Path control and optimization in segment-routed networks

Using the Nokia Network Services Platform as an external controller for traffic engineering

Application note

Abstract

Cloud services and 5G are driving interest in path control and optimization. Cloud providers have created new standards for service velocity and flexibility. The introduction of 5G is creating new business opportunities and service differentiation with new use cases that require low latency and highly reliable services.

At the same time, network operators need to stay profitable. There is a continuing focus on minimizing OPEX through higher operational agility, responsiveness and efficiency, which can be achieved through automation and operational simplification. In addition, CAPEX can also be lowered through a more efficient use of network resources.

This document examines how the Nokia Network Services Platform (NSP) path control and optimization helps operators achieve these objectives.

Contents

Networking technology advancements enabling evolution

High performance and reliability plus OPEX and CAPEX savings objectives have been important to network operators since long before cloud services and 5G. Beyond these business drivers, the recent interest in path control and optimization solutions is enabled by the maturation of several key technologies:

- Segment routing, which is a much more scalable traffic engineering solution than Resource Reservation Protocol – Traffic Engineering (RSVP-TE)
- Real-time telemetry, which enables the path control engine to react to changes in utilization and latency in near-real time
- The standardization of control protocols, which enables multivendor interoperability support.

The benefits of traffic engineering have long been recognized, but until recently the only traffic engineering solution was RSVP-TE. Unfortunately, RSVP-TE does not scale very well. Every label switched path (LSP) needs to be signaled, and every core router that traffic transits maintains the state of every LSP. Consequently, operators use RSVP-TE, but they use it to create large pipes.

There is broad consensus in the industry that segment routing is the future (see Figure 1). It scales much better than RSVP-TE because core routers do not maintain state and there is no need for signaling protocols to establish paths through the networks, such as happens when using RSVP or Label Distribution Protocol (LDP).

Instead, the path of the network is captured in the label stack that is added to a packet at the ingress of the network. The result is the increased network scalability that enables LSP creation on a per-application basis.

Figure 1. Segment routing principles and benefits

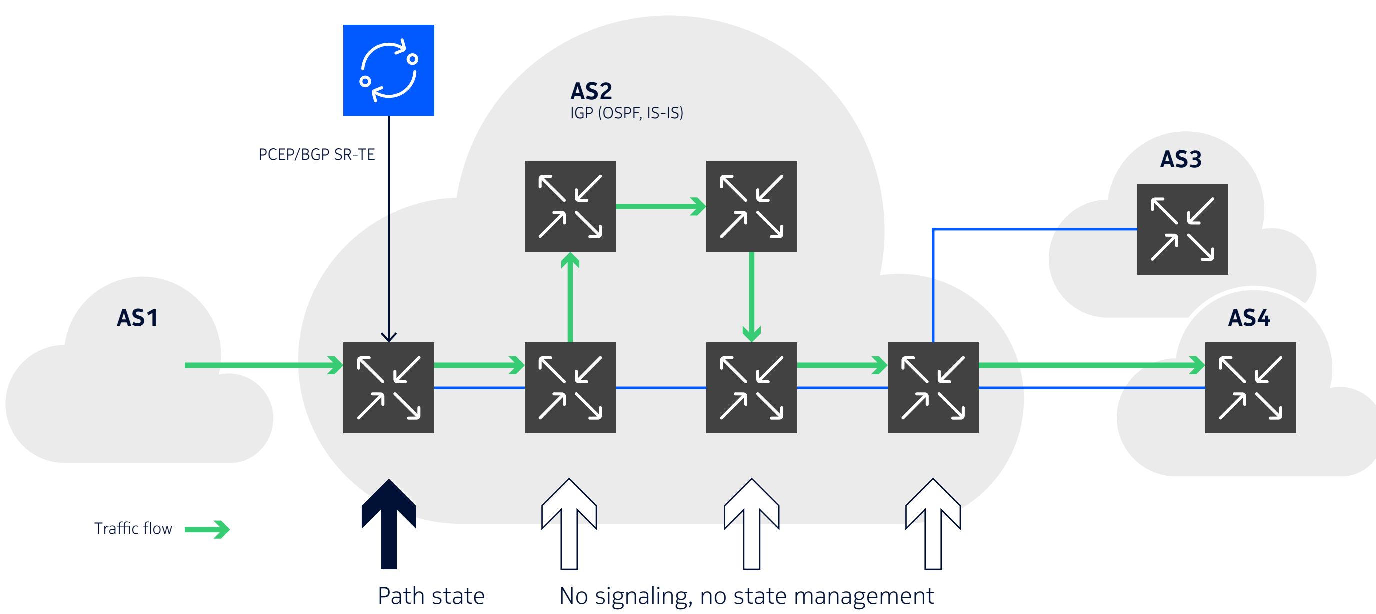

New technologies such as real-time telemetry (see Figure 2) have also enabled network monitoring to detect issues as they happen and trigger optimization in near-real time to:

- Minimize latency
- Avoid congestion and packet loss
- Route around network failures.

This optimization allows the network to meet stricter performance and SLA guarantees under continuously changing network conditions.

Figure 2. Real-time telemetry principles and benefits

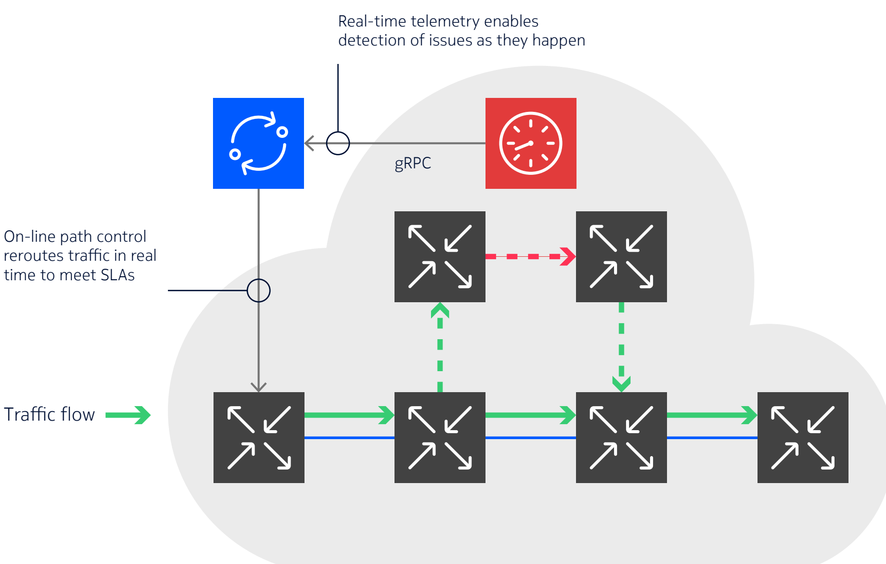

Real-time telemetry relies on new protocols such as gRPC Remote Procedure Call (gRPC). Using gRPC, routers stream performance data to a controller in near-real time. Generally, they send updates every few seconds.

In the past, management systems used Simple Network Management Protocol (SNMP) to retrieve statistics every 5 or 15 minutes. This had two shortcomings:

- Performance data was aggregated over a longer period, resulting in a limited view of what was happening in the network because sudden traffic spikes were filtered out.
- To the extent that the aggregate data provided actionable intelligence, it arrived late. The management system would detect congestion and inform operations support systems (OSS) minutes after the fact.

By contrast, with real-time telemetry a controller can detect congestion or changes in latency within seconds.

The benefits of a controller to complement segment routing

In principle, segment routing can be deployed without controllers. Routers generally support a certain level of path selection intelligence, often referred to as constraint-based shortest path first (CSPF).

Furthermore, the IETF has defined the IGP Flexible Algorithm (FlexAlgo), which extends router capabilities to support latency-based optimization or path selection based on admin tags.

FlexAlgo is a useful tool, but without an additional controller acting as an external Segment Routing – Path Computation Element (SR-PCE), its functionality is limited.

Automation using a controller to reroute paths (see Figure 3) provides:

- Bandwidth guarantees and utilization-based optimization
- Distribution of traffic to achieve higher network utilization.

These functions cannot be achieved using FlexAlgo through a distributed set of routers.

Figure 3. Complementing segment routing with a controller

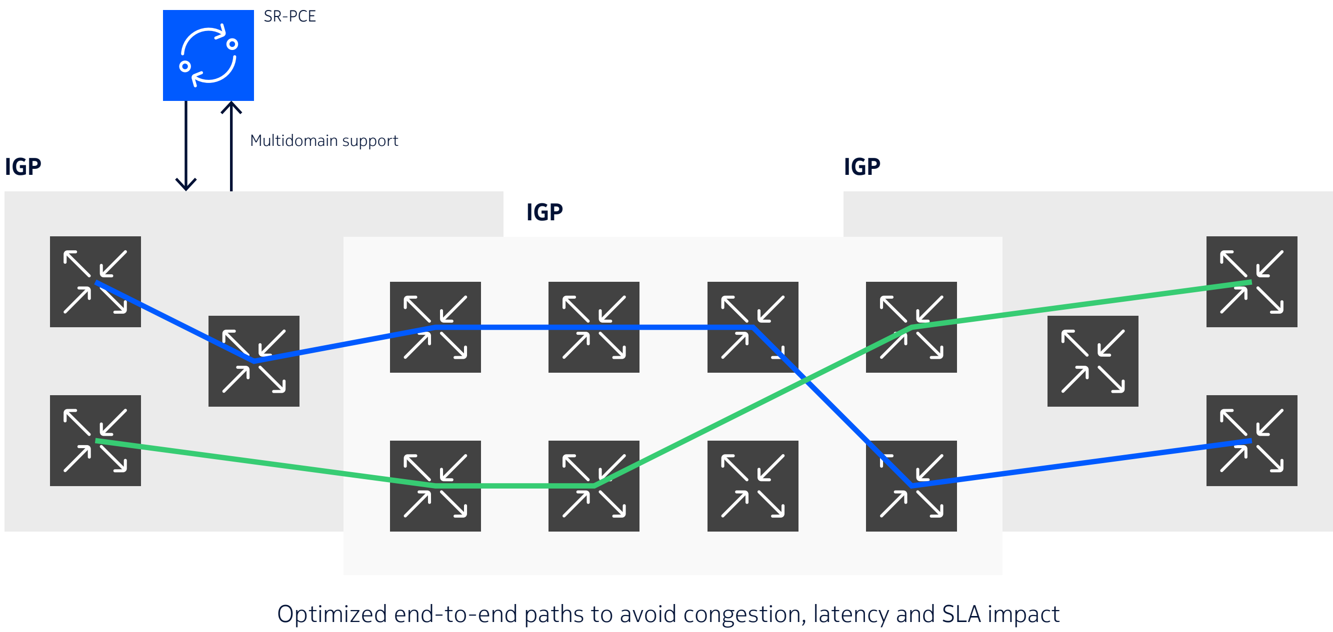

A router-based solution can route based on latency and admin tags, but any consideration based on link utilization (congestion avoidance, SLA guarantees, or optimal network utilization) requires a controller.

Without an external controller there is no mechanism to:

- Continuously optimize path placement and achieve greater network utilization
- Trigger a full automated network re-optimization for all paths (or a subset of paths) for global concurrent optimization (GCO).

A controller provides:

- **Visibility** into the network topology and LSPs to enable graphic presentation of a topology
- **Improved SLA adherence** through path routing (or rerouting) based on bandwidth guarantees or routing-based actual utilization
- **Informed decisions** on which LSPs to reroute when congestion occurs and where to reroute them
- **Traffic optimization** to run the network more efficiently.

Network operators can also trigger automation for path re-optimization and to reroute paths to prevent issues resulting from maintenance actions.

Within a single domain, a router-based solution provides only limited path diversity support. A router can establish two diverse LSPs that start and terminate in the same provider equipment. However, a router-based solution cannot guarantee diversity of LSPs that start and end on different provider equipment unless the operator uses admin groups to create independent planes (and this process has its own limitations).

An external controller simplifies supporting path diversity and enhances resiliency support. It also allows operators to simulate how LSPs would be routed before they are created.

Without a controller there is no way to control paths for traffic engineering in a multidomain network. With an external controller, paths can be optimized end to end within single domains or across multiple domains to avoid congestion, latency and SLA impact.

Introducing a controller can seem daunting for operators because they are used to a network where routers act autonomously, and this provides a certain level of resiliency and fault tolerance. However, operators can introduce an external controller in a gradual manner (see Figure 4).

For example, they can start by deploying an external controller just to visualize what is happening in the network; this provides many operational benefits.

After operators get comfortable with the controller, they can start to use it to control LSPs, to calculate paths and reroute when necessary. This can be done gradually by starting with a subset of LSPs and using the controller to manage just those LSPs until operators are ready over time to increase the number of LSPs that are managed by the controller.

Figure 4. Gradual introduction of path control and optimization

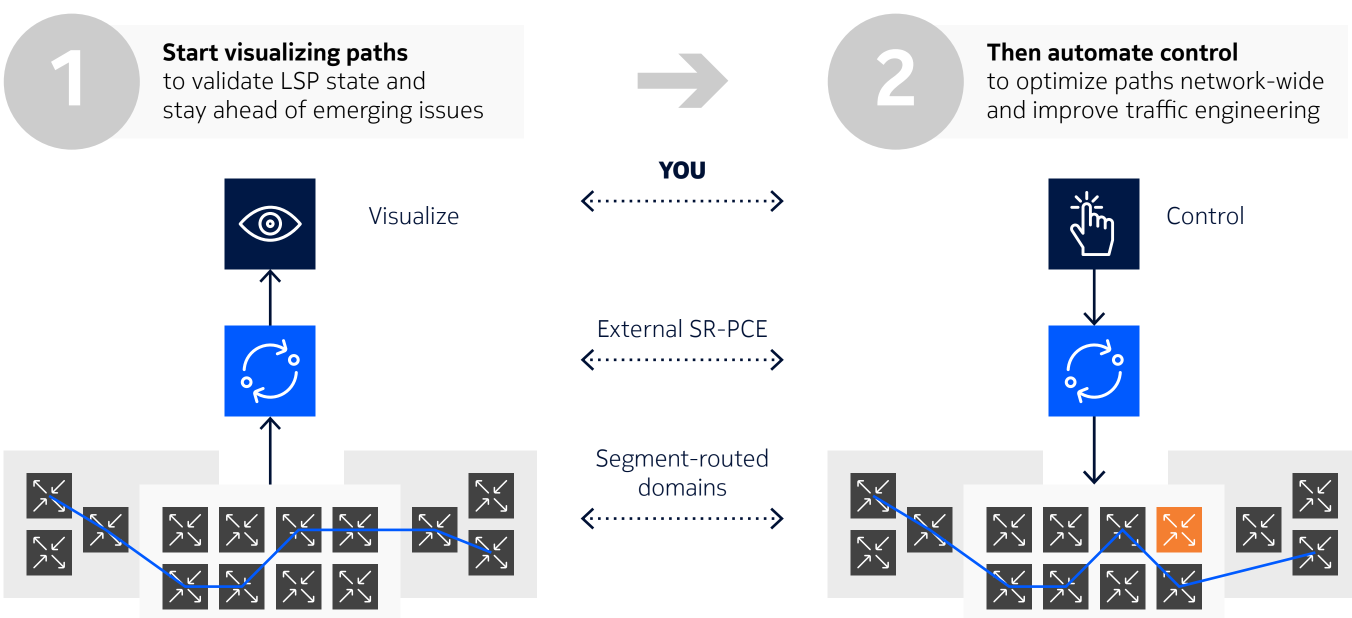

When operators are confident that the SR-PCE's automated control capabilities, stability and maturity are proven on selected LSPs in a network by using representative traffic, they can transition to a more network-wide deployment to optimize and improve traffic engineering for end-to-end paths in single domains and also across multiple network domains.

Path control and optimization with the Nokia NSP

Nokia's solution for path control and optimization is the Network Services Platform (NSP). This functionality is part of the NSP and natively integrates with all the other capabilities that the NSP provides for the configuration, provisioning and assurance of the network.

The NSP was designed to control traffic in real time with high performance (see Figure 5) through efficient and scalable path selection algorithms. This real-time and high-performance traffic control is especially important when network failures occur because hundreds of LSPs may need to be rerouted in the shortest time possible.

Figure 5. Path control and optimization using the Nokia NSP

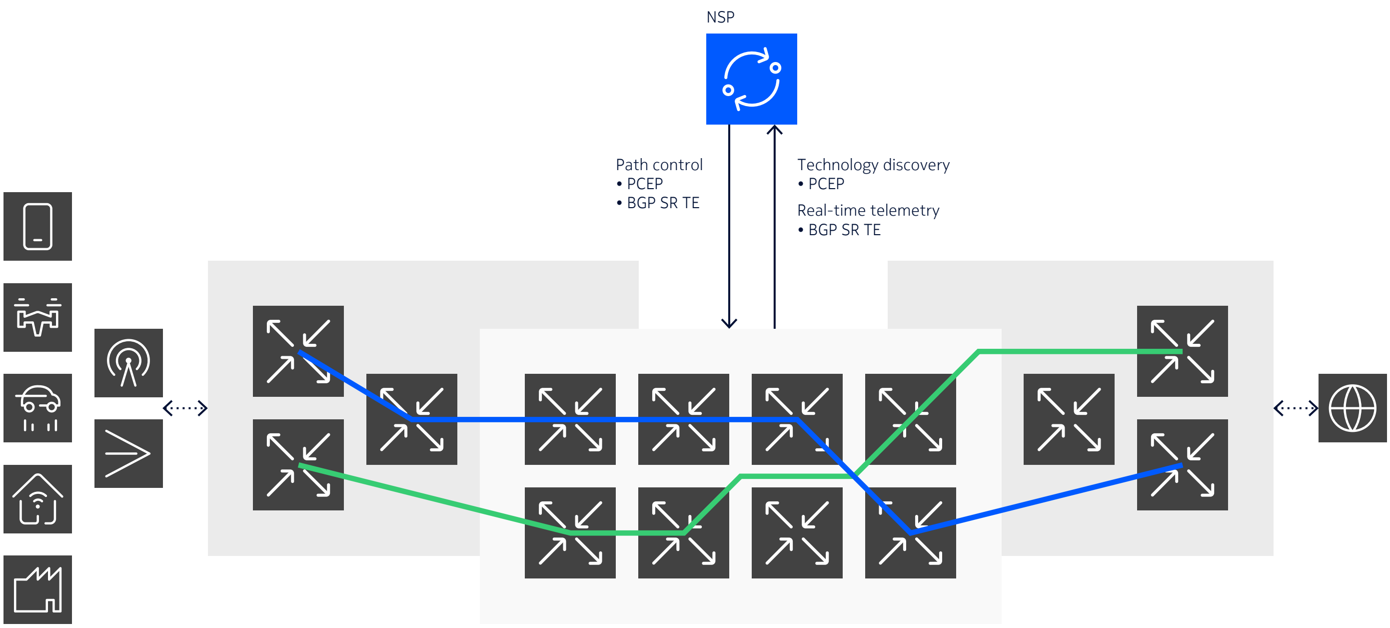

By contrast, some competitors' solutions are derived from off-line simulation tools that are more than a decade old, and a significant difference can be seen in design and performance compared to the NSP.

The NSP acts as a PCE as defined by the IETF. As such, it provides the full set of features one would expect from a PCE for path control and optimization.

Tailor-made for real-time, high-performance network control, the NSP provides a rich feature set, including:

- RSVP-TE and segment routing support
- PCC- (Path Computation Client) and PCE-initiated LSP creation
- Multi-instance, multidomain support
- Multivendor support
- Closed-loop optimization based on real-time utilization and latency
- Simulation and capacity planning.

In addition, the NSP supports several innovative use cases that go beyond anything competitive solutions provide, as discussed in the "Operator case studies" section.

Operator case studies

Network operators from Asia Pacific (APAC) to Europe and North America are using the Nokia NSP for every-thing from 5G automation and network slicing to traffic optimization and reliable interconnection services.

5G automation at a large operator in APAC

A large network operator in APAC with approximately 6 million subscribers needed to enable the rollout of 5G services. The operator required a new 5G backhaul network with both IP and optical equipment, including thousands of routers.

The network needed to be scalable and to provide telemetry-based optimization with latency and utilization triggers.

The NSP delivers telemetry-based optimization with latency and utilization threshold triggers. The NSP can support up to 120,000 LSPs and delegation of up to 50,000 LSPs by acting as a PCE for path control and traffic engineering. When failures occur, the NSP can reroute 50 LSPs per second on average.

The NSP enhances path control with multi-layer coordination, providing advantages:

- Multilayer topology discovery
- Increased resiliency (by ensuring that LSPs which are supposed to be disjoint do not run through the same fiber)
- CAPEX savings by minimizing spare capacity across the IP and optical layers
- OPEX savings through faster and more accurate troubleshooting.

Enhancing path control with multilayer coordination is one example of a broader NSP benefit that competitive products do not provide. The NSP is a holistic platform that features seamless interworking between its applications.

Traffic optimization at a large operator in Europe

A large fixed and mobile network operator in Europe with approximately 14 million subscribers is building a new network. One of the operator's objectives is to use the network as efficiently as possible by limiting the total network capacity and running links hotter (at more than 50 percent).

The operator realized that it is vulnerable when link failures occur. Without a controller, when a high-capacity link fails, traffic could be rerouted over lower-capacity links, which would become overloaded.

The NSP's utilization-based optimization mitigates this risk. It monitors utilization on every link and LSP under its management. It then can provide congestion avoidance under all circumstances.

When traffic reaches a certain threshold, the NSP reroutes one or more LSPs to less utilized links to mitigate congestion. This functionality is especially useful under failure conditions, when traffic spikes may occur.

The NSP's ability to monitor link utilization in real time and its ability to reroute hundreds of LSPs in seconds minimizes connectivity interruptions.

The NSP does not simply reroute LSPs based on bandwidth considerations. Simultaneously, it takes latency requirements into account. Multiconstraint, telemetry-based optimization reroutes traffic to avoid congestion while still meeting latency constraints.

The NSP even provides this capability for partial loose-hop LSPs. The use of partial loose-hop LSPs is a way to overcome the fact that routers can put only a limited number of labels on a packet; this is referred to as the Maximum Stack Depth (MSD).

Segment routing MSDs

One complication in a segment-routed network is that there are limitations to the size of the label stack that routers can put on packets. For example, it is possible that 10 labels are required to specify a route through the network while the ingress router can add only 8 labels to the packet.

The NSP supports two solutions to solve this problem:

- The ability to use partial loose-hop routes, where the NSP identifies a subset of segment identifiers (SIDs). This, in combination with IGP routing, is enough to direct traffic along the selected path.
- The use of Binding SIDs.

If the number of labels required to explicitly define the end-to-end path is too large, the NSP identifies a subset of transit routers on the path and relies on shortest-path forwarding within the network to get from transit point to transit point. This obviously complicates the path monitoring and path selection algorithms. Yet, even with partial loose-hop LSPs, the NSP can use telemetry data to meet multiple constraints at the same time.

To maximize the ability to spread traffic evenly through the network, the NSP automatically splits LSPs that are becoming too large. This increases traffic granularity and minimizes the potential of congestion.

Periodically, the operator will need to increase network capacity to deal with traffic growth. The operator uses the NSP for capacity planning. The NSP monitors utilization on each link and on each LSP.

Through its analytics capabilities, the NSP identifies growth trends. Network topology, LSPs and utilization data are fed into the NSP's path control simulator, which the operator can use to perform a worst-case failure analysis. The operator can then understand utilization trends over time on various LSPs as well as estimate when the network will be running out of capacity and exactly where.

The operator decided to start NSP deployment with RSVP-TE but intends to move to segment routing as soon as possible. This is something we see often. Some operators start with segment routing but many others start with RSVP-TE, with the intention to move to SR soon.

Reliable interconnection services at a major international operator

An international network operator has been providing circuit-switched services to its customers. The operator is introducing segment routing in its network and wants to benefit from the flexibility that SR brings while maintaining the service characteristics and SLAs it has been offering.

The operator wanted to use their IP network's ability to load balance traffic across multiple equal-cost multi-path (ECMP) routing paths. The NSP perfectly meets the operator's needs by introducing the concept of TE ECMP, a bundle of multiple ECMP paths that are treated as if they form a single LSP.

With these multi-strand LSPs, the NSP can provide the resiliency options requested by the operator. Of special interest is the "dual diverse" option, which consists of a pair of LSPs that must follow diverse routes.

The diversity starts when traffic leaves the point of presence (POP). Within the POP, the two LSPs are then allowed to traverse the same spine switches. To the operator, load balancing across multiple paths was more important than strict, end-to-end path diversity.

Another requirement was that in the dual diverse option, the first LSP that was established would follow the shortest path. To keep the end-user experience as consistent as possible, this ordering also needed to be maintained after node failures and subsequent restoration.

A simplistic approach to meet the operator's requirements would rely heavily on the use of admin groups. For example, by including or excluding certain admin groups when setting up LSPs, dual diversity can be achieved.

The problems with this approach are:

- It requires operator involvement whenever a new service is established
- If the network grows and new switches are being installed, existing path definitions may need to be modified.

Instead, the NSP solution fully automates the path selection so there is no need of operator intervention to steer traffic. The NSP relies solely on its understanding of the IP topology and meets the operator's requirements algorithmically. This enables deterministic and predicable behavior where the primary path in a set will always be routing over the shortest path, including during failure conditions and after subsequent restoration.

In addition, the NSP provides blackholing prevention. In a multi-strand LSP, certain failures could result in blackholing of traffic. The NSP constantly monitors the network and when such failures occur, it redefines the LSP (e.g., by replacing an anycast SID by a single node SID to avoid sending traffic into a black hole.)

This solution answers the specific needs of this operator, but many aspects will be useful for other operators as well. For example:

- Leaf-spine architectures are being used more often throughout the network. TE ECMP has universal applicability.
- To achieve path diversity, the NSP takes only relevant network segments into account. In this case, it ignores the fact that within the data center the paths are not necessarily diverse. A similar solution could be useful in mobile backhaul networks or other access networks, where the first mile may be a single link and where diversity is possible only after the first aggregation router.

Summary

Cloud services and 5G have created demand for automation and 5G network slicing as well as dynamic interconnection and traffic engineering. As a result, momentum for controllers is growing because operators are seeing the advantages that can be achieved by path control and optimization, including:

- Improved performance
- Increased resiliency
- CAPEX savings
- Operational efficiency

The Nokia NSP delivers a path and optimization solution that is purpose-built for real-time, online deployments. The NSP delivers innovative features that are unique in the industry and it integrates seamlessly with a broader scope of NSP functions.

In addition, path selection and LSP creation can be deployed as integral parts of service fulfillment.

The telemetry data collected for closed-loop path control can also be captured in reports and can be analyzed through the real-time NSP analytics engine for trending and baselining. The output of the analytics engine can in turn be used by the PCE to drive optimization actions.

Interesting use cases can be created by complementing PCE actions with other commands in scripts executed by NSP's workflow manager.

The NSP is built for real-time, online deployments with:

- Scalability, high performance and reliability
- Telemetry-based optimization
- Visualization
- Multivendor support.

The NSP is unique in the industry through:

- Multiconstraint optimization
- Traffic engineered ECMP

The NSP is a holistic platform that allows for seamless extension of path control functionality with other features:

- Service management
- Service assurance
- IP/optical coordination
- Real-time analytics.

Learn more

To learn more about the Nokia NSP:

Visit the web page

Read the brochure

Read more case studies:

- LG U+ and Nokia NSP
- SDN Communications and NSP
- Automating IP services delivery over a multivendor network NSP

Abbreviations

AS	Autonomous System
BGP	Border Gateway Protocol
BGP-LS	Border Gateway Protocol – Link State
CAPEX	capital expenditures
EANTC	European Advanced Networking Test Center
gRPC	gRPC Remote Procedure Call
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IP-VPN	Internet Protocol – Virtual Private Network
IS-IS	Intermediate System to Intermediate System (protocol)
LDP	Link Discovery Protocol
LSP	label switched path
OSPF	Open Shortest Path First
PCEP	Path Computation Element Communication Protocol
PNF	physical network function
RSVP-TE	Resource Reservation Protocol – Traffic Engineering
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SR-LSP	Segment Routing – Link State Protocol
SR-PCE	Segment Routing – Path Computation Element
SR-TE	Segment Routing – Traffic Engineering
VXLAN	Virtual eXtensible local area network

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: 398237 (February) CID210433