Segment routing on Nokia routing platforms

Tools and applications for IP network architects

Application note

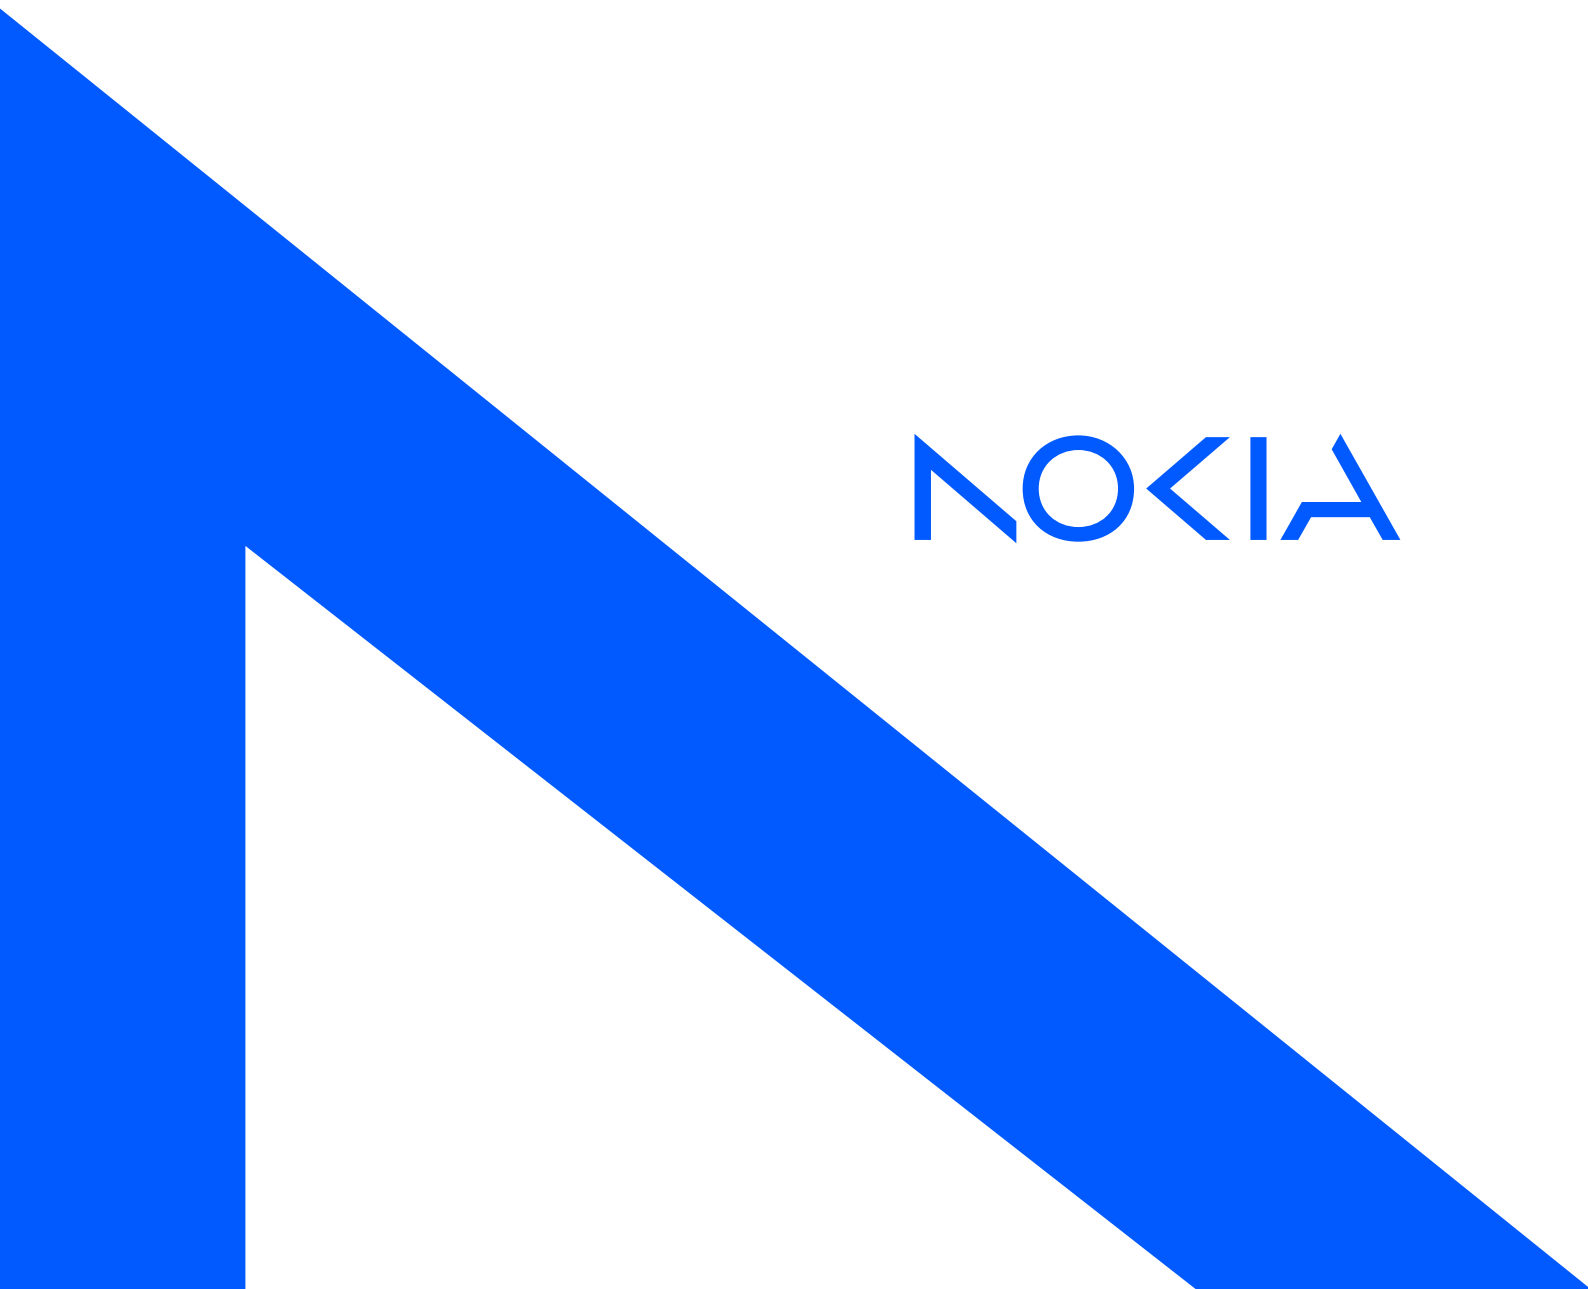

Contents

Introduction and overview

Segment routing (SR) is a powerful technology for protecting and engineering network traffic. This application note gives a high-level overview of the operation and capabilities of SR with insights into the suite of tools and applications supported by Nokia routing platforms to help network operators take full advantage of these capabilities.

SR is a packet steering technology that offers a scalable approach for establishing predefined forwarding paths in the IP network that override the default shortest path, while meeting specific constraints like available bandwidth, latency, protection or physical diversity. SR steers packets by encoding Segment Identifiers (SIDs) in their headers. SIDs contain the packet-processing instructions for each intermediate and destination router, which greatly reduces the need for a control plane to instantiate and maintain path state in the network, while simplifying network operations and reducing resource requirements. As a result, SR provides better scalability than Resource Reservation Protocol - Traffic Engineering (RSVP-TE) or Label Distribution Protocol (LDP).

There are two types of SID encoding:

- Segment Routing MPLS (SR-MPLS) – encodes a 32-bit SID and programs it as an MPLS label to provide a tunnel to an IPv4 or IPv6 destination.
- Segment Routing IPv6 (SRv6) – encodes a 128-bit SID and programs it as an IPv6 address to provide a tunnel to an IPv6 destination.

SR-MPLS combines all the proven attributes of MPLS protocols into a single framework. These attributes include an efficient label switching data plane, shortest path and source routing, source-routed fast reroute (FRR) protection path, traffic engineering, and bandwidth efficiency. SR-MPLS represents a mature and field-hardened choice for many brownfield and greenfield IPv4 and IPv6 transport networks. SRv6 is however emerging in edge telco cloud data center use cases as well as new IPv6 backbone use cases and provides a framework for the programmability of IPv6 networks that takes advantage of the large IPv6 address space. The interworking of SRv6 with SR-MPLS domains is therefore a crucial capability. Nokia's implementation of SR provides a comprehensive SR-MPLS to SRv6 interworking capability, deployed both in public interoperability tests and in real-world production environments.

Segment routing – traffic engineering (SR-TE)

SR is an ideal technology to engineer forwarding paths with granular policy constraints—for example, nodes and links to include or exclude in the path, physical diversity, administrative state, and path metrics such as available bandwidth, cumulative latency and maximum number of hops. SR takes a source-based routing approach, which only requires the ingress or headend router to maintain policy and state information about the path. In its simplest form, a segment route is a sequence of segments that must be traversed when forwarding packets along a constrained path that meets a given policy (Figure 1).

Figure 1. SR basic principles

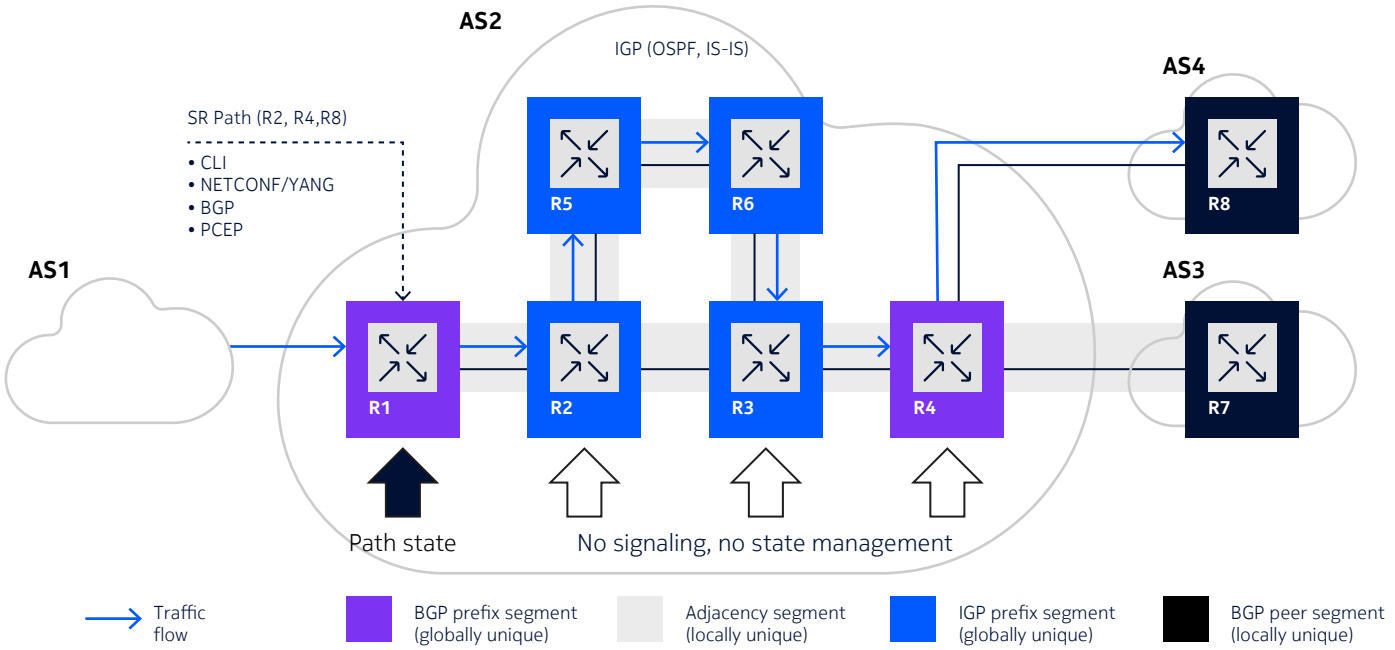

Traffic-engineered paths for segment routes can be conveyed to a headend router in various ways, such as model-driven CLI, NETCONF/YANG, Border Gateway Protocol (BGP) updates or the Path Computation Element Protocol (PCEP – RFC 5440). Segments can refer to different types of network objects:

- Adjacency segments refer to numbered router interfaces and are locally unique to a router.
- Adjacency sets refer to groups of adjacency segments between routers.
- IGP prefix segments refer to a subnet or routing node and are globally unique.
- BGP prefix segments refer to border nodes such as data center gateways.
- BGP peer segments refer to an outgoing peering interface or node, which allows traffic to be steered to a particular egress point.
- Anycast segment is a node prefix that is advertised by more than one node. The set of nodes advertising the same anycast-SID form a group called an anycast set.

SR supports multiple data plane options to facilitate network portability and evolution:

- SR for MPLS (SR-MPLS) represents segments as MPLS labels that are encoded in an MPLS label stack (RFC 8663).
- SR-MPLS over User Datagram Protocol (UDP). MPLS traffic is encapsulated over UDP (RFC 7510). This option is useful for IP-only data center fabrics that do not natively support MPLS.
- SR for IPv6 (SRv6) represents each segment as an IPv6 address construct encoded in an IPv6 segment routing header (SRH) extension (RFC 8754).

When deployed over an MPLS data plane, the segment identifiers (SIDs) are allocated from a reserved block in the MPLS label space (in the case of dynamic adjacency SIDs, allocation is from the dynamic range). To advertise segment reachability information within IPv4 or IPv6 routing domains, SR-MPLS and SRv6 can use interior gateway protocol (IGP) extensions for OSPF (RFC 8665) and IS-IS (RFC 8667). This enables a straight-forward migration that preserves all transport capabilities of LDP and RSVP-TE, while improving failure recovery. BGP is used to advertise BGP prefix and peer segments, and to advertise segment reachability information in IP fabrics without an IGP, such as data center fabrics.

SR policies and SR-TE tunnels

Segment routing supports traffic engineering using SR-TE label switched paths (LSPs) or SR Policies. Both provide a logical traffic tunnel that adheres to specific forwarding constraints. From a user perspective, SR-TE LSPs are similar to traditional traffic-engineered LSPs such as RSVP-TE and offer a natural migration path. They support primary and secondary paths with strict or loose hops that can be initiated by a router or by an external Path Computation Element (PCE) controller. Path level equal cost multi-path (ECMP) is achieved with multiple SR-TE LSPs, where each LSP independently operates active/standby path protection. The association of LSPs with an intent/color is performed via admin-tag policy.

SR Policies consist of multiple candidate paths, of which one is selected as active at a given time. Each candidate path has multiple traffic-engineered segment lists, which are all programmed for the active candidate path and across which packets are sprayed using equal cost multi-path (ECMP) load-balancing. The active/standby mechanism is achieved by switching among candidate paths.

Sub-address family extensions to BGP allow dynamic instantiations and advertisement of SR policies in an intra-domain or inter-domain context. These extensions allow unique identification of an SR policy by means of an abstract color (a 4-byte number carried in signaling messages to indicate an intent for traffic flow treatment), an end point address and a route distinguisher. Each SR Policy can be assigned a unique binding SID (BSID) for the purpose of steering traffic into the policy.

SR-TE with SDN, whether using SR-TE LSPs or SR Policies, allows dynamic provisioning of traffic paths on behalf of service requests and tracking these paths in a centralized Traffic Engineering Database. The BSID of an SR-TE LSP or an SR policy instantiated at a network boundary point can be recursively used in the segment list of an end-to-end SR-TE path across multiple network domains. BSID allows the compression of the segment list of the end-to-end path and hides the changes of the mid-point intra-domain path segments from the ingress PE.

Centralized path computation does not face the resource constraints of distributed routing algorithms and can potentially optimize resource use from a global network perspective. It also avoids collisions, re-tries, and packing problems that have been observed in networks using distributed TE path calculation, where head-ends make autonomous decisions.

However, where distributed path computation meets a service provider's operational objectives, IGP Flexible Algorithms (Flex-Algo) complement SR-TE by adding new prefix segments with specific optimization objectives and constraints that can be advertised by supporting router nodes. Pre-defined algorithms and metrics are referred to by an agreed-upon 7-bit identifier, where algorithms 0-127 are standard algorithms, such as the shortest path algorithm, and 128-255 are operator defined. Applicable Flex-Algo link metrics include the default IGP metrics, a minimum unicast link delay or the TE metric, and links to be excluded or included based on their administrative group link affinity/color or SRLG membership.

Local and end-to-end path protection and restoration

Besides engineering a working traffic path, segment routing can also be applied to pre-calculate backup paths that protect against failure of the primary path. Since intermediate routers do not have state information about the path, autonomous recovery mechanisms such as MPLS fast reroute will not work properly because a point of local repair (PLR) that is doing a local detour will not be able to identify the downstream nodes of the SR-TE path.

Therefore, where fast reroute is used in an SR context, an alternate next-hop path must be pre-computed so that when a failure is detected with the primary next hop, the alternate can rapidly be used until an SPF algorithm is run and a new primary next hop is installed. When a PLR attempts to pre-compute an alternate backup next hop, the backup next hop is generically called a loop-free alternate or LFA (RFC 5286). The existence of a suitable LFA—and therefore the percentage of fast reroute coverage that a given network can obtain, depends on the topology. Basic LFA calculations have good coverage for meshed topologies, but generally perform poorly in ring topologies, in which case they may create micro loops (see RFC 6571).

Remote LFA (RLFA – RFC 7490) and directed (DLFA) extend the basic LFA repair mechanism by extending the topology coverage to nearly 100 percent. If a link cannot be protected for a given destination with local, adjacent LFA neighbors, RLFA attempts to create a virtual LFA by using a tunnel to carry packets to a point in the network where they will not be looped back. However, when an RLFA repair tunnel uses RSVP or LDP, targeted LDP sessions must be tunneled to repair tunnel endpoints so that inner labels can be exchanged. Since the repair points can dynamically change along with the topology, the targeted LDP sessions must also be dynamically set up and torn down, but many operators do not favor such dynamic behavior.

Figure 2. Traffic protection and restoration

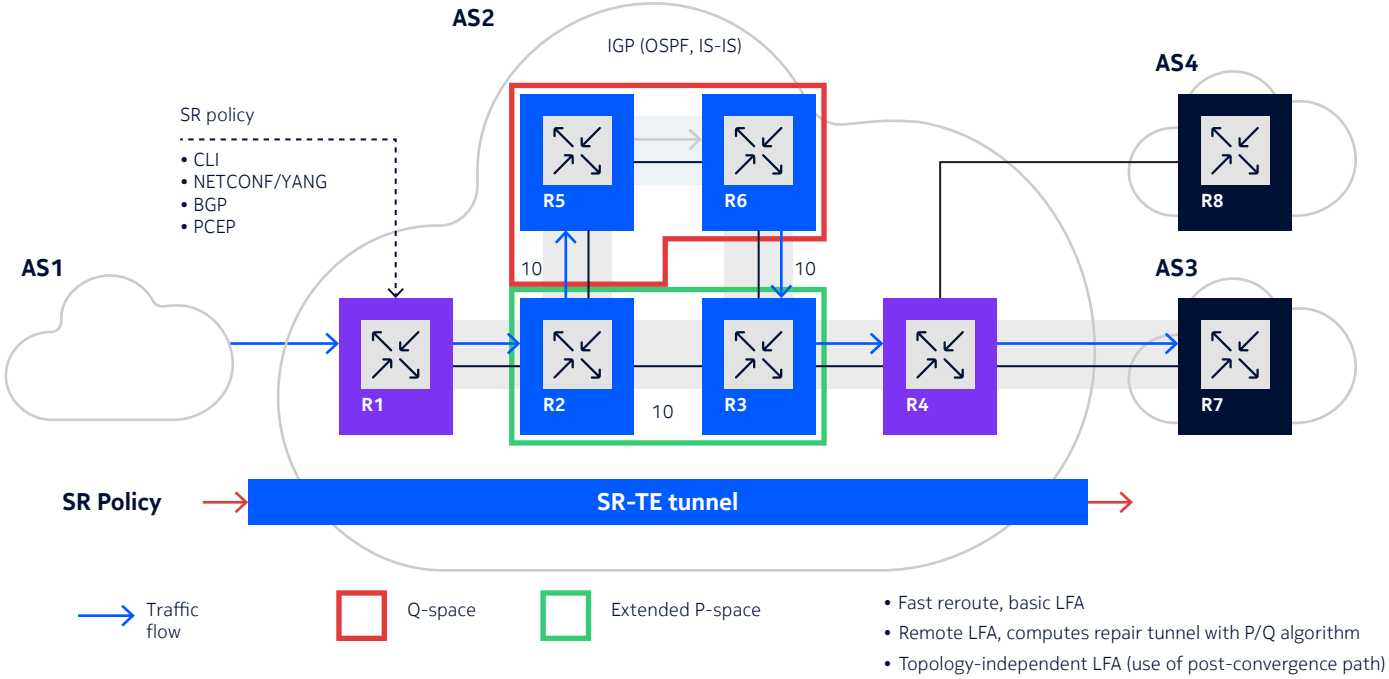

Another issue is that the backup path established through RLFA and DLFA repair mechanisms may be altered when IGP re-convergence occurs and reachability information is updated, which means a second transition that impacts the service. Topology independent LFA (TI-LFA) addresses this issue by using the so-called P/Q algorithm to engineer repair paths that reduce these two transitions to a single pre-convergence to post-convergence transition, to minimize the impact of failures on services (see Figure 2).

Seamless Bidirectional Forwarding Detection (SBFD – RFC 7880) is complementary to LFA and used for rapid and deterministic path failure detection in as little as 30 ms. SBFD sessions are established on every path of an SR-TE label switched path (LSP) or segment list of an SR Policy. If an SBFD session on the active path goes down, SR switches to a pre-programmed standby path or SR Policy. If there are multiple active segment lists in an SR Policy (ECMP), then SBFD can trigger the failed segment list to go out of service to prevent blackholing of traffic. SBFD can also detect “silent faults” that are not visible to the control plane. Since SBFD works on the end-to-end paths, it does not trigger an LFA.

Segment routing on Nokia routing platforms

Support for segment routing on Nokia routing platforms started in 2015 and has since evolved into a robust, comprehensive and versatile toolkit that has been validated and deployed by network operators for a variety of applications (see Table 1).

Table 1. Nokia routing platforms segment routing toolkit

Programmatic control	PCEP	BGP, BGP Link-State	NETCONF/YANG
Protection/assurance	LFA, TI-LFA, RLFA	Primary-secondary	LSP ping/trace, BFD
Traffic engineering	SR-TE	SR Policy	Flex-Algo
Control plane	IS-IS, OSPF, BGP or static, IPv4 or IPv6		
Data plane	SR-MPLS (IPv4, IPv6)	SRv6 (IPv6)	

Nokia routing platforms implement a comprehensive SR-MPLS feature set for IPv4 and IPv6 networks, and an advanced implementation of segment routing for IPv6 (SRv6).

Getting started with segment routing

Using the Nokia segment routing toolkit, operators can enhance traffic visibility, control and resiliency in their existing MPLS networks.

The scope and depth of the SR toolkit can be intimidating, but operators can incrementally add the SR features they need to an existing LDP/RSVP network with the options they're comfortable with. This eases the migration and enables operational experience to be gained before introducing more powerful features (see Figure 3).

Figure 3. SR applications using the Nokia toolkit

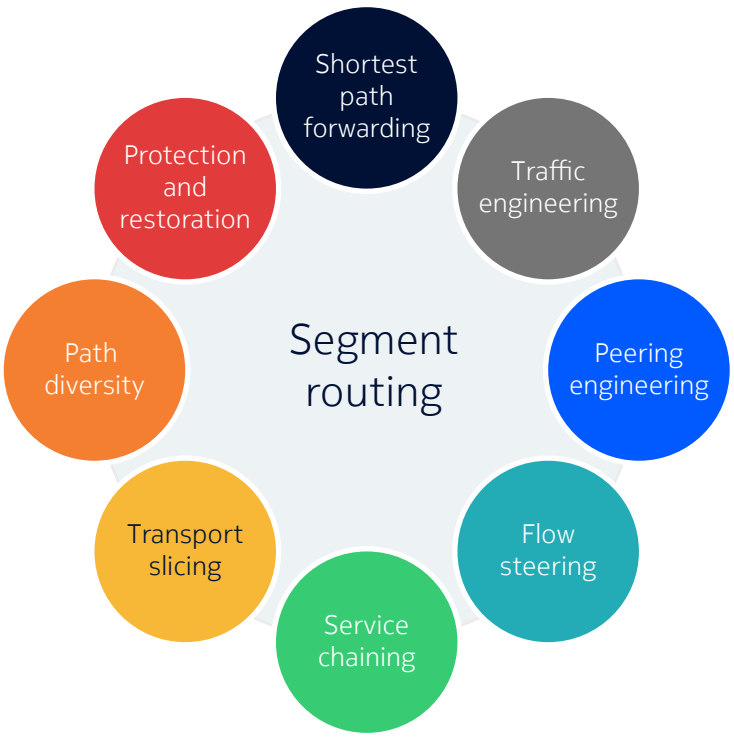

Shortest path routing is a good starting point for introducing SR as a replacement for LDP, bringing better protection coverage with TI-LFA source-routed repair tunnels. This step only requires control plane extensions for IS-IS or OSPF. Global SR label blocks and node and link adjacencies can be configured for each router, after which SR-IS-IS or SR-OSPF tunnels can be used by layer 2 and layer 3 services.

Constraint-based shortest path routing could be a logical next step, for example to engineer low-latency paths for delay-sensitive applications, or to ensure traffic flows are kept within a controlled set of links (data sovereignty). By augmenting the SR shortest path with an IGP flexible algorithm (Flex-Algo) it becomes possible to add topology constraints and alternative link-metrics to the shortest path calculation. This application invokes SR traffic engineering capabilities (SR Policy and Flex-Algo), and programmatic control (BGP and BGP-LS). SR-TE capabilities can be further extended with BGP and BGP-LS to include BGP prefix and peer segments, and steer traffic to a particular egress point (i.e., egress peer engineering).

Path diversity and end-to-end protection enable highly available, premium transport services as a more scalable alternative for MPLS fast reroute based on RSVP-TE. SR-TE LSPs with diverse primary and secondary paths are enabled by including SRLG constraints for the candidate paths. Seamless BFD is used for fast detection of failures, including silent failures that are not visible to the control plane, and to trigger a failover to the secondary path if the primary path fails.

Intent based multi-domain networking

A multi-domain network is a large network, typically controlled by one administration, that is composed of regional networks inter-connected at specific gateway locations (see Figure 4). An intent is a networking objective to be met for a specific set of traffic flows. For example, an intent may be applied to traffic that constitutes a transport slice. Lowest-delay and lowest-cost are examples of intent. The concept of “color” is important. In this context color is a numerical value carried in signaling messages to indicate intent. Color values and mappings are decided independently by each administration to suit the engineering and business needs of the network.

Multi-domain networks are increasingly common, even for national and regional networks with a relatively small number of routers that do not necessarily have extensive geographical reach. There are a number of potential benefits of a segmented, multi-domain design when compared to one large network:

- The responsibility for network administration can be divided across multiple teams for operational efficiency.
- Operators can deploy routers that have fewer capabilities and are less expensive due to reduced requirements on memory (to hold state for the entire network), next-hop scale and label stack depth. This can reduce equipment costs.
- Gateways provide a point of operational control for traffic moving between regions.
- There is a smaller affected user community if problems occur leading to an improved end user experience.
- The transition to new technologies, such as segment routing, can proceed domain by domain reducing operational risk exposure.

Figure 4. A typical multi-domain network

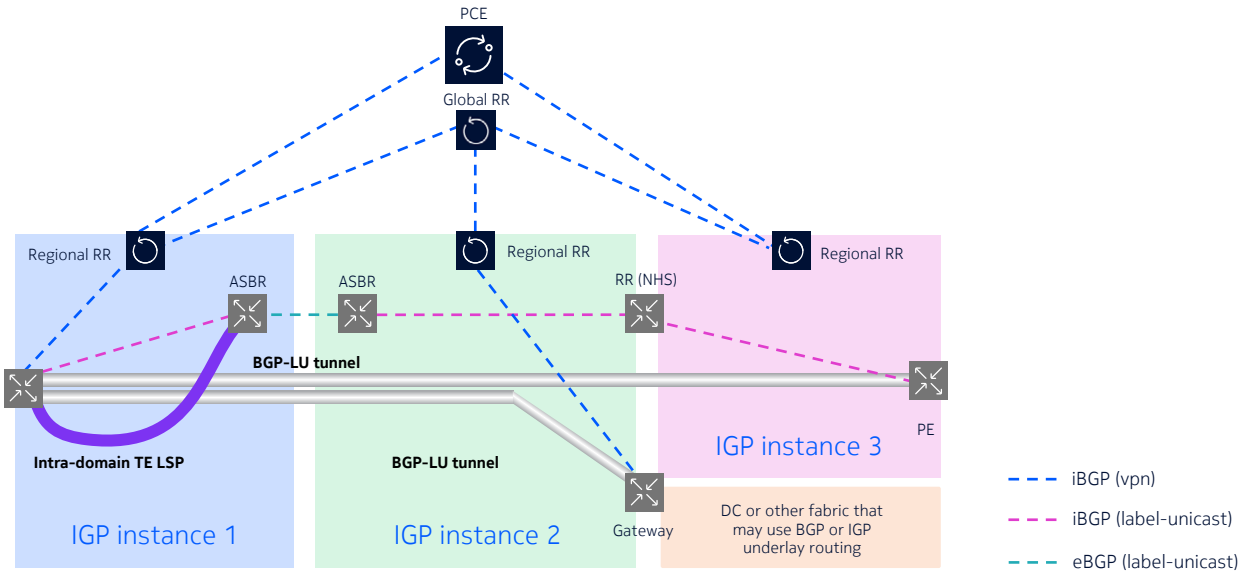

As illustrated in Figure 4, there are different IGP instances in each domain. BGP labeled unicast (BGP-LU) advertises PE loopback addresses and propagates them across domain boundaries using next-hop-self. Segment routing tunnels are used to resolve BGP-LU routes in local domains. Services gateways are selectively placed at domain boundaries when they are needed for route aggregation or service and signaling translation. BGP service routes are advertised PE-to-PE or PE-to-gateway with next-hop unchanged. A PCE can be used to meet more complex traffic engineering requirements.

The following are the building blocks of intent-aware multi-domain transport:

- Intent-aware capability in all transport domains individually
- End-to-end service overlay
- PCE control

Intent-aware transport enforces the intent from end to end using SR-TE LSP and SR Policy. Active and standby protection are provided together with ECMP load balancing. The intent crosses domain boundaries and uses a binding SID to reduce the label stack depth. A tunnel is associated with an intent either using a local admin-tag policy or implicitly via the color attribute. Distributed routing enforces the intent within each domain and is sufficient in many applications. Flex-Algo sets up tunnels that are shortest-path, subject to constraints (i.e., the intent). BGP-LU stitches the per-domain tunnels.

An end-to-end service overlay sits on the intent-aware transport. For example, a BGP EVPN or an IP-VPN creates virtual private (overlay) networks on the common underlay that spans domain boundaries. A VPN associates with a 5G slice, a data center tenant, an end user application or an end customer. An EVPN or IP-VPN route color extended community signals intent for flows matching the advertised route. The service gateway allows flexible transport options and scalable domain connectivity while maintaining intent end-to-end.

A PCE enhances the intent-based, multi-domain transport by, for example, establishing SR-TE path diversity across colored topologies or slices. It enforces intent within domains and across domain boundaries. The PCE also allows setting upper bounds on the accumulated delay budget of a computed SR-TE path. PCE redundancy is important, but with a fallback strategy in place the network will gracefully switch to distributed routing to preserve intent within each domain during any period of PCE unavailability.

Segment routing in Telco Cloud data centers

Segment routing is also being adopted for data center networks and DCI applications. The use case example shown in Figure 5 applies to SR-enabled data center fabrics and uses SR-TE to establish a reliable, low-latency transport service to interconnect two Ethernet VPN (EVPN) service instances in geographically separated data centers over a physically diverse, dual-plane core network. The EVPN endpoints are instantiated on leaf nodes with multiple uplinks to a cluster of spine routers. As the state of the uplink between the spine and core is not visible to the leaf nodes, protection from uplink failures is required to avoid blackholing traffic.

Figure 5. Interconnecting leaf-spine DC fabrics over a low-latency dual-core network

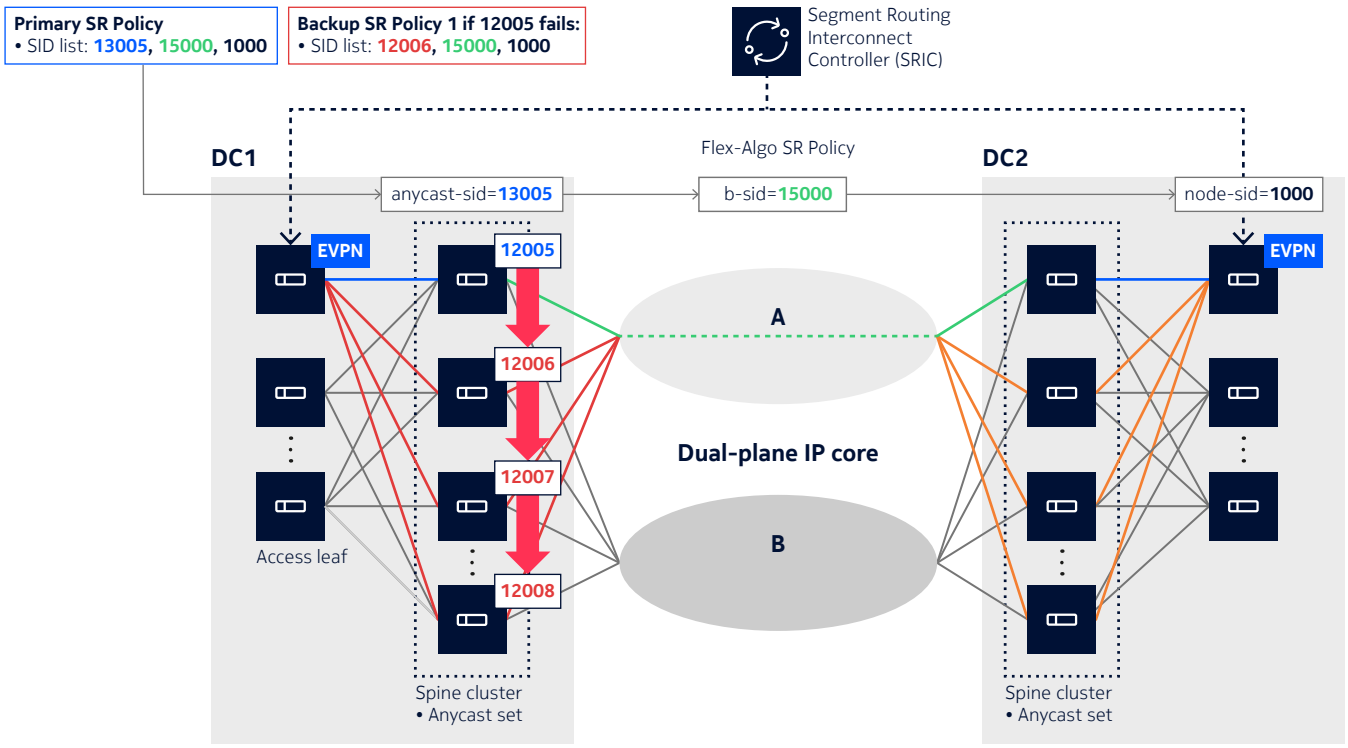

To accomplish this, all spine nodes of the connecting IP fabric are included in an anycast set with anycast SID 13005, which allows load balancing across the spine cluster nodes. The low-latency transport tunnel through the dual-plane core network segment is established using Flex-Algo and assigned with binding SID 15000. The destination leaf node hosting the remote EPVN instance is assigned SID 1000.

The primary SR path is implemented as a multi-path policy defined by the spine-anycast-SID (13005), the low-latency-core-SR-policy-BSID (15000) and the destination-node-SID (1000). The primary SR policy is backed up with an ECMP set of single path SR policies that identify the next alternative spine node to take over in case the uplink of the primary spine node fails. For example, if the core uplink of primary spine with node-SID 12005 fails, secondary spine-node 12006 will steer the EVPN leaf traffic into the low-latency SR core policy.

Legacy migration to segment routing

Many operators have already deployed traffic engineering and protection capabilities based on LDP or RSVP-TE. SR-MPLS offers a smooth and straight-forward migration path for these legacy networks and eases the evolution to next-generation networks using IPv6.

SR-MPLS introduction and LDP/RSVP-TE migration

Because SR-MPLS supports native IPv4 and IPv6 data planes without adding new hardware requirements, it can normally be deployed on existing hardware through a software upgrade. The capabilities of the SR toolkit can be gradually introduced and yield immediate service benefits, without sacrificing any transport capabilities offered by LDR and RSVP-TE:

- BGP base router and virtual routing and forwarding (VRF) routes can be migrated service-by-service using auto-bind-tunnels.
- Alternatively, the migration can be done PE-by-PE. Ingress PE can bind the same VRF to LDP/RSVP-TE or SR-IS-IS/SR-OSPF/SR-TE tunnel based on the capability of each remote PE.
- Migrating RSVP-TE to SR-TE may be impractical in some networks. In these cases, SR-TE LSPs can still be introduced on the same network, for example to support new services that require advanced capabilities such as path diversity or multi-domain SR-TE path.
- For new networks we recommend skipping RSVP-TE and deploy SR-IS-IS/SR-OSPF and SR-TE. Table 2 compares some of the important differences between RSVP-TE and SR to highlight the key benefits of legacy migration.

Table 2. Comparing segment routing with RSVP-TE

Aspect	RSVP-TE	Segment Routing
Data plane state	Locally significant LSP context	Global or local LSP context
Control protocols	Two: IS-IS or OSPF, RSVP	One: IS-IS or OSPF
Traffic engineering	Network-based: IGP-TE/RSVP-TE	Network and controller-based: SR-TE, SR policy, Flex-Algo IGP-TE with CSPF, SDN Controller
Inter-area/AS	Inter-AS TE (needs IGP across BGP boundaries)	Binding SID Egress Peer Engineering
Network programmability	PCEP, NETCONF PCC-init	PCEP, BGP, NETCONF PCC-init or PCE-init
Granularity and steering	Single level: Next hop or end-to-end paths ECMP into parallel end-to-end LSPs	Multi-level: Steer traffic across any kind of crafted path (node, adjacency, anycast, etc.) Can combine loose/ECMP hops within strict TE path
Protection/OAM	MPLS FRR, Active/standby LSP ping/trace/self-ping/LSP BFD Control plane-based re-optimization	LFA/RLFA/TILFA, Active/standby LSP ping/trace/seamless BFD IGP + SBFD-based re-optimization
Multicast support	Point-to-multipoint (P2MP) LSPs	Tree-SID (SR P2MP)

In summary, segment routing offers several operational benefits compared to LDP/RSVP-TE:

- Better scalability by only requiring ingress routers to keep state information
- Operationally simpler by avoiding the need for an additional signaling protocol (RSVP-TE and LDP)
- Greater control flexibility and scope by also supporting centralized TE controllers to facilitate inter-domain/inter-area traffic engineering applications (egress peering engineering, etc.)
- Granular multi-level traffic steering across any crafted path with strict and loose constraints and built-in ECMP load-balancing capabilities
- Better protection with full coverage for most deployed topologies through LFA, RLFA and TI-LFA
- Multicast-capable with the SR P2MP policy which requires a few extensions to the PCEP protocol used with P2P SR policy. There is no requirement for control protocols such as P2MP RSVP, Multicast Label Distribution Protocol, and Protocol-Independent Multicast, simplifying network operations.

SRv6 introduction and SR-MPLS interworking

The SRv6 standardization effort is far more recent than SR-MPLS and still evolving. The goal is to provide a programmable framework for IPv6 networks that tightly integrates overlay service functions and underlay transport functions (RFC 8986). SRv6 takes advantage of the larger address space to encode 128-bit SIDs as an IPv6 address:

- SRv6 shortest path routing encodes the destination SID in the Destination Address (DA) field of the outer IPv6 header.
- SRv6 source routing encodes the top SID in the DA field and the rest of the SIDs of nodes the packet must visit as SID list in the SRH.

Because the SRv6 service extensions impose additional data path requirements on IPv6 packet header processing, feature support is contingent on routing silicon and typically demands new hardware. For example, FP4 silicon or above is required to support the full SRv6 PE feature set on Nokia 7750 SR and 7950 XRS routers, and Jericho 2 silicon is required for 7250 IXR platforms.

Figure 6. SRv6 support and SR-MPLS interworking

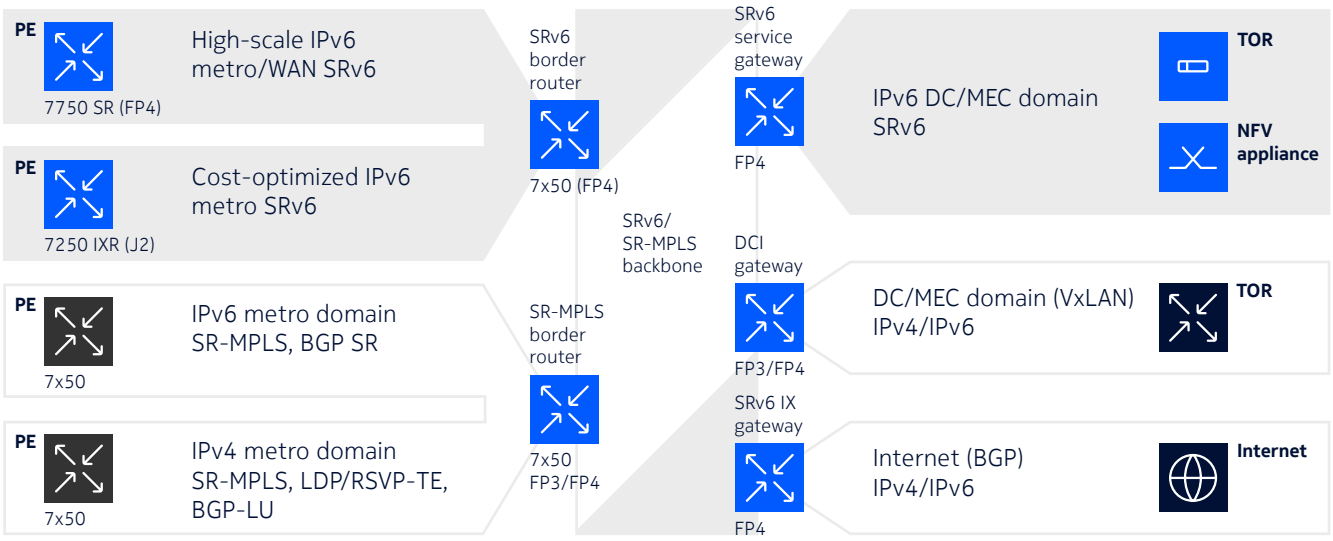

Segment Routing deployments are increasing and SR-MPLS is leading in overall numbers of deployments as the more mature and proven choice in most transport network use cases. Testing and trials of SRv6 applications are being conducted concurrently. Early adopters of SRv6 are targeting a few key deployment use cases. For example, many Telco Cloud edge data centers are deployed in hybrid or public cloud and an IPv6 underlay is becoming common. One initial use case for SRv6 is to simply provide the framework to overlay VPN services on the IPv6 transport network. Subsequently, more advanced applications such as traffic engineering, load-balancing, and service chaining can be deployed.

The second use case can be introduced with the deployment of a new IPv6 backbone, which provides a good incentive to introduce IPv6 underlay and an opportunity to experiment with SRv6. Because SRv6 uses IPv6 addresses as locators and SIDs, it can summarize locator prefixes at network domain boundaries and reduce the need to use a separate protocol, such as BGP-LU with multi-domain MPLS, for propagating end-to-end PE address reachability.

With the selective introduction of SRv6, MPLS/SRv6 gateways in the base router or in a VRF instance allow a seamless extension of services and connectivity to the rest of the network (Figure 7). This preserves prior investments made in the network.

The SRv6 standard is maturing with the addition of new features and enhancements. The introduction of SRv6 micro-segments to compress SIDs greatly reduces the bandwidth overhead of SRv6. Micro-segments are supported in SR OS. Public interoperability is also progressing well, while router hardware coverage is increasing and automation tools are being developed.

Nokia differentiators

Segment routing is a powerful and proven technology for deploying highly programmable IP services that meet deterministic service level objectives on cost, performance and reliability. Segment routing addresses all operational scalability issues of legacy traffic engineering and protection approaches using LDP or RSVP-TE and enables a wide range of new applications, especially when used in combination with the Nokia NSP.

Nokia routing platforms offer a comprehensive segment routing toolkit for IP/MPLS networks, next-generation IPv6 networks, and the transition between them. The segment routing capabilities that differentiate Nokia service routers include:

- Standards compliance and proven multivendor interoperability through Orange labs/EANTC
- A comprehensive SR-MPLS to SRv6 interworking capability, deployed both in public interoperability tests and in real world production environments
- Full SR support for IP and Ethernet VPNs, Virtual Private LAN, and virtual leased line services
- Complete SR policy management support through PCEP, SR-TE LSP and BGP SR policy
- Extensive protection and restoration through basic and remote LFA, topology independent LFA, link and node protection and advanced LFA policy for control of backup path
- The industry's largest label stack (LER push up to 12 labels, LSR hash up to 16 labels)
- The industry's most flexible SR gateway solution with capabilities to interwork and translate among different data and control planes: MPLS, SR-MPLS, SRv6, Virtual Extensible LAN, MPLS-over-UDP
- Programmability and flexibility of Nokia routing silicon to support variable SID lengths (32-bit, 16-bit or 128-bit) and new data path optimizations to compress segments in SRv6

Please consult our product documentation and user guides for detailed information on supported features and platforms, or contact Nokia Sales to learn more.

Abbreviations

BFD	bi-directional failure detection
BGP	Border Gateway Protocol
BGP-LS	Border Gateway Protocol Link-State
BSID	binding segment identifier
CLI	command line interface
DCI	data center interconnect
DLFA	directed LFA
ECMP	equal cost multi-path
EVPN	Ethernet VPN
IGP	Interior Gateway Protocol
IPFIX	IP Flow Information Export
IS-IS	Intermediate System-to-Intermediate System
LDP	Label Distribution Protocol
LER	label edge router
LFA	loop-free alternate
LSP	label switched path
LSR	label switch router
MPLS	Multiprotocol Label Switching
NFIX	Network Function Interconnect
NLRI	network layer reachability information
NSH	Network Service Header
NSP	Network Services Platform
OSPF	Open Shortest Path First
P2MP	point-to multipoint
PCE	path computation element
PCEP	Path Computation Element Protocol
PE	provider edge
PLR	point of local repair
RLFA	remote LFA
RR	route reflection
RSVP	Resource Reservation Protocol
RSVP-TE	Resource Reservation Protocol - Traffic

RSVP-TE	Resource Reservation Protocol - Traffic Engineering
SBFD	Seamless Bidirectional Forwarding Detection
SF	Service Function
SFC	Service Function Chain or Chaining
SFF	Service Function Forwarder
SFP	Service Function Path
SID	segment identifier
SLA	service level agreement
SR	Service Router
SR	segment routing
SR OS	Service Router Operating System
SRIC	Segment Routing Interconnect Controller
SRLG	Shared Risk Link Group
SR-MPLS	segment routing for MPLS
SR-TE	segment routing – traffic engineering
SRv6	segment routing for IPv6
TI-LFA	topology-independent LFA
UDP	User Datagram Protocol
VRF	virtual routing and forwarding

Standards and recommendations

RFC 3209	RSVP-TE: Extensions to RSVP for LSP Tunnels
RFC 4655	A Path Computation Element (PCE)-Based Architecture
RFC 5286	Basic Specification for IP Fast Reroute: Loop-Free Alternates
RFC 5440	Path Computation Element (PCE) Communication Protocol (PCEP)
RFC 6020	A Data Modeling Language for the Network Configuration Protocol (NETCONF)
RFC 6241	Network Configuration Protocol (NETCONF)
RFC 6571	LFA applicability in service provider networks
RFC 7011	IPFIX Protocol Specification
RFC 7432	BGP MPLS-based Ethernet VPN
RFC 7490	Remote Loop-Free Alternate Fast Reroute (RLFA)
RFC 7510	Encapsulating MPLS in UDP
RFC 7665	Service Function Chaining Architecture
RFC 7752	Link-State Information Distribution Using BGP
RFC 7854	BGP Monitoring Protocol
RFC 8277	BGP and Labeled Address Prefixes
RFC 8402	Segment routing architecture
RFC 8426	Recommendations for RSVP-TE and segment routing LSP coexistence
RFC 8661	Segment routing interworking with LDP
RFC 8663	MPLS Segment Routing over IP
RFC 8665	OSPF Extensions for Segment Routing
RFC 8667	IS-IS Extensions for Segment Routing
RFC 8754	IPv6 Segment Routing Header (SRH)
RFC 8986	Segment Routing over IPv6 (SRv6) Network Programming
RFC 9085	Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing
RFC 9252	BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)
RFC 9259	Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)
RFC 9352	IS-IS Extension to Support Segment Routing over IPv6 Dataplane
Flex-Algo	IGP Flexible Algorithm

About Nokia

At Nokia, we create technology that helps the world act together. As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

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

Document code: CID210750 (December)