscriber Type value is SAP, SDP Binding, or Business Transit Sub, configure the System ID parameter by specifying the system IP address of the host NE.

9

To delete a policy, select the Delete on save check box beside the policy.

10

Click Save Configuration. The configuration is saved.

End of steps

5.29 To configure an AA application or protocol filter

5.29.1 Steps

1

Open the NSP Flow Collector web UI. The Collection Policy page is displayed.

2

Click on the Application / Protocol Filters tab.

3

Click Add. A new table row is displayed.

4

Configure the Filter Type parameter.

5

Configure the Application / Protocol Name parameter.

6

To delete a filter, select the Delete on save check box beside the filter.

7

Click Save Configuration. The filters are applied to the next scheduled collection.

END OF STEPS

NSP Flow Collector Controller administration

5.30 Introduction

5.30.1 Description

The following procedures describe NSP Flow Collector Controller administration operations.

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

5.31 To start or stop an NSP Flow Collector Controller

5.31.1 Purpose

The following steps describe how to start or stop an NSP Flow Collector Controller that is on a dedicated station, or collocated with an NSP Flow Collector.

CAUTION

Data loss

Stopping an NSP Flow Collector Controller may affect the statistics collection of the associated NSP Flow Collectors.

Perform the procedure only under the guidance of technical support during a scheduled maintenance period.

5.31.2 Steps

1

Log in as the nsp user on the NSP Flow Collector Controller station.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/flow/fcc/bin ↵
```

4

To stop the NSP Flow Collector Controller, enter the following:

Note: If the NSP Flow Collector Controller is collocated with an NSP Flow Collector,

stopping the NSP Flow Collector Controller also stops the Flow Collector.

```
bash$ ./flowCollectorController.bash stop ↲
```

The NSP Flow Collector Controller stops.

5

To start the NSP Flow Collector Controller, enter the following:

Note: If the NSP Flow Collector Controller is collocated with an NSP Flow Collector, starting the NSP Flow Collector Controller also starts the Flow Collector.

```
bash$ ./flowCollectorController.bash start ↲
```

The NSP Flow Collector Controller starts.

6

Close the console window.

END OF STEPS

5.32 To display the NSP Flow Collector Controller status or release level

5.32.1 Purpose

The following steps describe how to view the operational status or software release of an NSP Flow Collector Controller that is on a dedicated station, or collocated with an NSP Flow Collector.

5.32.2 Steps

1

Log in as the nsp user on the NSP Flow Collector Controller station.

2

Open a console window.

3

Enter the following:

```
bash$ cd /opt/nsp/flow/fcc/bin ↲
```

4

To view the NSP Flow Collector Controller status, enter the following:

```
bash$ ./flowCollectorController.bash status ↲
```

The NSP Flow Collector Controller status is displayed.

If the NSP Flow Collector Controller is running, the line that begins with flow-collector-controller includes Started.

If the NSP Flow Collector Controller is not running, the following is displayed:

```
nspos-karaf.service is not running
```

5

To view the NSP Flow Collector Controller release level, enter the following:

```
bash$ ./flowCollectorController.bash version ↵
```

The NSP Flow Collector Controller release level is displayed

6

Close the console window.

END OF STEPS

5.33 To open the NSP Flow Collector Controller web UI

5.33.1 Purpose

The following steps describe how to open the NSP Flow Collector Controller web UI for Flow Collector Controller configuration.

5.33.2 Steps

1

Use a browser to open the following URL:

https://*server*:8443/fcc/admin

where *server* is the NSP Flow Collector Controller IP address or hostname

2

Enter the required user credentials and click OK. The NSP Flow Collector Controller web UI opens.

END OF STEPS

5.34 To force an NSP Flow Collector Controller to extract a network data snapshot

5.34.1 Purpose

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

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

CAUTION

Service Disruption

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

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

5.34.2 Steps

1

Open the NSP Flow Collector Controller web UI, as described in 5.33 “To open the NSP Flow Collector Controller web UI” (p. 114).

The NFM-P Configuration tab is displayed.

2

Click on the Operations tab.

3

To force the snapshot extraction for AA Cflowd statistics collection, click Force AA Snapshot Extraction.

The extraction begins.

4

To force the snapshot extraction for system Cflowd statistics collection, click Force SYS Snapshot Extraction.

The extraction begins.

5

Close the NSP Flow Collector Controller web UI.

END OF STEPS

6 NSP database administration

6.1 NSP database administration overview

6.1.1 Introduction

The following procedures describe how to:

- back up and restore the Kubernetes etcd database
- schedule or manually perform backups of the following:
  - Neo4j, PostgreSQL, and file server databases in an NSP cluster
  - Tomcat database of IP resource control
- restore the backed-up NSP system data in the event of a failure

Note: The following RHEL CLI prompts in command lines denote the active user, and are not to be included in typed commands:

- # —root user
- bash$ —nsp user

Note: NSP Analytics application data, such as the report repository contents, are stored in the PostgreSQL database, so are included in the database backup and restore operations described; no separate backup or restore process is required for Analytics application data.

Shared-mode deployments

In a shared-mode NSP deployment, you must synchronize the backup and restore operations among the systems in the deployment. See the backup and restore documentation for integrated systems such as the NFM-P and NFM-T, as required.

NSP database failure alarms

The NSP raises the following alarms in the Fault Management application response to a suspected PostgreSQL database failure:

Note: The alarms are not auto-clearing, so must be cleared manually.

- Critical—the leader database is unresponsive
- Major—at least one follower database is unresponsive

Identifying the source of a database alarm

In a shared-mode NSP and NFM-P deployment, the NSP and NFM-P raise similar alarms in response to a database failure.

Before you take action to respond to the alarm, you must identify the system that has raised the alarm, and which database instance is at fault.

The Source Type field of a database failure alarm indicates which system, NSP or NFM-P, has raised the alarm.

The Site ID and Site Name fields identify the following:

- NFM-P alarm—the NFM-P main server that raised the alarm
- NSP alarm—the faulty PostgreSQL database instance

Note: Regardless of the source system, the Additional Text field contains the IP address of the database instance that is at fault.

For example, the Source Type field of a standby database failure alarm contains “NFM-P”. An operator views the Site Name field, which identifies the NFM-P main server that has reported the fault. The operator then views the Additional Text field, and learns that the standby database associated with the main server has failed.

When a similar NSP alarm is raised, the operator has to view only the Site ID or Site Name field to identify which PostgreSQL database instance is at fault.

6.2 To configure scheduled NSP backups

6.2.1 Purpose

Perform this procedure to configure scheduled backups of the following NSP databases:

- NSP cluster Neo4j and PostgreSQL databases
- Tomcat database of IP resource control

Scheduled backups are enabled by default, and scheduled to run daily at 12:30 AM UTC.

Note: By default, the NSP retains the three most recent scheduled backups.

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

6.2.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Open the following file with a plain-text editor such as vi:

/opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

3

Locate the section that begins with the following:

```
backups:
```

4

Configure the following parameters:

Note: If the schedule value is an empty string, no scheduled backup is performed.

Note: See the RHEL cron man page for information about defining a crontab schedule.

```
schedule: "definition"
retained: n
```

where

definition is a UNIX crontab schedule definition; for example, “30 0 * * *” specifies the default backup schedule of 12:30 a.m. daily

n is the number of backups to retain

5

Close the open console windows.

END OF STEPS

6.3 To back up the databases in a hybrid NSP deployment

6.3.1 Purpose

Perform this procedure to manually back up the contents of the following databases:

- NSP cluster Neo4j and PostgreSQL databases
- Tomcat database of IP resource control

Note: The NSP performs scheduled daily IP resource control database backups, which are stored in the following directory for up to seven days:

/opt/nsp/backup/scheduled

A maximum of four backups can be saved for up to one month. The backup schedule is defined in the following file:

/opt/nsp/scripts/db/nsp-backup.conf

Note: If the NSP is an HA deployment, you must perform the backup on the active cluster member in each NSP cluster.

6.3.2 Steps

Back up traditionally-deployed component databases

1

Log in as the nsp user on the primary cross-domain resource control or IP resource control server.

2

Open a console window.

3

Enter the following:

```
bash$ nspdctl --host server backup -d backup_directory ↵
```

where

server is the server IP address or hostname

backup_directory is the name of a new directory that is to hold the database backup file set; if the directory already exists, the backup fails

The NSP backs up the databases.

4

Enter the following to verify that the backup completed successfully.

```
bash$ nspdctl --host server backup status ↵
```

where *server* is the server IP address or hostname

5

Transfer the backup files from *backup_directory* to a secure location in a separate facility for safekeeping.

Note: It is strongly recommended that for the greatest fault tolerance, you transfer the backup files to a secure facility that is outside the local data center.

6

Close the console window.

Back up NSP cluster databases

7

Perform 6.5 “To back up the NSP cluster databases” (p. 127).

END OF STEPS

6.4 To restore the databases in a hybrid NSP deployment

6.4.1 Purpose

Perform this procedure to restore the following NSP databases:

- NSP cluster Neo4j and PostgreSQL databases
- Tomcat database of IP resource control

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

6.4.2 Steps

Stop IP resource control, cross-domain resource control

1

You must stop each IP resource control and cross-domain resource control server.

Log in as the nsp user on an IP resource control or cross-domain resource control server, as required.

2

Open a console window.

3

Stop each server cluster.

Note: In a DR deployment, you must stop the standby server cluster first.

Enter the following:

Note: In an HA deployment, you must enter the command once for each cluster member.

```
bash$ nspdctl --host server_IP stop ↵
```

where *server_IP* is the cluster member IP address

Restore Tomcat database

4

Enter the following to switch to the root user:

```
bash$ su - ↵
```

5

Enter the following:

```
# cd NSP_installer_directory/tools/database ↵
```

where *NSP_installer_directory* is the directory that contains the extracted NSP software bundle

6

Restore the Tomcat database.

Note: In a DR deployment, you must perform this step once for each NSP cluster using the --target option to specify an NSP host in the cluster.

Note: The --target option is required only in a DR deployment.

Note: The --target option is not required in an HA deployment; the database is automatically restored on each NSP host in the cluster.

1. Enter the following:

   ```
   # ./db-restore.sh --target target_IP ↵
   ```

   where *target_IP* is one of the following, depending on the settings in the NSP configuration file:
   - the ansible_host value, if specified
   - the private IP address, if no ansible_host value is specified
   - the public IP address, if no advertised_address or ansible_host value is specified

   The following message and prompt are displayed:

   ```
   Verifying prerequisites...
   Starting database restore ...
   Backupset file to restore (.tar.gz format):
   ```

2. Enter the following and press ↵:

   path/nsp-tomcat_backup_*timestamp*.tar.gz

   where

   path is the absolute path of the Tomcat backup file

   timestamp is the backup creation time

   The following messages and prompt are displayed:

   ```
   PLAY [all] *********************************************
   [dbrestore : pause]
   Do you want to restore the nsp Tomcat db from file:
   path/nsp-tomcat_backup_timestamp.tar.gz? Press return to continue,
   or Ctrl+C to abort:
   ```

3. Press ↵.

   Messages like the following are displayed:

   ```
   TASK [dbrestore : Running nspdctl stop] *********************
   changed: [server_IP]
   ```

```
TASK [dbrestore : Perform database restore] *****************
changed: [server_IP]
TASK [dbrestore : Delete temporary directory] ***************
changed: [server_IP]
PLAY RECAP **********************************************
server_IP      : ok=n    changed=n    unreachable=n    failed=n
```

4. If the `failed` value is greater than zero, a restore failure has occurred; contact technical support for assistance.

7

Enter the following to switch back to the nsp user:

```
# su - nsp ↵
```

Start IP resource control, cross-domain resource control

8

Start each server cluster.

Note: In a DR deployment, you must start the primary cluster first.

Enter the following:

Note: In an HA deployment, you must enter the command once for each cluster member.

```
bash$ nspdctl --host server_IP start ↵
```

where *server_IP* is the cluster member IP address

The server starts.

9

Close the console window.

Prepare to restore NSP cluster databases

10

If you are restoring the data on an existing NSP cluster rather than on new NSP cluster VMs, you must stop the cluster and remove the cluster data.

For each existing NSP cluster in the deployment, perform the following steps on the NSP configurator VM in the cluster.

Note: For the existing clusters in a DR deployment, you must perform the steps first in the standby data center.

1. Log in as the root user on the NSP configurator VM.
2. Open the following file with a plain-text editor such as vi:

/opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

3. Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

   ```
   deleteOnUndeploy:true
   ```

4. Save and close the file.
5. Enter the following:

   ```
   # /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --undeploy ↵
   ```

   The NSP cluster is undeployed.

11

Perform the following steps on the NSP configurator VM in each data center.

Note: In a DR deployment, you must perform the steps first in the data center that you want to start as the standby data center.

1. Log in as the root user on the NSP configurator VM.
2. Open the following file with a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

3. Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

   ```
   deleteOnUndeploy:false
   ```

4. Save and close the file.
5. Enter the following to enter restore mode:

   ```
   # /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --restore ↵
   ```

6. Enter the following periodically to display the cluster status:

   ```
   # kubectl get pods ↵
   ```

   The cluster is ready for the restore when the status of the following pods is Running:

   - nsp-backup-storage-*n*
   - nspos-neo4j-core-default-*n*
   - nspos-postgresql-primary-*n*

Note: You must not proceed to the next step until the cluster is ready.

12

Perform Step 14 to Step 22 on the NSP configurator VM in each data center.

Note: In a DR deployment, you must perform the steps first in the data center that you want to start as the primary data center.

13

Go to Step 23.

Restore PostgreSQL database

14

Enter the following to copy the PostgreSQL backup file from local storage to the PostgreSQL backup pod:

```
# kubectl cp path/nspos-postgresql_backup_timestamp.tar.gz
nspos-postgresql-primary-0:tmp/restoreData ↵
```

where

path is the absolute path of the PostgreSQL backup file

timestamp is the backup creation time

15

Enter the following:

Note: The command may generate error messages about roles that exist; the messages can be ignored.

```
# kubectl exec -it nspos-postgresql-primary-0 -c nspos-postgresql
-- /opt/nsp/os/pgsql/scripts/pg-restore.sh -C -Q
-f /tmp/restoreData/PostgreSQL_backup_file ↵
```

where *PostgreSQL_backup_file* is the PostgreSQL backup file name

16

Enter the following:

```
# kubectl exec -it nspos-postgresql-primary-0 -c nspos-postgresql
-- rm -f /tmp/restoreData/PostgreSQL_backup_file ↵
```

Restore Neo4j database

17

Enter the following to uncompress the backup file:

```
# tar zxf local_dir/nspos-neo4j_backup_timestamp.tar.gz ↵
```

where

local_dir is the local directory that contains the Neo4j backup file

timestamp is the backup creation time

18

Enter the following for each cluster member:

```
# kubectl cp local_dir/graph.db namespace/nspos-neo4j-core-dc_
name-pod#:/tmp/restoreData ↵
```

where

local_dir is the local directory that contains the uncompressed Neo4j backup files

namespace is the Kubernetes namespace

dc_name is the data center name

pod# is the Neo4j pod number

19

Transfer the backup file in *local_dir* to the same directory on an NSP host in the other data center.

20

If the NSP cluster is an HA deployment, enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-pod# --
/var/lib/neo4j/bin/neo4j-admin unbind ↵
```

21

Enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-pod# --
/var/lib/neo4j/bin/neo4j-admin restore --force --database=graph.db
--from=/tmp/restoreData/graph.db ↵
```

22

Enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-pod# -- rm -rf
/tmp/restoreData/graph.db ↵
```

Start NSP system

23

Perform the following steps on the NSP configurator VM in each data center.

Note: In a DR deployment, you must perform the steps first in the data center that you want to start as the primary data center.

1. Open the following file with a plain-text editor such as vi:

   /opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

2. Edit the following line in the **platform** section, **kubernetes** subsection to read as shown below:

   ```
   deleteOnUndeploy:false
   ```

3. Save and close the file.

4. Enter the following to exit restore mode and terminate the restore pods:

   ```
   # /opt/nsp/NSP-CN-release-ID/bin/nsp-config.bash --undeploy ↵
   ```

5. Enter the following:

   ```
   # kubectl get pods ↵
   ```

The pods are listed.

6. If either of the following restore pods is listed, the pod is not terminated; return to substep 5.
   - nspos-neo4j-core-default-*n*
   - nspos-postgresql-primary-*n*

   Note: You must not proceed to the next step if a restore pod is listed.
7. Enter the following:

   # **/opt/nsp/NSP-CN-*release-ID*/bin/nsp-config.bash --deploy** ↵

   The NSP initializes using the restored data.
8. Enter the following periodically to display the cluster status:

   # **kubectl get pods** ↵

   The cluster is operational when the status of each pod is Running.

24

Close the open console windows.

END OF STEPS

6.5 To back up the NSP cluster databases

6.5.1 Purpose

Perform this procedure to manually back up the following databases in an NSP Kubernetes cluster:

- etcd
- Neo4j
- PostgreSQL

Note: *release-ID* in a file path has the following format:
R.r.p-rel.*version*
where
R.r.p is the NSP release, in the form *MAJOR.minor.patch*
version is a numeric value

6.5.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

If a common backup storage location is defined in the NSP configuration, go to Step 8.

3

Open the following file with a plain-text editor such as vi:

/opt/nsp/NSP-CN-*release-ID*/config/nsp-config.yml

4

To use an existing PVC, perform the following steps.

Note: The PVC must support ReadWriteMany semantics.

1. Locate the section that begins with the following:

   ```
   kubernetes:
   ```

2. Configure the following parameter in the section:

   ```
   rwxClass: "class"
   ```

   where *class* is the storage class

3. Locate the section that begins with the following:

   ```
   backups:
   ```

4. Configure the following parameter in the section:

   ```
   existingClaim: "store"
   ```

   where *store* is the name of the PVC store

5. Go to Step 8.

5

To use an existing NFS server, perform the following steps.

1. Locate the section that begins with the following:

   ```
   backups:
   ```

2. Configure the following parameters in the following subsection:

   ```
   nfs:
     server: "server"
     path: "path"
   ```

   where

   server is the NFS server IP address

   path is the path of the exported file system on the server

3. Go to Step 8.

6

To use an existing storage class that supports ReadWriteMany semantics, perform the following steps.

1. Locate the section that begins with the following:

   ```
   backups:
   ```

2. Configure the following parameters in the following subsection by adding the lines in boldface type:

```
storage:
  create:
    storageClass: class
    capacity: size
```

where

class is the storage class

size is the storage class capacity

3. Go to Step 8.

7

To use an existing storage class that supports only ReadWriteOnce semantics, perform the following steps.

Note: The NSP configurator provisions an NFS ReadWriteMany layer on the ReadWriteOnce storage for storing the backups.

1. Locate the section that begins with the following:

```
backups:
```

2. Configure the following parameter in the following subsection:

```
storage:
  capacity: size
```

where *size* is the expected backup storage requirement

8

Enter the following commands:

```
# kubectl create job etcd --from cronjob/nsp-etcd-backup ↵
# kubectl create job neo4j --from cronjob/nspos-neo4j-backup ↵
# kubectl create job postgresql --from cronjob/nspos-postgresql-backup
↵
```

The database backups begin.

9

To display the status of a backup job, enter the following:

```
# kubectl get pod -o wide | grep job ↵
```

where *job* is the backup job—etcd, neo4j, or postgresql

The backup job is finished when the status is Complete.

Note: You must not proceed to the next step until all backup jobs are finished.

10

If you use local storage, enter the following set of commands to copy the backup files to a local folder:

```
# kubectl cp $(kubectl get pods | awk '/nsp-backup-storage/ {print $1;exit}'):tmp/backups/nsp-etcd path/nsp-etcd ↵
# kubectl cp $(kubectl get pods | awk '/nsp-backup-storage/ {print $1;exit}'):tmp/backups/nspos-neo4j path/nspos-neo4j ↵
# kubectl cp $(kubectl get pods | awk '/nsp-backup-storage/ {print $1;exit}'):tmp/backups/nspos-postgresql path/nspos-postgresql ↵
```

where *path* is the absolute path of the local folder

11

Transfer each backup file to a remote, secure location for safekeeping.

Note: It is strongly recommended that for the greatest fault tolerance, you transfer the backup files to a secure facility that is outside the local data center.

12

Enter the following for each backup to delete the backup job:

```
# kubectl delete jobs.batch job ↵
```

where job is the backup job—etcd, neo4j, or postgresql

The backup job is deleted.

13

Close the open console windows.

END OF STEPS

6.6 To restore the etcd database in an NSP cluster

6.6.1 Purpose

CAUTION

System Data Corruption

Attempting to restore an etcd database backup from one NSP cluster in a different NSP cluster causes the NSP cluster restore to fail, and renders the cluster unrecoverable.

You must restore only an etcd database backup from the same NSP cluster; you cannot move an NSP cluster configuration to a different cluster, or restore a cluster configuration in a new cluster.

An etcd database backup is a snapshot of all Kubernetes objects and associated critical information. In an NSP cluster, an etcd database backup is created once each day.

Perform this procedure to help recover a failed NSP cluster by restoring the etcd database from a backup.

6.6.2 Steps

Obtain and distribute snapshot

1

Log in as the root user on an NSP cluster member in the primary data center.

2

Enter the following to copy the snapshot from the backup pod to an empty directory on the local file system:

```
# kubectl cp namespace/nsp-backup-storage-0:
/tmp/backups/nsp-etcd/nsp-etcd_backup_timestamp.tar.gz local_path ↵
```

where

namespace is the Kubernetes namespace

timestamp is the snapshot creation time

local_path is the empty local directory

3

Identify the cluster members on which etcd is running.

1. Open the /etc/etcd.env file for viewing.
2. Locate the following parameters:
   - ETCD_INITIAL_CLUSTER
   - ETCD_INITIAL_CLUSTER_TOKEN
   - ETCD_INITIAL_ADVERTISE_PEER_URLS
3. Record the parameter values.
4. Close the file.

4

Enter the following as the root user on each cluster member identified in Step 3:

Note: After you perform this step, the cluster is unreachable.

```
# systemctl stop etcd ↵
```

5

Transfer the snapshot file obtained in Step 2 to each cluster member.

Restore database on members

6

Perform Step 8 to Step 16 on each member station.

7

Go to Step 17.

8

Log in as the root user.

9

Navigate to the directory that contains the transferred snapshot file.

10

Enter the following:

```
# tar xzf path/nsp-etcd_backup_timestamp.tar.gz ↵
```

where

path is the absolute path of the snapshot file

timestamp is the snapshot creation time

The snapshot file is uncompressed.

11

Enter the following:

```
# ETCDCTL_API=3 etcdctl snapshot restore etcd.db --name etcdN
--initial-cluster cluster --initial-cluster-token token
--initial-advertise-peer-urls URL ↵
```

where

N is the member number in the recorded ETCD_INITIAL_CLUSTER value, for example, 1, 2, or 3

cluster is the recorded ETCD_INITIAL_CLUSTER value

token is the recorded ETCD_INITIAL_CLUSTER_TOKEN value

URL is the recorded ETCD_INITIAL_ADVERTISE_PEER_URLS value

The etcd database is restored.

12

Enter the following to create a directory in which to store the previous database:

```
# mkdir path/old_etcd_db ↵
```

where *path* is the absolute path of the created directory

13

Enter the following to move the previous database files to the created directory:

```
# mv /var/lib/etcd/* path/old_etcd_db ↲
```

where *path* is the absolute path of the directory created in Step 12

14

Enter the following:

```
# mv ./etcd1.etcd/* /var/lib/etcd/ ↲
```

The uncompressed snapshot files move to the /var/lib/etcd directory.

15

Enter the following:

```
# systemctl start etcd ↲
```

The etcd service starts.

16

Enter the following:

```
# systemctl status etcd ↲
```

The etcd service status is displayed.

The service is up if the following is displayed:

```
Active: active (running)
```

17

When the service is up, close the open console windows.

END OF STEPS

6.7 To back up the NSP file service application data

6.7.1 Purpose

Perform this procedure as required to manually back up the application data of the NSP file service; for example, after you import an MDM device adaptor or customize an MDM device mapping, or before a system upgrade.

6.7.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Perform one of the following to change to the file service pod.

a. For an NSP system at Release 20.11 or 21.3, enter the following:

```
# kubectl exec -it nsp-file-service-app-0 -c nsp-file-service-app
bin/bash ↲
```

b. For an NSP system at Release 21.6 or later, enter the following:

```
# kubectl exec -it nsp-file-service-app-0 bin/bash ↲
```

4

Enter the following:

```
# cd /opt/nsp/containers/nspvolume/fileservice ↲
```

5

Enter the following:

```
# tar -cvf fileServiceData.tar * ↲
```

The files that are backed up are listed, and added to a backup file in the current directory called fileServiceData.tar.

The files may include MDM adaptor suite and device mapping files.

6

Perform one of the following to copy the backup file to the local file system.

a. For an NSP system at Release 20.11 or 21.3, enter the following:

```
# kubectl cp nsp-file-service-app-0:
/opt/nsp/containers/nspvolume/fileservice/fileServiceData.tar
fileServiceData.tar -c nsp-file-service-app ↲
```

b. For an NSP system at Release 21.6 or later, enter the following:

```
# kubectl cp nsp-file-service-app-0:
/opt/nsp/containers/nspvolume/fileservice/fileServiceData.tar
fileServiceData.tar ↲
```

7

Transfer the backup file from the current directory to a secure location in a separate facility for safekeeping.

Note: It is strongly recommended that for the greatest fault tolerance, you transfer the backup file to a secure facility that is outside the local data center.

8

Close the console window.

END OF STEPS

6.8 To restore the NSP file service application data

6.8.1 Purpose

Perform this procedure as required to restore the application data of the NSP file service; for example, in the event that the data is lost during a file service or system failure.

6.8.2 Steps

1

Log in as the root user on the NSP configurator VM.

2

Open a console window.

3

Enter the following to copy the backup data file to the file-service-pod:

```
# kubectl cp fileServiceData.tar nsp-file-service-app-0:
/opt/nsp/containers/nspvolume/fileservice/fileServiceData.tar ↲
```

4

Enter the following to change to the file service pod:

```
# kubectl exec -it nsp-file-service-app-0 bin/bash ↲
```

5

Enter the following:

```
# cd /opt/nsp/containers/nspvolume/fileservice ↲
```

6

Enter the following:

```
# tar -xvf fileServiceData.tar ↲
```

The files are restored.

When the operation is complete, the MDM adaptor suites and customized device mapping files load automatically.

7

Close the console window.

END OF STEPS