NOKIA

Network Services Platform Release 2.0 R3

Release Description

3HE-11080-AAAC-TQZZA

Issue 1

November 2016

Legal notice

Contents

About this document **4**

1 Release overview **5**

1.1 Release 2.0 at a glance **5**

1.2 Compatibility and support **9**

1.3 References **11**

2 NSP feature descriptions **13**

2.1 NSP platform features **13**

2.2 NSD features **15**

2.3 NRC-P features **24**

2.4 NRC-T features **29**

About this document

Purpose

The *NSP Release Description* lists the features introduced in the NSP Release 2.0. This document is intended to assist network planners and administrators by providing high-level feature descriptions for NSP Release 2.0, along with the schedule for delivery. For more detailed information and procedures related to the features described in this Release Description document, refer to the NSP customer documentation suite delivered with the NSP product.

Document support

Customer documentation and product support URLs:

Customer documentation welcome page

- https://infoproducts.alcatel-lucent.com/cgi-bin/doc_welc.pl

Technical support

- http://support.alcatel-lucent.com

How to comment

Documentation feedback

- Documentation Feedback

1 Release overview

1.1 Release 2.0 at a glance

1.1.1 Introduction

This chapter provides an overview of the feature rollout within the release, as well as platform requirements and licensing and compatibility information. See Chapter 2, “NSP feature descriptions” for more detailed descriptions of the features listed in the 1.1.5 “Feature summary” (p. 6) section. For additional information and procedures related to the features described in this document, refer to the NSP customer documentation suite.

1.1.2 Target schedule

Three releases with new content are planned in the NSP Release 2.0 stream, as summarized below:

- Release 2.0 R1 — June 2016
- Release 2.0 R2 — August 2016
- Release 2.0 R3 — November 2016

1.1.3 Packaging and documentation

The following NSP resources are distributed through OLCS:

- NSP Release 2.0 R3 software installer
- *NSP Installation and User Guide*
- *NSP API Programmer Guide*
- *NSP Planning Guide*
- *NSP API Differences Guide*
- *NSP Release Notice*
- Downloadable, offline representations of the NSP APIs
- Postman collection of NSP API examples

1.1.4 Licensing information

Nokia personnel provides NSD, NRC-P, and NRC-T licenses, which must be requested through the Nokia support organization. Customers must specify which licenses they require, as well as whether or not the licenses are to be used for a standalone NSP, or an NSP in HA mode. A license that is generated for an NSP in HA mode will include three UUIDs. The installation file is provided during installation or can be updated later during a reconfiguration.

1.1.5 Feature summary

The table below lists the features planned for the NSP Release 2.0. See Chapter 2, “NSP feature descriptions” for more detailed feature descriptions.

Release	Key	Summary	Area
2.0 R3	NSP-635	Generic topology NBI REST API	NRC-P
2.0 R3	NSP-656	External application notification — base platform	NSP platform
2.0 R3	NSP-1024	Support service modification in 1350 OMS	NRC-T
2.0 R3	NSP-1106	Support STAR load-balancing algorithm	NRC-P
2.0 R3	NSP-1108	SROS ISIS feature addition	NRC-P
2.0 R3	NSP-1110	Support fir RSVP PCC-initiated LSPs	NRC-P
2.0 R3	NSP-1122	External application notification — object state change	NSP platform
2.0 R3	NSP-1190	Multi-domain L3 VPN service provisioning from L2 endpoints	NSD
2.0 R3	NSP-1217	REST API enhancement: filtering	NSP platform
2.0 R3	NSP-1235	IGP topology discovery via vSROS NSP/PCE Traffic-Engineering Database	NRC-P
2.0 R3	NSP-1237	Support manual resignalling for delegated LSPs	NRC-P
2.0 R3	NSP-1241	Bandwidth management for SR-TE and MPLS LSPs on NRC-P	NRC-P
2.0 R3	NSP-1243	Path disjoint algorithm accounts for SRLG	NRC-P
2.0 R3	NSP-1269	Support of area colors	NRC-P
2.0 R3	NSP-1537	Show link utilization for reserved bandwidth on NSD application	NRC-P

Release	Key	Summary	Area
2.0 R3	NSP-1668	LSP GCO for select LSPs	NRC-P
2.0 R3	NSP-1892	Association between primary and secondary LSP paths	NRC-P
2.0 R3	NSP-1940	Support for IRO object	NRC-P
2.0 R3	NSP-1749	Standalone external topology visualization	NSP platform
2.0 R3	NSP-1897	Support of brownfield LSP and SDP tunnels	NSD
2.0 R3	NSP-2040	LLDP remote peer information	NSD
2.0 R3	NSP-2137	Support of L2 service on 1830 PSS with endpoints on 11QCE12X	NSD
2.0 R3	NSP-2196	Service topology and node type visualization	NRC-T
2.0 R3	NSP-2231	Support of multi-vendor RSVP-TE LSPs on NSD	NSD
2.0 R3	NSP-2237	Support of brownfield service resources	NSD
2.0 R3	NSP-2319	Service CAC at access interface level	NSD
2.0 R3	NSP-2775	Support of non-GMPLS network topologies	NRC-T
2.0 R3	NSP-3061	Support of brownfield E-Line services	NSD
2.0 R2	NSP-326	Base support for assigned roles and licenses	NSP platform
2.0 R2	NSP-441	NRC-T support for 1350 OMS	NRC-T
2.0 R2	NSP-744	L0 GMPLS Integration: SBR and OPSA LSPs via 260SCX2 and 130SNX10 cards	NRC-T
2.0 R2	NSP-1024	Service Consistency: Support service modification in 1350 OMS	NRC-T

Release	Key	Summary	Area
2.0 R2	NSP-1150	Support 260SCX2 card	NRC-T
2.0 R2	NSP-1228	Optical service parameters for the NBI	NRC-T
2.0 R2	NSP-1391	REST API Enhancement: Patch (phase 1)	NSD
2.0 R2	NSP-1400	Support for LAGs as service endpoints	NSD
2.0 R2	NSP-1541	Support for 100G muxponders: 1350 OMS Support	NRC-T
2.0 R2	NSP-1570	Links between routers and 1830 nodes	NSD
2.0 R2	NSP-1739	Link disjoint computation for two OTU4 services	NRC-T
2.0 R2	NSP-1956	1350 OMS HA support	NRC-T
2.0 R2	NSP-2109	Startup service display selection	NSP platform
2.0 R1	NSP-330	HSP high availability	NSP platform
2.0 R1	NSP-583	Secondary LSP path creation	NSD
2.0 R1	NSP-985	Tunnel steering parameters for controlled service path search	NSD
2.0 R1	NSP-1160	Model customization and SAM templates phase 1	NSD
2.0 R1	NSP-1182	QoS Scheduler model phase 1	NSD
2.0 R1	NSP-1213	REST API enhancement: endpoints as dedicated resource	NSD
2.0 R1	NSP-1239	Internationalization NSP	NSP platform
2.0 R1	NSP-1251	NSD support for administrative state	NSD

Release	Key	Summary	Area
2.0 R1	NSP-1255	QoS Policer model	NSD
2.0 R1	NSP-1257	L3VPN extensions	NSD
2.0 R1	NSP-1267	HTTPS encryption of internal and external communication	NSP platform
2.0 R1	NSP-1619	Model enhancements	NSD
2.0 R1	NSP-1779	Service access QoS for 7210 SAS	NSD

1.2 Compatibility and support

1.2.1 Compatible network management products

The NSP must be deployed alongside the following network management products:

- 1350 OMS
- 5620 SAM
- 5650 CPAM
- 7701 CPAA
- vSROS for NSP/PCE
- vSROS for PCC

These components must be deployed with compatible software releases. See the *NSP Release 2.0 R3 Release Notice* for software compatibility information.

1.2.2 Supported network elements

The NSP supports the following nodes/cards:

Node type	Node	Card
IP nodes	7210 SAS	—
	7450 ESS	
	7705 SAR	
	7750 SR	

Node type	Node	Card
Optics nodes	1830 PSS-8 1830 PSS-16 1830 PSS-16 II 1830 PSS-32	130SCX10
		130SNX10
		260SCX2 – 100G
		260SCX2 – OTU4
	1830 PSS-4	—

1.2.3 Supported Optical configurations

The NSP supports the following Optical configurations on the specified cards/nodes:

Card	Node(s) supporting ROADM Directional-Colored configuration	Node(s) supporting Config-D configuration	Node(s) supporting Config-D" configuration	Node(s) supporting Config-D' configuration	Node(s) supporting CDC-F configuration	Nodes(s) supporting GMPLS configuration
11QPA4	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	Yes
11QPEN4	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	Yes
11QPA4 + 130SNX10	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	Yes
11QPEN4 + 130SNX10	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	Yes
11QCE12X	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	N/A
130SCX10	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32
130SCX10 w/OPSA (OCHP)	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32
130SNX10	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32
130SNX10 w/OPSA (OCHP)	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32
260SCX2 – 100G	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32

Card	Node(s) supporting ROADM Directional-Colored configuration	Node(s) supporting Config-D configuration	Node(s) supporting Config-D” configuration	Node(s) supporting Config-D' configuration	Node(s) supporting CDC-F configuration	Nodes(s) supporting GMPLS configuration
260SCX2 – 100G w/OPSA (OCHP)	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	N/A	PSS-8/16/ 16II/32
260SCX2 – OTU4	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32	PSS-8/16/ 16II/32

1.3 References

1.3.1 Referenced documentation

The following documents are referenced:

- *5620 SAM User Guide*, 3HE 10708 AAAC TQZZA 01
- *5620 SAM Planning Guide*, 3HE 10698 AAAC TQZZA 01
- *5650 CPAM User Guide*, 3HE 10840 AAAC TQZZA 01
- *7701 CPAA Setup and Installation Guide*, 3HE 10842 AAAC TQZZA 01
- *NSP Installation and User Guide*, 3HE 11081 AAAC TQZZA 01
- *NSP Release Notice*, 3HE 11084 0003 TQZZA 01

2 NSP feature descriptions

2.1 NSP platform features

2.1.1 [NSP-326] Base support for assigned roles and licenses

Release introduced: 2.0 R2

An NSP instance can have different roles with different functionality. This is governed by licenses. This feature introduces support for NSD and NRC-T licenses. For non-HA deployment, a single UUID is required, whereas an HA deployment requires three UUIDs.

2.1.2 [NSP-330] NSP high availability

Release introduced: 2.0 R1

NSP Release 2.0 R1 introduces a cluster-based high-availability (HA) solution where the NSP and all provided services can survive one server failure, given that three servers are installed.

This solution can also be configured for disaster recovery (DR), provided certain criteria are met. Contact your Nokia support personnel to enable disaster recovery. Switching to the DR site is not yet automated:

- The Primary site has two instances of NSP Core servers, while the DR site has one. All persistent content is automatically replicated among those servers.
- Single server failure in the Primary site is automatically handled and the standby server there will become the active.
- Complete failure on the Primary site (two server failure) will require a manual procedure to bring the DR site acting as the active (and standalone) server.

2.1.3 [NSP-656] External application notification - base platform

Release introduced: 2.0 R3

The NSP provides a base platform for asynchronous event notifications to external applications, such as orchestrators. These notifications are transported using HTTP Server Side Events (SSE) according to the IETF RESTCONF protocol specification. Notifications are defined in the YANG modeling language and encoded in JSON format. This base platform is used by NSP components to realize different types of notifications.

2.1.4 [NSP-1122] External application notification - object state change

Release introduced: 2.0 R3

Clients of the NSP northbound interface receive notifications whenever the state of an NSP-managed object changes. This simplifies synchronization with the NSP, as periodic polling of the REST API is avoided. Notifications are provided for the operational and administrative status of services and endpoints.

2.1.5 [NSP-1217] REST API enhancement: filtering

Release introduced: 2.0 R3

When using the NSP REST API, GET calls can return long lists of objects. This feature introduces support for optional filters in the URI of GET calls, which allow only objects with attributes that match the given expression to be returned. For instance, a filter can be used to query only services of a certain type. The URI syntax supports expressions for many object attributes, and also allows logical combinations of more than one filter criteria.

2.1.6 [NSP-1239] Internationalization NSP

Release introduced: 2.0 R1

This feature enables the NSP to support GUI languages other than the default, English. The translation from English to a desired language can be done after the SW GA.

2.1.7 [NSP-1267] HTTPS encryption for internal and external communication

Release introduced: 2.0 R1

In order to protect the integrity and confidentially of data, the NSP uses HTTPS to secure northbound communications to the REST API and NSP applications. TLS/SSL is also enabled on the southbound communication to the 5620 SAM. In addition, TLS/SSL encrypts the communication between Neo4j databases in a cluster, and all communications with the PostgreSQL database.

Starting with NSP Release 2.0 R1, the NSP installer automatically configures the TLS /SSL infrastructure. Self-signed certificates are created during the installation. On the NSP UI and on the REST API, only HTTPS is supported. HTTP requests will be redirected to HTTPS.

2.1.8 [NSP-1749] Standalone external topology visualization

Release introduced: 2.0 R3

A rendering of the NSD's service topology map can be embedded in any third-party web page that has sufficient access rights. The user has only read-only access to the objects in this map rendering. For example, this means that the position of the objects within the map rendering can be modified, but that these changes are not stored. Third-party web pages embed this map rendering using an HTML 'iframe' tag. The map rendering is then retrieved using a well-defined URI.

2.1.9 [NSP-2109] Startup service display selection

Release introduced: 2.0 R2

A small extension is introduced to the GUI URL that allows users to identify a specific service for visualization upon opening a new browser window/tab.

2.2 NSD features

2.2.1 [NSP-583] Secondary LSP path creation

Release introduced: 2.0 R1

When an LSP is auto-created by the NSP, the NSP creates a primary LSP path using the existing mechanism. In addition, it creates a secondary LSP path with zero bandwidth, referring to a Provisioned Path with no hop information.

The support for secondary path creation is enabled by default in the NSP system. This can be disabled from the Tunnel Creation policy.

2.2.2 [NSP-985] Tunnel steering parameters for controlled service path search

Release introduced: 2.0 R1

Steering Parameters are used to steer the selection of service tunnels for the connections of a given service. All other logic to determine the best selected set of tunnels remain the same. In NSP Release 2.0 R1, MPLS is the only supported tunnel type.

Step 1: Create Steering Parameters by carrier operators (with the appropriate scope of commands and span of control). A steering parameter has a unique string and a unique bit position (0..31). The NSD supports a maximum of 64 parameters.

Step 2: Assign steering parameters to tunnels (also by operators) 'Canadian Tunnels' to tunnel.

Step 3: Create a Tunnel Selection Profile, which has lists of included and excluded steering parameters.

Step 4: When a service is created, a tenant or operator can specify a Tunnel Selection Profile to be used. The NSD Service Connection Orchestrator will then combine the service connection objective, constraints, and the Tunnel Selection Profile in its search for the best possible service tunnels. NSD service path search must only use tunnels whose steering parameter list is a super set of the service's 'included' list and does not have any parameter that is part of the service's 'excluded' list

2.2.3 [NSP-1160] Model customization and SAM templates phase 1

Release introduced: 2.0 R1

The Network Services Director (NSD) uses abstract models to define services. New custom model attributes and scripting support provide flexibility and programmability for the NSD object models.

Custom attributes are a new concept on the REST northbound interface. Additional parameters can be used to extend the NSD models and change the standard provisioning behavior towards devices. In requests to create or modify services or endpoints, the additional customized attributes can be provided as name/value string lists. The NSD passes these parameters transparently to the device mediation layer, which can use these key/value pairs when configuring services and service endpoints. In this release of the NSP, custom attributes are supported for mediation by the 5620 SAM.

In the default setup of the NSP, custom attributes can reference the 5620 SAM object model for services and endpoints. Therefore, a client of the NSP NBI can manipulate service and endpoint parameters in the 5620 SAM mediation layer, even if the objects are not defined in the abstract NSD models. This enables the deployment of highly customized services.

Additional processing logic for custom attributes can be implemented by installing scripts in the 5620 SAM mediation layer. These system scripts can realize even complex provisioning operations, and they can be uploaded on the fly without restarting the 5620 SAM. The combination of custom attributes and scripts is therefore a very powerful method to achieve flexible and customized service provisioning within the NSD. The 5620 SAM scripts can be further fine-tuned by professional service teams for any future needs.

In this NSP Release 2.0 R1, custom attributes are supported on service and endpoint objects for all IP services, including ELINE, ELAN and L3VPN. Custom attributes can be set and modified via the REST API by users with an admin role. Custom attributes can also be persistently configured using NSD service templates, and can then be applied by non-admin tenants. Custom attributes cannot be configured in the NSD application.

2.2.4 [NSP-1182] QoS scheduler model phase 1

Release introduced: 2.0 R1

This feature extends the NSP QoS model to support multi-tier scheduler policies on endpoints. The NSP leverages the Generic Queue Profiles (GQP), which are enhanced by ingress and egress scheduler policies. These GQP scheduler policies have to be configured in the 5620 SAM.

When an endpoint is configured in the NSP, the NSP will trigger the 5620 SAM mediation layer, and the 5620 SAM will use scheduler policies associated to a GQP to provision the scheduler configuration to the node. This feature is supported on ELINE, ELAN, and L3 VPN on nodes that support scheduler policies.

This feature allows the NSP to configure aggregate CIR and PIR parameters for endpoints in order to enforce CIR and PIR limits for all traffic passing a scheduler. The aggregate CIR and PIR parameters can override the 5620 SAM GQP scheduler policy per direction (ingress and/or egress). Aggregate CIR and PIR can be configured both in the NSP NBI or in the NSP GUI.

Additionally, specification of absolute rate and relative percentages is supported. In this release, the conversation from percent to absolute rate is done inside the NSP. Percentage values are calculated by the rate of the port. Setting of percent values is not supported on the GUI in NSP Release 2.0 R1.

2.2.5 [NSP-1190] Multi-domain L3 VPN service provisioning from L2 endpoints

Release introduced: 2.0 R3

Multi-domain L3 VPN services from L2 metro areas are supported. These services are created between PE routers on metro areas, however, because some PE routers are not L3 capable, the NSD performs the path search across the network, from L2 metro areas to L3 core, and finds the best exiting routers from metro to core. Then, the NSP provisions L2 E-Line services on all metro areas and L3 VPN services in the core. Finally, the services are stitched together by the NSP using VLAN hand-off.

The intra-domain tunnels must be created in advance, and all metro domains are interconnected via Ethernet links (VLAN handoff) to the core. Since none of the routers on L2 metro domains are L3 VPN capable, the NSP uses this property to run the path search algorithm. This property can be set from the NSP REST API.

The NSD uses L2 and L3 service templates to define the common attributes for the autocreated services. Profiles are used for QoS and the auto-assignment of L3 RD/RT. The NSP also uses the tunnel selection profile to include and exclude specific tunnels during path search. The path search objectives (such as minimizing hop or cost) and other values specific to the VPN (such as the IP addresses of the L3 access points) are defined either from the NSD application or the NSP REST API. The NSP uses the QoS CIR values to book the bandwidth on tunnels.

2.2.6 [NSP-1213] REST API enhancement: endpoints as dedicated resource

Release introduced: 2.0 R1

This feature extends the REST API. Endpoints of services are modeled as separate resources. As a result, there are new individual REST API calls to add, update, and delete endpoints. This feature is supported on services with multiple endpoints (L3VPN, ELAN). It is also possible to create a service first without any endpoint definitions. Endpoints can then be added via future modifications.

2.2.7 [NSP-1251] Support for administrative state

Release introduced: 2.0 R1

The REST API and the UI of the Network Services Director (NSD) allow for the explicit configuration of the administrative state for services and endpoints. This feature is currently supported for L3VPN, ELAN, and ELINE services. By default, services and endpoints have an administrative state of "up". If a template redefines the value as "down", services and endpoints will need to be explicitly administratively enabled after service provisioning is completed.

2.2.8 [NSP-1255] QoS Policer model

Release introduced: 2.0 R1

This feature adds support for QoS policers to the NSP service endpoint QoS model. Policer schedulers, which are also known as arbiters, are an alternative to queues, and can be configured on ingress and/or egress. The NSP leverages a corresponding enhancement in the 5620 SAM Generic Queue Profile (GQP), which can now distinguish between queues and policers. With the GQP extension, the NSP can now use GQP profiles with policer configuration during service creation, either directly or via QoS endpoint templates. When the NSP creates a service endpoint, the corresponding policer configuration will be created on the nodes by the 5620 SAM mediation layer. The web GUI and the REST API has been extended, with some functionality only being available on the REST API.

The NSP can also configure aggregate PIR parameters if a service endpoint will use a QoS policier. The aggregate PIR parameter will then be provisioned to the node by overriding the policer control policy on that service endpoint. In NSP Release 2.0 R1, configuration of QoS policers is supported on 7x50 and 7210 nodes.

2.2.9 [NSP-1257] L3VPN extensions

Release introduced: 2.0 R1

This feature adds the following additional parameters to the L3VPN model within the NSP:

1. Service MTU: When this parameter is configured for an L3VPN service, its value will be applied to all existing and new endpoints in the service. The parameter can be configured via the NSP web GUI, via service templates, or via the REST API.
2. Loopback address: A user can add a 'loopback' endpoint with a /32 IP address to an L3VPN. This allows the user to set the router ID of that L3VPN site.
3. Static route preference: With this parameter, a user can modify the preference of a static route. The default preference is 5 and the range is 0..255.

2.2.10 [NSP-1391] REST API Enhancement: Patch (phase 1)

Release introduced: 2.0 R2

The REST Patch command allows for incremental updates to be applied to one or more of the basic attributes of an IP service (E-Line/E-LAN/L3 VPN) and/or their corresponding endpoint(s). Any attribute that can be configured using an IP service /endpoint REST command can also be updated using the REST Patch command. For complex attributes (such as siteServiceQosProfile in L2EndpointRequest, routeTargets, routingBgp, routingStatic, and secondaryAddresses in L3EndpointRequest) the REST Patch command will replace the attribute's previously-provisioned value with the new value supplied by the user. There are no restrictions to the number of attributes that can be part of the REST Patch payload.

2.2.11 [NSP-1400] Support for LAGs as service endpoints

Release introduced: 2.0 R2

The NSD supports LAGs as service endpoints. Users can select a LAG port as an endpoint in the GUI or REST API, and all functions applicable to endpoints (add, delete, modify) are also applicable to LAGs.

2.2.12 [NSP-1570] Links between routers and 1830 nodes

Release introduced: 2.0 R2

This feature addresses the manual creation of physical links between two nodes when there is no dynamic LLDP to discover the link. The NSP creates a new NBI API to create a physical link between two physical ports of two nodes. For example, this can be used to create physical links in the NSP between an 1830 port and an SR port. This entity will only be stored in the NSP database and nothing will be provisioned to the 5620 SAM or the 1350 OMS. The NSP will also provide an NBI API to delete such manually-created links.

This can be done on the GUI by right-clicking on the network element in the map (physical map). This provides an option to delete the link. These links are not pushed

down to the 1350 OMS or the 5620 SAM.

The Operational state of the link will be determined by the operational state of one (or both) of these ports. This might require updates to the link when the operational state of the port changes.

2.2.13 [NSP-1619] Model enhancements for ports and NEs

Release introduced: 2.0 R1

This is a small extension of the data models used by the REST API. First, the REST API now returns the port type for ports, using the IANA definitions in RFC 7224. For Ethernet ports, additional encapsulation information is also returned. Second, the REST API can now be queried to obtain basic information about network elements (NEs) that are managed and controlled by the NSP. The NSP returns NE platform information according to the YANG model in RFC 7317.

2.2.14 [NSP-1779] Service access QoS for 7210 SAS

Release introduced: 2.0 R1

This feature allows to configure templates and profiles for access QoS with CIR, PIR, MBS and CBS on 7210 SAS nodes. The services supported in this feature include ELAN (VPLS), ELINE (EPIPE), and L3 VPN (VPRN). An NSP user can create Generic QoS Profiles (GQP) in the 5620 SAM and apply the profiles to Endpoint QoS templates and Service templates within the NSP. The templates can then be applied when creating a service. CIR, PIR, MBS and CBS values from the template can also be overridden.

This feature provides the access QoS functionality on 7210 SAS nodes. On 7210 SAS nodes, QoS enforcement is limited to ingress traffic. The change does not affect NSP users.

2.2.15 [NSP-1897] Support of brownfield LSP and SDP Tunnels

Release introduced: 2.0 R3

The NSD is capable of discovering LSP and SDP tunnels created previously within the 5620 SAM, including multi-vendor LSP and SDP tunnels, with the following exceptions:

- A single SDP tunnel using multiple LSPs
- Multiple SDP tunnels using the same LSP
- SR (Segmented Routing) type LSPs

The NSP will discover, and will allow users to create services with bandwidth constraints on service tunnels created previously within the 5620 SAM. The NSP will operate with initial allocated bandwidth on these tunnels and will keep track of used bandwidth for all the services created by the NSP. It is assumed that the NSP is the only entity creating

services on these tunnels. The NSP cannot be used to delete, resize the allocated bandwidth, or modify the LSPs associated with service tunnels previously created within the 5620 SAM.

The NSP will discover, and will allow users to create service tunnels on RSVP-TE LSPs created previously within the 5620 SAM. The NSP will operate with initial allocated bandwidth on these LSPs and will keep track of used bandwidth for all the service tunnels created on these LSPs by the NSP. It is assumed that NSP is the only entity creating service tunnels on these LSPs. The NSP cannot be used to delete, resize the allocated bandwidth, or modify LSPs previously created within the 5620 SAM.

When the reserved bandwidth of a previously-discovered LSP is modified, the NSP receives an event, and will update the both initial and available bandwidth on the LSP and SDP tunnel. This case should apply to all LSP and SDP tunnels managed by the NSD, regardless of their origin, with the following exceptions:

- A single SDP tunnel that uses multiple LSPs
- Multiple SDP tunnels that use the same LSP
- SR (Segmented Routing) LSPs

When an LSP is used by an SDP tunnel, but is not yet bound to any service, that LSP's initial and available bandwidth will be updated. However, since the LSP is used by an SDP tunnel, the SDP tunnel will take the entire bandwidth. As there are no services using the SDP tunnel, the available bandwidth should be equal to current bandwidth.

When an LSP is used by an SDP tunnel and there are services bound to the SDP tunnel, the SDP tunnel will take all of the LSP's current bandwidth. The LSP's available bandwidth should be 0, and depending on the services bound to the SDP tunnel, the available bandwidth of the SDP tunnel will be adjusted to reflect the current bandwidth, minus the total bandwidth of all services running on that SDP tunnel.

2.2.16 [NSP-2040] LLDP remote peer information

Release introduced: 2.0 R3

The NSP's REST interface provides LLDP remote peer information for router ports that have LLDP enabled. This allows a client of the REST interface to retrieve the LLDP status and LLDP remote peer information as far as it is known to routers managed by the NSP.

2.2.17 [NSP-2137] Support of L2 services on PSS 1830 nodes with endpoints on 11QCE12X

Release introduced: 2.0 R3

This feature adds support for E-Line Optical Ethernet services between two 1830 PSS nodes when the access side is on 11QCE12X Carrier Ethernet cards. The tunnel

between the 1830 PSS nodes will be L0 "OCh Service". During the L2 service provisioning, the NSD will assume the existence of the L0 tunnels. The NSD will neither resize the existing trails, nor create a new one. Therefore, the underlay L0 service should exist.

During NSP E-Line service creation, the user will select the client ports (C ports on the card) that are to be used for the service. The NSP will then select the best network ports (X ports on the card) and links between nodes.

2.2.18 [NSP-2231] NSD support of multi-vendor RSVP-TE LSPs

Release introduced: 2.0 R3

This feature extends the support of RSVP-TE LSP tunnels to multi-vendor nodes. The NSP supports the following multi-vendor endpoint combinations for E-Line services:

- Cisco-Nokia
- Juniper-Nokia
- Juniper-Nokia
- Cisco-Juniper
- Cisco-Cisco
- Juniper-Juniper

Cisco LSP names must be in the format of Tunnel<*number*>, where <*number*> is an integer between 0 and 65535. Standby paths are not supported by Cisco or Juniper, only secondary paths. Therefore, in instances where Cisco or Juniper endpoints are used and the Tunnel Creation Template has the Protection Type set to Standby, secondary paths will be created instead. Cisco LSP-Path Bindings contain a property called Path Option. This property will be set to 1 for primary and 2 for secondary.

When creating an E-Line service on multi-vendor nodes, the NSP will attempt to find a tunnel based on the criteria specified in the Tunnel Selection Profile (TSP). If no tunnel exists, and the TSP specifies that new tunnels should be created, the NSP will create MPLS RSVP-TE tunnels, including the Dynamic LSP and LSP-Path Bindings.

E-LAN services can be created on Cisco nodes. When this is done, the NSP configures a property called bridgeDomainId during site creation. E-LAN services are not supported on Juniper nodes.

For L3VPN services, the NSP supports the RSVP-TE option, since multi-vendor nodes do not support SDP tunnels. As a result, if an L3 VPN service is created on a multivendor node, the NSPs algorithm will try to find or create RSVP-TE tunnels and always set the auto-bind property to RSVP-TE on the multi-vendor nodes.

2.2.19 [NSP-2237] Support of brownfield service resources

Release introduced: 2.0 R3

The Global Cache enables the NSP to track resources being used by the network, including the resources of services that originate from the 5620 SAM. In order for the NSP to discover such services, they must have their “NSD-managed” flag enabled within the 5620 SAM. Once this is done, the usage of VLAN IDs, L3 VPN Route Distinguishers (RD), and L3 VPN Route Targets (RT) can be tracked across 5620 SAM-NSD managed networks. When the NSD requests one of these resources, the Global Cache verifies their availability before assignment. Only freed resources are considered available for usage. All services created using the NSD will be validated for resource usage, and therefore will not infringe upon the resources of existing services.

2.2.20 [NSP-2319] Service CAC at access interface level

Release introduced: 2.0 R3

The NSD can perform bandwidth CAC and validation on access ports. Every port available for use in E-Line, E-LAN, and L3 VPN services will have their available ingress bandwidth and available egress bandwidth displayed as read-only properties in the NSD application and the NSP's REST APIs. When any of these ports are discovered, available bandwidth is initialized to port speed. In some cases, such as the 60-port 10 /100 card when the port is operationally down, the port speed is zero. On fixed port speed cards, the port speed is populated, allowing services with bandwidth to be configured even when the port is down. Service CAC is not available on the variable-speed SFP-based cards.

Any changes to port speed will be reflected in the displayed available ingress bandwidth and available egress bandwidth. This may result in these fields displaying a negative value. No alarms or notifications will occur but a WARN level log will be generated.

A formula is used to calculate both the ingress and egress aggregate bandwidth of all endpoints used by E-Line, ELAN, and L3 VPN services. The formula yields the sum of the CIR values, which is based on each of the configured queues and the scheduler policy of the QoS. This same value is used for E-Line service tunnel bandwidth calculation. No overbooking is applied to the formula. When the NSP creates a service on one of these endpoints, the validation code will make sure that the sum of the formula is less than, or equal to, the current available bandwidth on the port, otherwise the service will not be created and an error is returned.

The bandwidth is only booked after the traversal operation is run to match with the current behavior of the core bandwidth. It is possible that between the validation check and the traversal operation, the port bandwidth was consumed by another service. In this case, the OLC state is changed to Routing Failed, and the user is told that either the access port ingress or egress bandwidth was exceeded. Modifying the CIR will reinitiate the traversal operation. Similar operations occur when adding endpoints to an existing service and modifying endpoints. In the latter case, it is the bandwidth delta which is applied to the available ingress or egress bandwidth. Upon deletion of an endpoint or service, the available ingress or egress bandwidth is increased by the bandwidth of the endpoints.

Service CAC is available on both access and hybrid ports. If there are network interfaces on hybrid ports, these are not tracked as part of the available ingress or egress bandwidth. When an upgrade is performed, the available ingress and egress bandwidths will be calculated based on all existing services within the NSD. This may result in negative values. When in an overbooked state, any request that will not cause a change to bandwidth reservation, or that will cause a shrink in bandwidth reservation, will be permitted.

2.2.21 [NSP-3061] Support of brownfield E-Line services

Release introduced: 2.0 R3

E-Line services created within the 5620 SAM can be managed by the NSP. In order for the NSP to discover these services, their "NSD-managed" flag must be enabled within the 5620 SAM. Once discovered by the NSP, these services will function the same as E-Line services created within the NSP itself, provided that they meet the NSD requirements. Any change made to these services within 5620 SAM after discovery will be propagated to the NSD, provided the change impacts the topology of the service.

2.3 NRC-P features

2.3.1 [NSP-635] Generic topology NBI REST API

Release introduced: 2.0 R3

The NRC-P exposes a TE-based topology to any interested applications. This topology is discovered via IGP or BGP-LS. Each TE topology is defined within a domain/network. These networks should not have any interconnections. Therefore, a node or link within a network must only exist in that network. The NRC-P can export all such domains /networks, as well as details such as network topology and details of specific network nodes/links. The TE topology is exported as JSON data structures and is modeled in the NRC-P based on the following IETF yang specifications:

- https://tools.ietf.org/html/draft-ietf-i2rs-yang-network-topo-04
- https://tools.ietf.org/html/draft-ietf-teas-yang-te-04
- https://tools.ietf.org/html/draft-ietf-teas-yang-te-topo-05

The JSON objects that are returned correspond to the objects that are defined within the yang modules in the above drafts. However, the dashes (-) are removed from the object names, and the first letter of each word is capitalized. All objects are identified by their UUID. The following JSON data structures are returned:

- Networks - A List of networks that are each an OSPF or ISIS topology within a BGP-LS domain.
- Network - A network containing a Network-ID (UUID), a list of nodes, a list of links, as well as network attributes.

- Node - A node contains a Node-ID (UUID), a list of termination points, as well as node attributes. A node corresponds to an OSPF or ISIS router or subnet.
- Termination Point - A termination point contains a TP-ID (UUID), as well as termination point attributes. A termination point corresponds to an OSPF or ISIS interface.
- Link - A link contains a Link-ID (UUID), a Source Node-ID, a Source TP-ID, a Destination Node-ID, a Destination TP-ID, as well as link attributes. A link corresponds to an OSPF or ISIS adjacency.

There are five NSP REST API commands that return the above data structures.

- GET: /api/v3/ietf/te/networks
- GET: /api/v3/ietf/te/network/ {networkId}
- GET: /api/v3/ietf/te/node/ {nodeId}
- GET: /api/v3/ietf/te/termination-point/ {tpId}
- GET: /api/v3/ietf/te/link/ {linkId}

2.3.2 [NSP-1106] Support of STAR load-balancing algorithm

Release introduced: 2.0 R3

The NRC-P provides a load-balancing and optimal-path-placement algorithm, known as the STAR algorithm. This algorithm uses an internal metric, calculated from the current value of the TE bandwidth reservation, to route the CSPF paths. Every path that is allocated on a TE link changes the internal metric for both the link and the overall path. Initially, all links have the same star weight, or metric, so the first path requests for CSPF traversal will choose the shortest path that satisfies all constraints. If there are multiple paths that satisfy the user constraints, then a path will be chosen randomly. This behavior is the same for normal CSPF.

Subsequent requests will choose paths that possess the least star weight, thereby ignoring the path that the normal CSPF algorithm would have chosen. The calculation of the star weight is based on a formula that uses the current link reservation. The user constraints are still satisfied. This balances the overall network utilization.

The STAR algorithm is invoked per LSP by associating that LSP to a path profile. The path profile template is defined in the NRC-P and requires setting the objective to use STAR WEIGHT. The path profile is specified with the LSP definition and is conveyed to the NRC-P via a PCE request message.

2.3.3 [NSP-1108] SROS ISIS feature addition

Release introduced: 2.0 R3

The NRC-P constructs an inter-area IP topology based on ISIS TE. This topology is obtained via ISIS or BGP-LS. The NSP SROS VM agent connects to the individual areas via ISIS instances. These instances are defined on the agent because, without defining the instances, the NRC-P cannot discern between multiple ISIS L1 areas (all ISIS L1

areas appear as one large L1 area). The NRC-P requires that each local L1 instance on the agent connect to an ISIS L1 area. The NRC-P also requires the configuration of the IGP identifier to be the same for each L1 instance.

2.3.4 [NSP-1110] Support RSVP PCC-initiated LSPs

Release introduced: 2.0 R3

The NRC-P supports RSVP TE PCC-initiated LSPs and provides paths for both primary and secondary standby path requests. The NRC-P also maintains the bandwidth of the LSP and routes the its paths when they are delegated by the PCC.

The NRC-P maintains the active path in case both the primary and secondary paths are signaled, and also when the primary path is down. The NRC-P also maintains the shared explicit behavior when the primary and secondary paths share common link resources.

The NRC-P also indicates the active path between the primary and secondary pair.

2.3.5 [NSP-1235] IGP topology discovery via vSROS NSP/PCE Traffic-Engineering Database

Release introduced: 2.0 R3

The NRC-P supports using the vSROS NSP/PCE Traffic-Engineering Database to import an IGP topology, in addition to importing it natively. The IGP domain is defined in the vSROS NSP/PCE Traffic-Engineering Database as a set of nodes in which each node is assigned a unique IGP representation using a combination of Area-ID, Router-ID, Protocol, Topology-ID, and Instance ID.

2.3.6 [NSP-1237] Support manual resignalling for delegated LSPs

Release introduced: 2.0 R3

The NRC-P supports the ability to select multiple LSPs and resignal the LSPs. The LSPs are resignaled in the order selected. This applies to both RSVP and SR TE LSPs when they are delegated to the NRC-P. LSPs not delegated to the NRC-P will not be resignaled.

The resignal applies to a single LSP, however, if that LSP is part of both a profile and a path group, and these contain other LSPs, then the profile definition and network state will determine whether some, or all of the LSPs in that set are resignaled/rerouted. If the LSP is associated to a path profile that has the disjoint attribute (regardless of the type), then the resignal may cause the disjoint algorithm to run on all the LSPs in the set (common profile and group), provided that there are existing disjoint LSPs. If there are

no disjoint LSPs in the disjoint group, or the profile doesn't have the disjoint attribute set but has the bidirectional requirement, then both the directional LSP and its bidirectional pair are rerouted.

2.3.7 [NSP-1241] Bandwidth management for SR-TE and MPLS LSPs

Release introduced: 2.0 R3

The NRC-P manages the LSP bandwidth consumption on the TE links for both stateless and stateful PCC configurations. In a stateless configuration, the NRC-P receives TE updates from the network as LSPs are signaled, thereby mimicking the TE DB bandwidth consumption on the nodes. This allows for accurate LSP path computation without maintaining state on the NRC-P. In a stateful case, wherein the reports are sent to the NRC-P from the PCC, the bandwidth is again communicated by the PCC to the NRC-P via the bandwidth object. Here, the NRC-P will reconcile the TE update with the specific LSP bandwidth update via the report. Therefore, the NRC-P maintains full LSP state along with the consumption on the TE links for these LSPs only.

It is possible that existing brownfield LSPs will not request paths from the NRC-P, and therefore, will have no state on the NRC-P. The NRC-P will not show these LSP reservations on the TE links. For a mixture of LSPs that are PCE-reported and non-PCEreported, the NRC-P will track and show the actual TE consumption on a TE link in addition to the LSP reservation for PCE-reported LSPs.

A bandwidth value that is specified on an LSP has no significance on the PCC/router because the SR TE does not maintain any state on the intermediate or destination routers. Therefore, no bandwidth tracking is done in the local TE DB. The bandwidth has to be tracked by the NRC-P if the LSP is configured to report bandwidth. Bandwidth tracking on the NRC-P is done only after a valid PCE report message is generated by the PCC. The NRC-P tracks the bandwidth reservation for SR TE LSPs separate from RSVP TE LSPs.

A loose hop SR LSP whose bandwidth is specified and computed locally will not be tracked by the NRC-P, even with the PCE report option enabled. The NRC-P only tracks SR TE LSP paths computed by the NRC-P itself.

2.3.8 [NSP-1243] Path disjoint algorithm needs to account for SRLG

Release introduced: 2.0 R3

The specification for SRLG applies over the node and link disjoint specification. The algorithm finds the suitable disjoint paths and then selects which of those are also SRLG disjoint.

2.3.9 [NSP-1269] Support of area colors

Release introduced: 2.0 R3

Area links discovered in the NRC-P are displayed in a variety of colors, which can be modified by the user.

2.3.10 [NSP-1537] Show link utilization for reserved bandwidth in NSD application

Release introduced: 2.0 R3

The NSD application displays the TE link reservation on IP links in a separate table. When TE LSPs are admitted. removed. or re-routed in the network, the TE consumption of the IP links changes in response. The IP link consumption is shown with the highest utilized links at the top.

2.3.11 [NSP-1668] LSP GCO for select LSPs

Release introduced: 2.0 R3

The NRC-P also supports optimizing the paths of existing LSPs by applying an optimization algorithm. This algorithm extracts the current resource availability on the current topology and re-routes the selected LSP paths such that the overall network consumption is minimized. The result is to utilize more network links, but also reduce the consumption on the links. LSPs must be delegated to the NRC-P and must be preselected. Profiles do not have to be associated to the paths in order to use this algorithm. The LSPs to be optimized are selected manually on from the NSD application.

LSPs that have a profile with the disjoint option enabled are excluded.

2.3.12 [NSP-1892] Association between primary and secondary LSP paths

Release introduced: 2.0 R3

The NRC-P maintains the active path in case both the primary and secondary paths are signaled, and also when the primary path is down. The NRC-P also maintains the shared explicit behavior when the primary and secondary paths share common link resources.

The NRC-P also indicates the active path between the primary and secondary pair.

2.3.13 [NSP-1940] Support of IRO object

Release introduced: 2.0 R3

The NRC-P supports the IRO object specification within a PCC request. The NRC-P computes a CSPF path from the source to the IRO object, and another CSFP path from the IRO object to the destination. If the second CSPF path visits any of the nodes in first CSPF path, the path computation fails.

When used with a path profile that contains the bidirectional disjoint specification, a forward LSP and its matching reverse LSP must share the same IRO configuration. This means that the list of addresses in the IRO path must be the same, but their order reversed. This is because the disjoint algorithm is natively bidirectional strict. If the reverse LSP contained IROs that did not exist in the forward path, no path found would be found, because it would no longer be bidirectional strict.

2.4 NRC-T features

2.4.1 [NSP-441] NRC-T support for 1350 OMS

Release introduced: 2.0 R2

This feature allows the NSP to work with the 1350 OMS, with the 1350 OMS acting as the southbound mediation layer for the NSP. The interface between the NSP and the 1350 OMS allows for the following:

- discovery of existing services in the network through the 1350 OMS and service consistency
- discovery of the supported network elements and TE links from the 1350 OMS
- passing requests for service configuration on the NE, including connectivity setup according to service parameters specified in the NBI: unprotected, protected with OPRC, restorable, and route diverse
- allowing for notifications and re-synchronization between the 1350 OMS and the NSP, specifically for notification of failure to be propagated from the 1350 OMS to the NSP
- allowing for the NSP to retrieve information about the current/nominal and backup path in the network, especially after failure conditions and GMPLS-led restoration

2.4.2 [NSP-744] L0 GMPLS Integration: SBR and OPSA LSPs via 260SCX2 and 130SNX10 cards

Release introduced: 2.0 R2

This feature allows for operation of the NSP together with L0 GMPLS. The supported configurations and NEs are PSS23 ROADM with config D/D'/D", and 130SNX10 and 260SCX2 cards. The requirements for services are unprotected, protected with OPSA, and restorable with SBR. In particular, PRC is supported with OPSA cards for protection and SBR for restoration. After a restoration event where PRC is activated, the new current/nominal and backup paths for the LSP are retrieved and are displayed in the GUI for the associated service, hence the display always shows the current and backup path of the service.

2.4.3 [NSP-1024] Service Consistency: Support service modification in 1350 OMS

Release introduced: 2.0 R2

Provides full consistency between services deployed by the NSP and the 1350 OMS. To achieve consistency in support of network troubleshooting, NSP-created services should be visible and identifiable in the 1350 OMS. In addition, all services (in any layer) created using the 1350 OMS should be correctly uploaded into the NSP.

2.4.4 [NSP-1150] Support 260SCX2 card

Release introduced: 2.0 R2

This feature allows for operation of the 200G flexible rate muxponder 260SCX2 card in PSS32 colorless, directionless ROADM in configuration D/D'/D". The card and its capabilities are discovered through the 1350 OMS and they are displayed in the GUI. Point and click allows for establishment of ODU4 services on the client side, and it triggers path computation and wavelength establishment on the line side. The service can be set up according to service requirements: unprotected, protected through PRC (with OPSA cards for protection and SBR for restoration), and route diverse. Note that the card is used only in 100G mode on the line side, 200G mode is not supported.

2.4.5 [NSP-1228] Optical service parameters for the NBI

Release introduced: 2.0 R2

This feature allows a user to specify the parameters of the requested service in the NBI. The parameters are: unprotected, protected through OPSA, restoration with SBR, and restoration service with GR. In addition, reversion of the service can be manual- or auto-reversion. There is also the possibility to specify the port type on the client port.

2.4.6 [NSP-1541] Support for 100G muxponders: 1350 OMS Support

Release introduced: 2.0 R2

This feature allows the NRC-T to setup a service on the 130SNX10 card. The card is discovered by the NRC-T through the 1350 OMS, including the card capabilities, and is made available on the NSP GUI. Point and click operation allows for creation of an ODU2 service from client port to client port, muxing client signals into 100G line and triggering automatic computation of a path and wavelength assignment. The service creation is subject to service attributes: unprotected, protected through PRC (with OPSA cards for protection and SBR for restoration), and route diverse. This feature is supported on PSS32 shelf with colorless, directionless ROADM with configurations D/D'/D".

2.4.7 [NSP-1739] Link disjoint computation for two OTU4 services

Release introduced: 2.0 R2

This feature allows the user to specify a service name and endpoints for a service, as well as to request that the service be link disjoint to an already existing service represented by its name. This is usually referred to as LSP diverse path computation.

2.4.8 [NSP-1956] 1350 OMS HA support

Release introduced: 2.0 R2

This feature allows the NSP to work with the 1350 OMS in a high availability scenario. This allows for:

- detection of the primary 1350 OMS and the secondary 1350 OMS
- identification of the active 1350 OMS
- detection of a 1350 OMS failure, followed by a switch to the active 1350 OMS
- handling of 1350 OMS switchover times

2.4.9 [NSP-2196] Service topology and node type visualization

Release introduced: 2.0 R3

This feature adds visualizations of optical services to the NSD application's multi-layer topology map. The physical network topology, consisting of multiple nodes and the links between them, is extracted from the underlying physical network via the NSP's REST API. 1830 PSS ROADM nodes and ILA nodes are distinguished from one another using different graphical icons. The network topology can be exported via the NSP's REST API for integration with hierarchical controllers/orchestrators.

2.4.10 [NSP-2775] NRC-T support of non-GMPLS network topologies

Release introduced: 2.0 R3

NRC-T support of optical services on 260SCX2, 130SNX10, and 130SCX10 cards of 1830 PSS nodes is extended to Managed Plane (non-GMPLS) network configurations.