NOKIA

NSP
Network Services Platform

Release 20.11

System Administrator Guide

3HE-16052-AAAD-TQZZA

Issue 4

January 2022

Legal notice

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

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 system 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 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 topics:

- TLS administration, such as version support and certificate expiry; see 2.3 “NSP TLS administration” (p. 10)
- system security management; see 2.4 “NSP security administration procedures” (p. 12)

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

2.2 User password policies

2.2.1 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.

2.3 NSP TLS administration

2.3.1 NSP TLS administration overview

The NSP uses TLS to secure the external interfaces between NSP components and with client applications. Additionally, if all NSP system components are at Release 19.6 or later, you can enable TLS on some internal nspOS subsystems and services.

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.

Note: By default, the NSP uses only TLS 1.2.

Note: NSP component communication is performed using IPv4 only; IPv6 communication is not supported.

Kubernetes infrastructure TLS

The Kubernetes infrastructure in a container-based NSP deployment is secured using TLS. The associated TLS certificates expire one year from installation, at which point the certificates must be replaced, or communication between Kubernetes nodes or services may cease.

See "Kubernetes infrastructure TLS certificate management" (p. 11) for information about updating the certificates.

Internal NSP 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 an upgrade from a Release earlier than 19.6, 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.

External NSP 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.

2.3.2 Managing NSP TLS certificates

NSP 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

2.6 “To replace the NSP TLS certificates” (p. 17) describes how to replace the internal TLS certificate, the external certificate, or both, in an NSP system.

Kubernetes infrastructure TLS certificate management

The TLS certificates for the Kubernetes infrastructure expire one year after installation. The expiry date is embedded in Kubernetes, and cannot be changed.

Note: The NSP is unable to provide notification of the certificate expiry; the recommended best practice is to update the certificates somewhat frequently to avoid closely approaching the expiry date.

See 2.5 “To replace the Kubernetes infrastructure TLS certificates” (p. 12) for certificate renewal information.

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: A certificate expiry alarm is not self-clearing, so must be manually cleared.

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.

2.4 NSP security administration procedures

2.4.1 Introduction

The following procedures describe NSP security administration operations.

Note: The NSP provides no notification of Kubernetes certificate expiry.

Note: NSP PKI server operation is described in the *NSP Deployment and Installation Guide*.

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 To replace the Kubernetes infrastructure TLS certificates

2.5.1 Purpose

Perform this procedure to replace the expiring TLS certificates used by the Kubernetes infrastructure in a Nokia container NSP deployment.

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.

2.5.2 Steps

1

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

2

You must identify the master nodes in the cluster; enter the following:

```
# kubectl get nodes -o wide ↲
```

A VM list 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

Record the name of each node whose role is master.

4

Perform Step 6 to Step 17 on each master node.

5

Go to Step 18.

6

Log in as the root user.

7

Open a console window.

8

Enter the following to display the certificate expiry information:

```
# kubeadm alpha certs check-expiration --config=
"/etc/kubernetes/kubeadm-config.yaml" ↲
```

Output like the following is displayed; the EXPIRES value is the expiry date, and the RESIDUAL value is the number of days until the expiry:

```
CERTIFICATE                EXPIRES                  RESIDUAL TIME
EXTERNALLY MANAGED
admin.conf                 Oct 15, 2022 15:19 UTC   364d            no
apiserver                  Oct 15, 2022 15:19 UTC   364d            no
apiserver-kubelet-client   Oct 15, 2022 15:09 UTC   364d            no
controller-manager.conf    Oct 15, 2022 15:10 UTC   364d            no
front-proxy-client         Oct 15, 2022 15:09 UTC   364d            no
scheduler.conf             Oct 15, 2022 15:10 UTC   364d            no
```

9

Enter the following sequence of commands to back up the current files:

```
# mkdir /root/certs_backup/ ↵
# cp /etc/kubernetes/*.conf /root/certs_backup/ ↵
# cp /etc/kubernetes/pki/* /root/certs_backup/ ↵
# cp /root/.kube/config /root/certs_backup/ ↵
```

10

Enter the following sequence of commands to renew the certificates:

```
# kubeadm alpha certs renew admin.conf --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
# kubeadm alpha certs renew apiserver --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
# kubeadm alpha certs renew apiserver-kubelet-client --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
# kubeadm alpha certs renew controller-manager.conf --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
# kubeadm alpha certs renew front-proxy-client --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
# kubeadm alpha certs renew scheduler.conf --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
```

The TLS certificates for the NSP cluster Kubernetes infrastructure are renewed.

11

Perform the following steps.

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

   /etc/kubernetes/kubelet.conf

2. Remove the following lines:

   ```
   client-certificate-data
   client-key-data
   ```

3. Add the lines in **boldface type** below to the position shown:

   ```
   users:
   - name: system:node:node_name
     user:
       client-certificate:
   /etc/kubernetes/pki/apiserver-kubelet-client.crt
       client-key: /etc/kubernetes/pki/apiserver-kubelet-client.key
   ```

4. Save and close the file.

12

Enter the following to restart the required containers, which then import the new certificate files.

```
# docker ps -af 'name=k8s_POD_
(kube-apiserver|kube-controller-manager|kube-scheduler)-*' -q | xargs
docker rm --force ↵
```

13

Enter the following to restart the kubelet service:

Note: The kubectl command is unavailable during the restart, which takes approximately one minute.

```
# systemctl restart kubelet ↵
```

The kubelet service restarts, and creates the following file:

/etc/kubernetes/admin.conf

14

Enter the following:

```
# cp -i /etc/kubernetes/admin.conf /root/.kube/config ↵
```

15

Perform the following steps.

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

   /root/.kube/config

2. Remove the following lines:

   ```
   client-certificate-data
   client-key-data
   ```

3. Add the lines in **boldface type** below to the position shown:

   ```
   users:
   - name: kubernetes-admin
     user:
       client-certificate:
   /etc/kubernetes/pki/apiserver-kubelet-client.crt
       client-key: /etc/kubernetes/pki/apiserver-kubelet-client.key
   ```

4. Save and close the file.

16

Enter the following to display the certificate expiry information:

```
# kubeadm alpha certs check-expiration --config=
"/etc/kubernetes/kubeadm-config.yaml" ↵
```

Output like the following is displayed; the EXPIRES value is the expiry date, and the RESIDUAL value is the number of days until the expiry:

```
CERTIFICATE                EXPIRES                  RESIDUAL TIME
EXTERNALLY MANAGED
```

```
admin.conf                Oct 15, 2022 15:19 UTC   364d            no
apiserver                 Oct 15, 2022 15:19 UTC   364d            no
apiserver-kubelet-client  Oct 15, 2022 15:09 UTC   364d            no
controller-manager.conf   Oct 15, 2022 15:10 UTC   364d            no
front-proxy-client        Oct 15, 2022 15:09 UTC   364d            no
scheduler.conf            Oct 15, 2022 15:10 UTC   364d            no
```

17

Enter the following to verify that the kubectl command is available:

```
# kubectl get pods -A ↵
```

If the command is available, all pods are listed.

Optional verification actions

18

To verify that Kubernetes is using the correct certificate, perform the following steps.

1. Enter the following as the root user on the NSP deployer VM:

   ```
   # netstat -nap | grep kube-schedule | grep -i tcp ↵
   ```

   Output like the following is displayed:

   ```
   tcp   0   0 192.168.96.8:35914  IP:6443        ESTABLISHED
   4205/kube-scheduler
   tcp6  0   0 :::10259              :::*         LISTEN
   4205/kube-scheduler
   tcp6  0   0 :::10251              :::*         LISTEN
   4205/kube-scheduler
   ```

2. Enter the following:

   ```
   # openssl s_client -connect IP:6443 -servername IP 2>/dev/null |
   openssl x509 -noout -dates ↵
   ```

   where *IP* is the *IP* value displayed in substep 1

   The certificate information is displayed.

19

To display the current certificate expiry date, enter the following:

```
# openssl x509 -in /var/lib/kubelet/pki/kubelet-client-current.pem
-text | grep Not ↵
```

The expiry date is displayed.

20

Close the open console windows.

END OF STEPS

2.6 To replace the NSP TLS certificates

2.6.1 Purpose

Perform this procedure to do one or both of the following:

- replace the internal certificate used by all NSP components
- replace or install the external certificate in an NSP system, or on individual components

Note: In a DR deployment, you must perform the procedure first in the standby data center, and then in the primary data center.

Note: You require root user privileges on each NSP component.

Note: The install.sh utility for RPM-based component deployment requires SSH access to each target station. To enable SSH access, you must do one of the following.

- Configure the required SSH key on each station.
- If each station has the same root user password, add the --ask-pass argument to the install.sh command as follows:

  `./install.sh --ask-pass --`*`option`*

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

2.6.2 Steps

1

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

2

Open a console window.

Stop NSP cluster

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. If you are changing the deployment, such as adding or removing a component, or changing a component address, update the configuration parameters as required.
4. Save and close the file.
5. Enter the following:

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

Configure and deploy NSP cluster

4

The PKI server on the NSP configurator VM requires the new CA certificate files.

Copy the following files to the /opt/nsp/NSP-CN-*release-ID*/tls/ca directory:

- ca_internal.key
- ca_internal.pem
- ca.key
- ca.pem

5

Enter the following to start the NSP cluster:

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

Stop additional components

6

If the system includes one or more NSP analytics servers, stop each analytics server in the local data center, as described in 5.2 “To start or stop an NSP analytics server” (p. 49).

7

If the system includes NSP Flow Collectors, stop each NSP Flow Collector in the local data center, as described in 5.7 “To start or stop an NSP Flow Collector” (p. 55).

8

If the system includes NSP Flow Collector Controllers, stop each NSP Flow Collector Controller in the local data center, as described in 5.12 “To start or stop an NSP Flow Collector Controller” (p. 60).

9

If the system includes NFM-P auxiliary servers, stop each auxiliary server in the local data center, as described in the *NSP NFM-P Administrator Guide*.

10

If the system includes NFM-P main servers, stop each main server in the local data center, as described in the *NSP NFM-P Administrator Guide*.

Configure IP resource control, cross-domain resource control

11

If the NSP deployment includes IP resource control or cross-domain resource control, perform the following steps

1. On the NSP configurator VM, enter the following:

   # **cd /opt/nsp/NSP-CN-*release-ID*/tools/pki** ↵

2. Enter the following:

   # **cp /opt/nsp/NSP-CN-*release-ID*/tls/ca/ca* .** ↵

3. Enter the following:

   # **./pki-server &** ↵

   The PKI server starts.

4. Log in as the root user on the station used for installing IP resource control or cross-domain resource control.

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 NSP installation files

6. If you are updating the PKI-generated TLS certificates, 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. If you are updating the custom-CA-signed TLS certificates, configure the following parameters in the **tls** section using the new values:
   - custom_keystore_password
   - custom_truststore_password
   - custom_key_alias
   - custom_keystore_path
   - custom_truststore_path

8. Save and close the file.

9. Enter the following to update the IP resource control and cross-domain resource control configuration:

```
# ./install.sh ↵
```

Configure remaining NSP components

12

Configure the remaining NSP components 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*.

Note: Auxiliary databases do not support TLS, so require no configuration.

13

If you are replacing the internal certificate, perform the following steps.

1. Enter the following as the nsp user on each NSP analytics server station:

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

2. If the NSP deployment includes the NFM-P, use the samconfig utility on each NFM-P main and auxiliary server station to set the regenerate-certs parameter to true; see the *NSP NFM-P Installation and Upgrade Guide* for information about using the samconfig utility.

14

If the NSP deployment includes the NFM-T, configure the NFM-T to obtain the new TLS configuration; see the *NSP Deployment and Installation Guide* for information.

Start remaining components

15

If the system includes NFM-P main servers, start each main server in the local data center, as described in the *NSP NFM-P Administrator Guide*.

16

If the system includes NFM-P auxiliary servers, start each auxiliary server in the local data center, as described in the *NSP NFM-P Administrator Guide*.

17

If the system includes one or more NSP Flow Collector Controllers, start each NSP Flow Collector Controller in the local data center, as described in 5.12 “To start or stop an NSP Flow Collector Controller” (p. 60).

18

If the system includes one or more NSP Flow Collectors, start each NSP Flow Collector in the local data center, as described in 5.7 “To start or stop an NSP Flow Collector” (p. 55).

19

If the system includes one or more NSP analytics servers, start each analytics server in the local data center, as described in 5.2 “To start or stop an NSP analytics server” (p. 49).

Stop PKI server

20

If the PKI server is running, enter CTRL+C in the console window to stop the PKI server.

Note: You must not stop the PKI server until each NSP component has obtained the required TLS artifacts from the PKI server.

21

Close the open console windows.

END OF STEPS

2.7 To change the nsp system user password

2.7.1 Purpose

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

2.7.2 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

2.8 To whitelist an analytics server for OSS report requests

2.8.1 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 whitelist are accepted as NSP clients. Perform this procedure to disable the blocking of report retrieval by an OSS.

Similarly, when an NSP deployment includes an NSP Analytics server and an NFM-P at a release earlier than 18.12, the CORS filtering may result in a 401 error. Performing the procedure resolves this issue also.

2.8.2 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 a whitelist target:

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

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

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

The NSP analytics server address is added to the whitelist.

5

Close the console window.

END OF STEPS

3 NSP application administration

3.1 NSP application administration overview

3.1.1 Introduction

This chapter describes NSP application access requirements, general application configuration settings, and best practices for application access.

3.2 NSP application access and browser support

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 “To configure global NSP application settings” (p. 28).

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

- Inventory Management
- Service Navigator
- 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 “To activate or deactivate NSP applications” (p. 31) for information.

3.2.2 Browser access to redundant NSP servers

If you open a browser to the primary NSP URL in a redundant 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 access to geo-redundant NSP system

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

See 4.12 “To enable single-address DR NSP system access” (p. 44) 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 Nokia Network Developer Portal for more information about OSS access to the NSP using a REST API.

3.2.4 Browser support

All NSP applications are supported on the latest version of Google Chrome. For information about additional supported browsers for NSP applications, see the NSP NFM-P Planning Guide.

Note: In order for the Apple Safari web browser to open the Analytics application, you must ensure that the following Safari privacy settings are configured, if present in your browser version:

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

Note: If you are using Chrome or Firefox on Windows 8.1 or Windows Server 2012, it is recommended that you enable ClearType Text for optimal viewing of fonts:

In the Windows Control Panel, open the Display settings, and enable the Turn on ClearType parameter under the Adjust ClearType text settings.

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

Note: 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.

3.2.5 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.

3.2.6 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.

3.2.7 Best practices for application access

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

- Although other browser types are supported, Chrome is the preferred browser.
- 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.

3.2.8 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
Ctrl + v Command + v in Apple/OSX	Paste
Enter	Open folder or expandable object such as tile Invoke action on button or menu item

Keystroke	Action
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 To 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

From the NSP Launchpad, click User, Settings.

3

Click System Settings.

4

Set the application polling time.

5

Choose a Language from the drop-down menu.

6

Type a security statement in the text field, and then select the check box to enable the security statement.

7

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 Installation and Upgrade Guide*, is recommended.

8

Select the Color row with severity in IP and Wireless applications parameter, if required.

9

If required, configure which timezone is used when displaying timestamps for alarms by selecting an option in the Alarm Time View Mode parameter.

Note: After changing this parameter, reload any active applications to update the displayed time stamps.

10

Configure settings for physical maps in the Network Supervision and Group Manager applications, if required:

- 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

11

Click Save.

END OF STEPS

3.4 To configure the NSP alarm-severity colors

3.4.1 Purpose

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

3.4.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose User, Settings from the NSP Launchpad.

3

Click System Colors.

4

Under Alarms, click on an alarm severity category and then click on the color you want to associate with the alarm severity category.

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

5

Select a text color.

6

Click Save.

END OF STEPS

3.5 To configure NSP 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 Launchpad.

3

Click Linked URLs.

4

Configure the Display Name and URL parameters.

5

Click Add.

6

To remove a linked URL, hover over the URL item in the list and click the Delete button at the end of the row.

END OF STEPS

3.6 To 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.

3

Click App Deployment Control.

4

Expand an application category, and then select or deselect the required check boxes to activate or deactivate applications in the category.

5

Select the check box to indicate that you understand the implications of the change.

6

Click Save.

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

Note: The on-product help of a reactivated application may not be available in the Help Center for up to 24 hours after the reactivation.

7

If you are deactivating any application, restart the NSP; see 4.5 “To start or stop IP resource control or cross-domain resource control” (p. 39) for information about how to perform the actions in the following steps.

1. If the NSP system is redundant, stop the standby NSP server.
2. Stop the primary or standalone NSP server.
3. Start the primary or standalone NSP server.
4. If the NSP system is redundant, start the standby NSP server.

END OF STEPS

3.7 To 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 Launchpad.

3

Click Event Logging Policy.

4

To enable event logging, select the Enable Event Logging parameter.

5

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

6

To delete all event records from the database, click Purge Event Records.

7

Click Save.

End of steps

3.8 To configure an e-mail server for alarm notifications

3.8.1 Purpose

Use this procedure to configure connection information to an e-mail server for use with Fault Management alarm e-mail policies.

3.8.2 Steps

1

Sign in to the NSP as an administrator.

2

Choose User, Settings from the NSP Launchpad.

3

Click E-mail Server Settings.

4

Specify the address of an e-mail server, and the user name and credentials of an administrative user.

5

Click Save.

End of steps

4 NSP system administration

4.1 NSP system management

4.1.1 Introduction

The procedures in this chapter describe system-level NSP management and configuration operations such as component startup and shutdown, status display, and enabling specific functions.

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

4.1.2 NSP Kubernetes cluster shutdown and startup

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 4.2 “Workflow to stop and start DR NSP Kubernetes clusters” (p. 36).

See the following procedures for information about stopping and starting an NSP Kubernetes cluster:

- 4.3 “To stop an NSP Kubernetes cluster” (p. 37)
- 4.4 “To start an NSP Kubernetes cluster” (p. 38)

4.1.3 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 raised the alarm. The Source Type field of a system alarm indicates which system has generated 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 contains “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 raised the alarm. However, for an NSP database, the alarm fields identify the PostgreSQL database instance that is at fault.

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

4.2 Workflow to stop and start DR NSP Kubernetes clusters

4.2.1 Description

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:

- 4.3 “To stop an NSP Kubernetes cluster” (p. 37)
- 4.4 “To start an NSP Kubernetes cluster” (p. 38)

4.2.2 Stages

1

Stop the standby cluster.

2

Perform the required maintenance on the standby cluster.

3

Start the standby cluster.

4

Perform an activity switch to change the standby cluster role to active, as described in 4.10 “To perform a manual switchover in a DR NSP deployment” (p. 43).

The standby cluster assumes the active role.

5

Stop the former active cluster.

6

Perform the required maintenance on the former active cluster.

7

Start the former active cluster.

8

If required, perform an activity switch to restore the initial active and standby roles, as described in 4.10 “To perform a manual switchover in a DR NSP deployment” (p. 43).

4.3 To stop an NSP Kubernetes cluster

4.3.1 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.

Perform this procedure 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 4.2 “Workflow to stop and start DR NSP Kubernetes clusters” (p. 36).

4.3.2 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, where *R_r* is the NSP release ID:

/opt/nsp/NSP_CN_*R_r*/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          9h
```

8

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

END OF STEPS

4.4 To start an NSP Kubernetes cluster

4.4.1 Purpose

Perform this procedure to start the Kubernetes software in an NSP cluster.

4.4.2 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

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

4.5.1 Purpose

Perform this procedure 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.

4.5.2 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

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

4.6.1 Purpose

Perform this procedure to view the operational status of an IP resource control or cross-domain resource control instance.

4.6.2 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

4.7 To apply an NSP license to an IP resource control instance

4.7.1 Purpose

Perform this procedure to apply a new or updated license to an IP resource control instance.

4.7.2 Steps

1

Copy the required license files to the *NSP_installer_directory*/license directory, where *NSP_installer_directory* is the directory contains the extracted NSP software bundle.

2

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

3

Open a console window.

4

Enter the following:

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

5

Enter the following to apply the license:

```
bash$ ./install.sh ↵
```

The license is distributed to each IP resource control instance.

6

Enter the following to switch to the root user:

```
bash$ su ↵
```

7

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

Note: In a DR deployment, you must perform the step on each IP resource control instance.

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

The web server restarts, and the license is applied.

8

Close the open console windows.

END OF STEPS

4.8 To identify the master node in an HA NSP cluster

4.8.1 Purpose

Perform this procedure to identify the master node in an HA NSP cluster.

4.8.2 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

4.9 To display the NSP cluster status

4.9.1 Purpose

Perform this procedure to display the status of the NSP clusters.

Note: You require nsp user privileges on each NSP cluster member in each data center.

4.9.2 Steps

1

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

2

In a standalone NSP deployment, issue the following REST API call:

```
GET /nsp-role-manager/address/status
```

where *address* is the NSP cluster IP address or hostname

3

For a DR deployment, issue the following REST API call:

`GET /nsp-role-manager/`*`address`*`/statusAll`

where *address* is the NSP cluster IP address or hostname

4

Log out of the station.

END OF STEPS

4.10 To perform a manual switchover in a DR NSP deployment

4.10.1 Purpose

Perform this procedure to switch the active and standby roles of the NSP server clusters in a DR deployment.

Note: You require nsp user privileges on each NSP cluster member in each data center.

4.10.2 Steps

1

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

2

The Role Manager supports the following REST APIs, which are used internally for inter-site calls, but are also available for external use:

- GET /nsp-role-manager/server/status—returns the current status
- GET /nsp-role-manager/server/statusAll—returns the current status for the given site and the DR peer, if present
- GET /nsp-role-manager/server/role?role=*role*&force=*force*—sets the role, where *role* is a string value, and *force* is a Boolean value
- POST/PUT /nsp-role-manager/server/role—sets the role; the expected payload is a JSON object that contains the *role* and *force* fields.

3

Issue one of the following REST API calls:

a. `GET /nsp-role-manager/`*`address`*`/role?role=`*`role`*`&force=`*`force`*

where

address is the NSP virtual IP address

role is active or standby

force is true or false, and determines whether the role change is forced

The NSP attempts to change the cluster roles.

b. `POST /nsp-role-manager/address/role`

where *address* is the NSP virtual IP address

Note: The call requires a JSON object that contains the *role* and *force* parameter values described in Step 3 b.

The NSP attempts to change the cluster roles.

4

Log out of the station.

END OF STEPS

4.11 NSP system configuration

4.11.1 Introduction

The following procedures describe general configuration operations for all NSP deployment types.

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

4.12 To enable single-address DR NSP system access

4.12.1 Purpose

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

You can implement a reverse proxy that presents only one IP address for system access. A 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.

4.12.2 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 VIP addresses of the NSP clusters

5

Close the console window.

END OF STEPS

4.13 To enable additional IP resource control functions

4.13.1 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.

4.13.2 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

4.14 To disable NSP websocket event notifications

4.14.1 Purpose

Websocket-based events are used by the NSP applications. Perform this procedure 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.

4.14.2 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

5 NSP component administration

5.1 NSP analytics server administration

5.1.1 Introduction

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

5.2 To start or stop an NSP analytics server

5.2.1 Purpose

Perform this procedure to start or stop the NSP analytics server software on a station.

5.2.2 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

5.3 To manage images on an analytics server

5.3.1 Purpose

Perform this procedure 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.

5.3.2 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

5.4 To enable and manage analytics server logging

5.4.1 Purpose

Perform this procedure 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.

5.4.2 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

5.5 To collect analytics-server log files

5.5.1 Purpose

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

5.5.2 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/<CUSTOMER_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

5.6 NSP Flow Collector administration

5.6.1 Introduction

The following procedures describe NSP Flow Collector 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

5.7 To start or stop an NSP Flow Collector

5.7.1 Purpose

Perform this procedure 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.

5.7.2 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

5.8 To display the NSP Flow Collector status or release level

5.8.1 Purpose

Perform this procedure 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.

5.8.2 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 relase level is displayed.

6

Close the console window.

END OF STEPS

5.9 To open the NSP Flow Collector web UI

5.9.1 Purpose

Perform this procedure to open the NSP Flow Collector web UI for Flow Collector configuration.

5.9.2 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

Enter the required user credentials and click OK. The NSP Flow Collector web UI opens.

END OF STEPS

5.10 To configure NSP Flow Collector statistics aggregation

5.10.1 Purpose

Perform this procedure to enable the aggregation of different statistics types on an NSP Flow Collector.

5.10.2 Steps

1

Open the NSP Flow Collector web UI, as described in 5.9 “To open the NSP Flow Collector web UI” (p. 57).

The Collection Policy tab 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

Configure the aggregations.

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
- 80 000 flows/s—5 minutes

- 40 000 flows/s—15 minutes

1. Use the Interval drop-down menus in the Aggregation Intervals panel to specify the aggregation interval for each statistic type, as required.
2. 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.
3. Click on the tab in the lower panel that corresponds to the statistic type.
4. Select or deselect aggregations, as required.

5

Configure the transfer of BB NAT records in CSV format to a file server, if required.

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.

1. Click on the NAT Transfer tab.
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

6

Click Save Configuration. The configuration is saved.

7

Close the NSP Flow Collector web UI.

END OF STEPS

5.11 NSP Flow Collector Controller administration

5.11.1 Introduction

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

5.12 To start or stop an NSP Flow Collector Controller

5.12.1 Purpose

Perform this procedure 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.

5.12.2 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

5.13 To display the NSP Flow Collector Controller status or release level

5.13.1 Purpose

Perform this procedure 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.

5.13.2 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

5.14 To open the NSP Flow Collector Controller web UI

5.14.1 Purpose

Perform this procedure to open the NSP Flow Collector Controller web UI for Flow Collector Controller configuration.

5.14.2 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

5.15 To force an NSP Flow Collector Controller to extract a network data snapshot

5.15.1 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.

Perform this procedure 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.

5.15.2 Steps

1

Open the NSP Flow Collector Controller web UI, as described in 5.14 “To open the NSP Flow Collector Controller web UI” (p. 62).

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

5.16 MDM administration

5.16.1 Introduction

The workflow and procedures in this section describe how to perform MDM administration operations.

5.17 Workflow to commission a device for model-driven management

5.17.1 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.

5.17.2 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 the 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.

5.18 To restart an MDM server instance

5.18.1 Purpose

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

5.18.2 Steps

1

Log in as the root user on an NSP cluster member in the active data center.

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 to restart an mdm-server instance:

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

where *n* is the mdm-server pod number

The MDM instance restarts.

5

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

6

Close the console window.

End of steps

5.19 To enable TLS for MDM telemetry and gNMI on_change support

5.19.1 Purpose

To enable TLS communication between MDM and managed NEs, 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.

Perform this procedure to secure the following NSP communication with NEs by importing an NE TLS certificate:

- MDM telemetry
- gNMI on_change notifications

Note: The gRPC certificates are separate from the certificates used for secure communication within the NSP system.

5.19.2 Steps

1

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

2

Transfer the NE certificate files to the /opt/nsp/tls/telemetry directory on the NSP deployer host.

3

Enter the following:

```
# /opt/nsp/bin/nsp-config.bash --config ↵
```

4

Enter the following:

```
# /opt/nsp/bin/nsp-config.bash --deploy ↵
```

The certificate files are imported to each MDM instance during the next MDM server pod restart.

Import certificates to MDM

5

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

a. Manually import the certificate.

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

1. Enter the following to copy the certificate files to each MDM server pod:

   # **kubectl cp *source_path pod*:*destination_path*** ↵

   where

   source_path is the local path of the certificate files

   pod is the MDM server pod name

   destination_path is the path of the directory on each MDM server pod that is to contain the copied files

2. Go to Step 6.

b. Restart the MDM server pod.

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

1. Perform 5.18 “To restart an MDM server instance” (p. 64) for each MDM server pod to stop and then restart each pod.
2. Go to Step 7.

6

Perform the following steps on each MDM server.

1. Log in as the root user on the MDM server.
2. Enter the following:

   # **cd /opt/nsp/os/jre/bin** ↵

3. Enter the following:

   # **./keytool -import -alias *alias* -keystore nsp.truststore -file *path*** ↵

   where

   alias is the TLS keystore alias

   path is the path of the TLS certificate files on the MDM server

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

7

Close the open console windows.

END OF STEPS

6 NSP database administration

6.1 NSP database administration overview

6.1.1 Introduction

The following procedures describe how to:

- preserve the NSP system data using scheduled or manual backups
- restore the NSP system data in the event of a system failure
- restore the Kubernetes etcd database in an NSP cluster

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:

- Neo4j database of IP resource control
- Neo4j database of containerized NSP cluster
- PostgreSQL database of containerized NSP cluster

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.

6.2.2 Steps

1

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

2

Open the /opt/nsp/NSP_CN_*R_r*/config/nsp-config.yml file with a plain-text editor such as vi, where *R_r* is the NSP release ID.

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:

- Neo4j database of IP resource control
- Neo4j database of containerized NSP cluster
- PostgreSQL database of NSP cluster

Note: The NSP performs scheduled daily 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 member of 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 for safekeeping.

6

Close the console window.

Back up containerized NSP databases

7

Perform 6.5 “To back up the NSP databases in a containerized deployment” (p. 77).

END OF STEPS

6.4 To restore the databases in a hybrid NSP deployment

6.4.1 Purpose

Perform this procedure to restore the contents of the following databases:

- Neo4j database of IP resource control
- Neo4j database of containerized NSP cluster
- PostgreSQL database of containerized NSP cluster

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 for containerized database restore

10

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 /opt/nsp/NSP_CN_*R_r*/config/nsp-config.yml file using a plain-text editor such as vi, where *R_r* is the NSP release ID.
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-CN-R_r/bin/nsp-config.bash --restore ↵
   ```

   where *R_r* is the NSP release ID
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.

11

Perform Step 13 to Step 15 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 active data center.

12

Go to Step 16.

Restore PostgreSQL database

13

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-n:tmp/restoreData ↵
```

where

path is the absolute path of the PostgreSQL backup file

timestamp is the backup creation time

n is a pod ID

14

Enter the following:

Note: The command may generate error messages about roles that exist; the messages can be ignored.

```
# kubectl exec -it pod_name -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

pod_name is the nspos-postgresql pod name

15

Enter the following:

```
# kubectl exec -it nspos-postgresql-primary-n -c nspos-postgresql
-- rm -f /tmp/restoreData/PostgreSQL_backup_file ↵
```

where

PostgreSQL_backup file is the PostgreSQL backup file name

n is a pod ID

Restore Neo4j database

16

Perform Step 18 to Step 24 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 active data center.

17

Go to Step 25.

18

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

19

Enter the following for each cluster member:

```
# kubectl cp local_dir/graph.db namespace/nspos-neo4j-core-dc_name-n:
/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

n is the pod number

20

Enter the following to copy the Neo4j backup file from local storage to the PostgreSQL backup pod:

```
# kubectl cp path/nspos-neo4j_backup_timestamp.tar.gz
nspos-neo4j-core-default-n:tmp/restoreData ↵
```

where

path is the absolute path of the Neo4j backup file

timestamp is the backup creation time

n is a pod ID

21

Transfer the backup file in *local_dir* to the same directory on an NSP host in the other data center.

22

If the NSP cluster is an HA deployment, enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-n --
/var/lib/neo4j/bin/neo4j-admin unbind ↵
```

where

dc_name is the data center name

n is the cluster member number

23

Enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-n --
/var/lib/neo4j/bin/neo4j-admin restore --force --database=graph.db
--from=/tmp/restoreData/graph.db ↵
```

where

dc_name is the data center name

n is the pod number

24

Enter the following for each cluster member:

```
# kubectl exec -it nspos-neo4j-core-dc_name-n -- rm -rf
/tmp/restoreData/graph.db ↵
```

where

dc_name is the data center name

n is the pod number

Start NSP system

25

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 active data center.

1. Open the /opt/nsp/NSP_CN_*R_r*/config/nsp-config.yml file using a plain-text editor such as vi.
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-CN-R_r/bin/nsp-config.bash --undeploy ↵
   ```

5. Enter the following periodically to display the cluster status:

   ```
   # kubectl get pods ↵
   ```

   The pods are listed; the following restore pods are terminated when they are no longer listed:

   - nsp-backup-storage-*n*
   - nspos-neo4j-core-default-*n*

- nspos-postgresql-primary-*n*

Note: You must not proceed to the next step if any of the pods is listed.

Make sure all pods Terminated and no pods listed before proceed to next step

6. Enter the following:

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

   The NSP initializes using the restored data.

7. Enter the following periodically to display the cluster status:

   # **kubectl get pods** ↵

   The cluster is operational when the status of all pods is Running.

26

Close the open console windows.

END OF STEPS

6.5 To back up the NSP databases in a containerized deployment

6.5.1 Purpose

Perform this procedure to manually back up the following NSP databases:

- Neo4j database of IP resource control
- Neo4j database of containerized NSP cluster
- PostgreSQL database of containerized NSP cluster

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 /opt/nsp/NSP_CN_*R_r*/config/nsp-config.yml file with a plain-text editor such as vi, where *R_r* is the NSP release ID.

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:

```
# kubectl create job job --from cronjob/database-backup ↵
```

where

job is a custom name for the manual backup job

database is the database to back up, and is one of the following:

- nsp-tomcat
- nspos-neo4j-core
- nspos-postgresql

The database backup begins.

9

To display the status of a backup job, enter the following:

```
# kubectl get pod pod ↵
```

where *pod* is the database backup pod name

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, copy each backup file from the database backup pod to a location on the local file system, where

directory is the absolute path of a directory on the local file system

timestamp is the backup creation time

1. Enter the following:

```
# kubectl cp default/nsp-backup-storage-0:
tmp/backups/nsp-tomcat/nsp-tomcat_backup_timestamp.tar.gz
directory/ ↲
```

2. Enter the following:

```
# kubectl cp default/nsp-backup-storage-0:
tmp/backups/nspos-neo4j-core/nspos-neo4j-core_backup_timestamp.
tar.gz directory/ ↲
```

3. Enter the following:

```
# kubectl cp default/nsp-backup-storage-0:
tmp/backups/nspos-postgresql/nspos-postgresql_backup_timestamp.
tar.gz directory/ ↲
```

11

Transfer each backup file to a remote, secure location for safekeeping.

12

Enter the following for each backup to delete the backup job:

```
# kubectl delete jobs.batch pod ↲
```

where

where *pod* is the database backup pod name

The backup job is deleted.

13

Close the open console windows.

END OF STEPS

6.6 To restore the NSP databases in a containerized deployment

6.6.1 Purpose

Perform this procedure to restore the following databases in a containerized NSP system:

- Neo4j database of IP resource control
- Neo4j database of containerized NSP cluster
- PostgreSQL database of containerized NSP cluster

6.6.2 Steps

1

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

2

Create a temporary local directory for the backup files; record the directory location for use in subsequent steps.

3

Retrieve the Neo4j database backup from the backup pod; enter the following:

```
# kubectl cp default/nsp-backup-storage-0:tmp/restoreData/nspos-neo4j_backup_timestamp.tar.gz temp_dir ↵
```

where

site is the site that hosts the nspos-neo4j database

temp_dir is the absolute path of the temporary local directory

timestamp is the backup creation time

The Neo4j backup file is copied to the temporary local directory.

4

Retrieve the PostgreSQL database backup from the backup pod; enter the following:

```
# kubectl cp default/nsp-backup-storage-0:
tmp/restoreData/nspos-postgresql/nspos-postgresql_backup_timestamp.
tar.gz temp_dir ↵
```

where

temp_dir is the absolute path of the temporary local directory

timestamp is the backup creation time

The PostgreSQL backup file is copied to the temporary local directory.

5

Retrieve the Tomcat database backup from the backup pod; enter the following:

```
# kubectl cp default/nsp-backup-storage-0:
tmp/restoreData/nsp-tomcat/nsp-tomcat_backup_timestamp.tar.gz temp_dir
↵
```

where

temp_dir is the absolute path of the temporary local directory

timestamp is the backup creation time

The Tomcat backup file is copied to the temporary local directory.

6

If the NSP deployment is HA, perform the following steps.

1. Perform Step 7 to Step 15 in the standby data center.
2. Perform Step 7 to Step 15 in the active data center.
3. Go to Step 17.

7

Log in as the root user on the NSP configurator VM if you are not logged in.

8

Enter the following to enter restore mode:

```
# /opt/nsp/NSP_CN_R_r/bin/nsp-config.bash --restore ↵
```

The system is shut down, except for the database pods, which are brought up in restore mode.

9

Enter the following periodically to display the cluster status:

```
# kubectl get pods ↵
```

The cluster is operational when the status of each pod is Running or Completed.

Note: You must not proceed to the next step until the cluster is operational. If the status of any pod is Error, enter the following to restart that pod:
`# kubectl delete pod pod-name ↵`
Where *pod-name* is the name of the errored pod.

10

Record the pod number of the following pods:

- nspos-neo4j-core
- nsp-tomcat
- nrcx-tomcat

11

If you are restoring the databases on a newly installed NSP cluster, enter the following commands to create the required backup-pod directories:

```
# kubectl exec -it nsp-backup-storage-0 -c nspos-postgresql -- mkdir
/tmp/restoreData/nspos-postgresql ↵
# kubectl exec -it nsp-backup-storage-0 -- mkdir
/tmp/restoreData/nsp-tomcat ↵
# kubectl exec -it nsp-backup-storage-0 -- mkdir
/tmp/restoreData/nspos-neo4j ↵
```

12

Navigate to the temporary local directory that contains the backup files.

13

Restore the PostgreSQL database.

Note: You can use the following command to display the namespaces:
`kubectl get pods -A ↵`

1. Enter the following:

   # kubectl cp ./nspos-postgresql_backup_*timestamp*.tar.gz *namespace*/nspos-postgresql-0:/tmp/restoreData ↵

   where

   timestamp is the backup creation time

   namespace is the namespace in which the PostgreSQL database is deployed

2. Enter the following:

   # kubectl exec -it nspos-postgresql-0 -c nspos-postgresql -- /opt/nsp/os/pgsql/scripts/pg-restore.sh -C -Q -f /tmp/restoreData/nspos-postgresql/nspos-postgresql_backup_*timestamp*.tar.gz ↵

   where *timestamp* is the backup creation time

3. Enter the following:

   # kubectl exec -it nspos-postgresql-0 -c nspos-postgresql -- rm -f /tmp/restoreData/nspos-postgresql_backup_*timestamp*.tar.gz ↵

   where *timestamp* is the backup creation time

14

Restore the Neo4j database.

1. Enter the following:

   # kubectl cp *namespace*/nsp-backup-storage-0:/tmp/backups/nspos-neo4j/nspos-neo4j_backup_*timestamp*.tar.gz *local_dir* ↵

   local_dir is a temporary local directory

   namespace is the namespace in which the nspos-neo4j database is deployed

   where *timestamp* is the backup creation time

2. Enter the following:

   # tar zxf *local_dir* /nspos-neo4j_backup_*timestamp*.tar.gz *local_dir* ↵

3. Enter the following:

   # kubectl cp graph.db *namespace*/nspos-neo4j-core-*DR_site*-*pod#*:/tmp/restoreData ↵

   where

   DR_site is the nsp.dr.dcName value in the /opt/nsp/config/nsp-config.yml file

   namespace is the namespace in which the nspos-neo4j-core database is deployed

   pod# is the pod number

   Note: If NSP is a DR deployment, this substep must also be performed on any standby pods.

4. Enter the following:

```
# kubectl exec -it nspos-neo4j-core-<site_name>-0 -- tar zxf
/tmp/restoreData/nspos-neo4j/nspos-neo4j_backup_<date_time>.tar.gz
-C /tmp/restoreData/nspos-neo4j ↵
```

5. Enter the following:

```
# kubectl exec -it nspos-neo4j-core-<site_name>-0 --
/var/lib/neo4j/bin/neo4j-admin unbind ↵
```

Note: If NSP is a DR deployment, this substep must also be performed on any standby pods.

6. Enter the following:

```
# kubectl exec -it nspos-neo4j-core-<site_name>-0 --
/var/lib/neo4j/bin/neo4j-admin restore --force --database=graph.db
--from=/tmp/restoreData/nspos-neo4j/graph.db ↵
```

Note: If NSP is a DR deployment, this substep must also be performed on any standby pods.

Enter the following:

7. `# kubectl exec -it nspos-neo4j-core-<site_name>-0 -- rm -rf /tmp/restoreData/nspos-neo4j/graph.db ↵`

where

site_name is the name of the pod on which the database restore is being performed

date-time is the date and time that the backup being restored was taken

15

If required, restore the nsp-tomcat database. Execute the following on the standalone or primary pod:

1. Enter the following:

```
# kubectl cp <local-path> <k8s-namespace>/nsp-tomcat-0:
/tmp/restoreData/ nsp-tomcat_backup_timestamp.tar.gz ↵
```

2. Enter the following:

```
# tar zxf <local-path>/ nsp-tomcat_backup_timestamp.tar.gz
<local-path> ↵
```

3. Enter the following:

```
# kubectl cp graph.db <k8s-namespace>/nsp-tomcat-0:/tmp/restoreData
↵
```

NOTE: If NSP was installed in a redundant deployment, this substep must also be performed on any standby pods.

4. Enter the following:

```
# kubectl exec -it nsp-tomcat-0 -- tar zxf
/tmp/restoreData/nsp-tomcat/nsp-tomcat_backup_timestamp.tar.gz -C
/tmp/restoreData/nsp-tomcat ↵
```

5. Enter the following:

```
# kubectl exec -it nsp-tomcat-0 --
/opt/nsp/server/replication/bin/neo4j-admin restore --force
--database=graph.db --from=/tmp/restoreData/nsp-tomcat/graph.db ↵
```

NOTE: If NSP was installed in a redundant deployment, this substep must also be performed on any standby pods.

6. Enter the following:

```
# kubectl exec -it nsp-tomcat-0 -- rm -rf
/tmp/restoreData/nsp-tomcat/graph.db ↵
```

Where *date-time* is the date and time that the backup being restored was taken.

16

Enter the following to start the NSP:

Note: In a DR deployment, you must perform the step in the active data center first.

Note: Ensure that the "deletePvcOnUndeploy" parameter in the nsp-config.yml file is set to false.

```
# bin/nsp-config.bash --undeploy --deploy ↵
```

17

Close the open console windows.

END OF STEPS

6.7 To restore the etcd database

6.7.1 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.

6.7.2 Steps

Obtain and distribute snapshot

1

Log in as the root user on an NSP cluster member in the active 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:

```
# unzip 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 URL of member *N* in the 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