Segment routing

A brief overview

Antonio Prado
https://www.prado.it

Segment-routing 1/3

SR utilizes the source routing concept. A **node** guides a packet along a predetermined series of steps, known as "**segments**." Whether it is topological or service-centric, a segment can serve as a representation for any instruction. Within an SR **domain**, a segment can possess either a local or global semantic. SR enables the restriction of a flow to a particular topological path, while also maintaining per-flow information solely at the ingress node(s) of the SR domain.

RFC 8402 Segment Routing Architecture

Segment-routing 2/3

Implementing SR in the **MPLS** architecture requires no alterations to the forwarding plane. An MPLS label is used to encode a segment. The encoding of an ordered list of segments is depicted as a stack of labels. The top item on the stack is the segment to be processed. Once a segment is finished, the corresponding label is removed from the stack.

RFC 8402 Segment Routing Architecture

Segment-routing 3/3

The **IPv6** architecture can utilize SR through a newly introduced routing header. An IPv6 address serves as the encoding for a segment. A sequential sequence of segments is translated into a sequential sequence of IPv6 addresses within the routing header. The packet's Destination Address (DA) reveals the active segment. A pointer in the new routing header denotes the upcoming active segment.

RFC 8402 Segment Routing Architecture

Segment-routing elements

Segment Identifier (SID)	Each node in the network is assigned a unique identifier known as a Segment Identifier (SID)
Label Stack	A packet carries a label stack, which consists of a series of SIDs. These SIDs represent the nodes that the packet will visit.
Source Node	The source node determines the path the packet should take through the network and assigns the appropriate SIDs to the packet's label stack.
Forwarding Decision	As the packet travels through the network, each node makes a forwarding decision based on the top label in the stack. It pops the top label and forwards the packet to the next node indicated by the popped label.
Node Behavior	Nodes in the network can perform various actions based on the SIDs in the label stack. These actions may include simple forwarding, traffic engineering, or service chaining.
Destination Node	The packet continues to traverse the network until it reaches the destination node, where the label stack is empty, and the payload is delivered to the final destination.

Segment-routing example

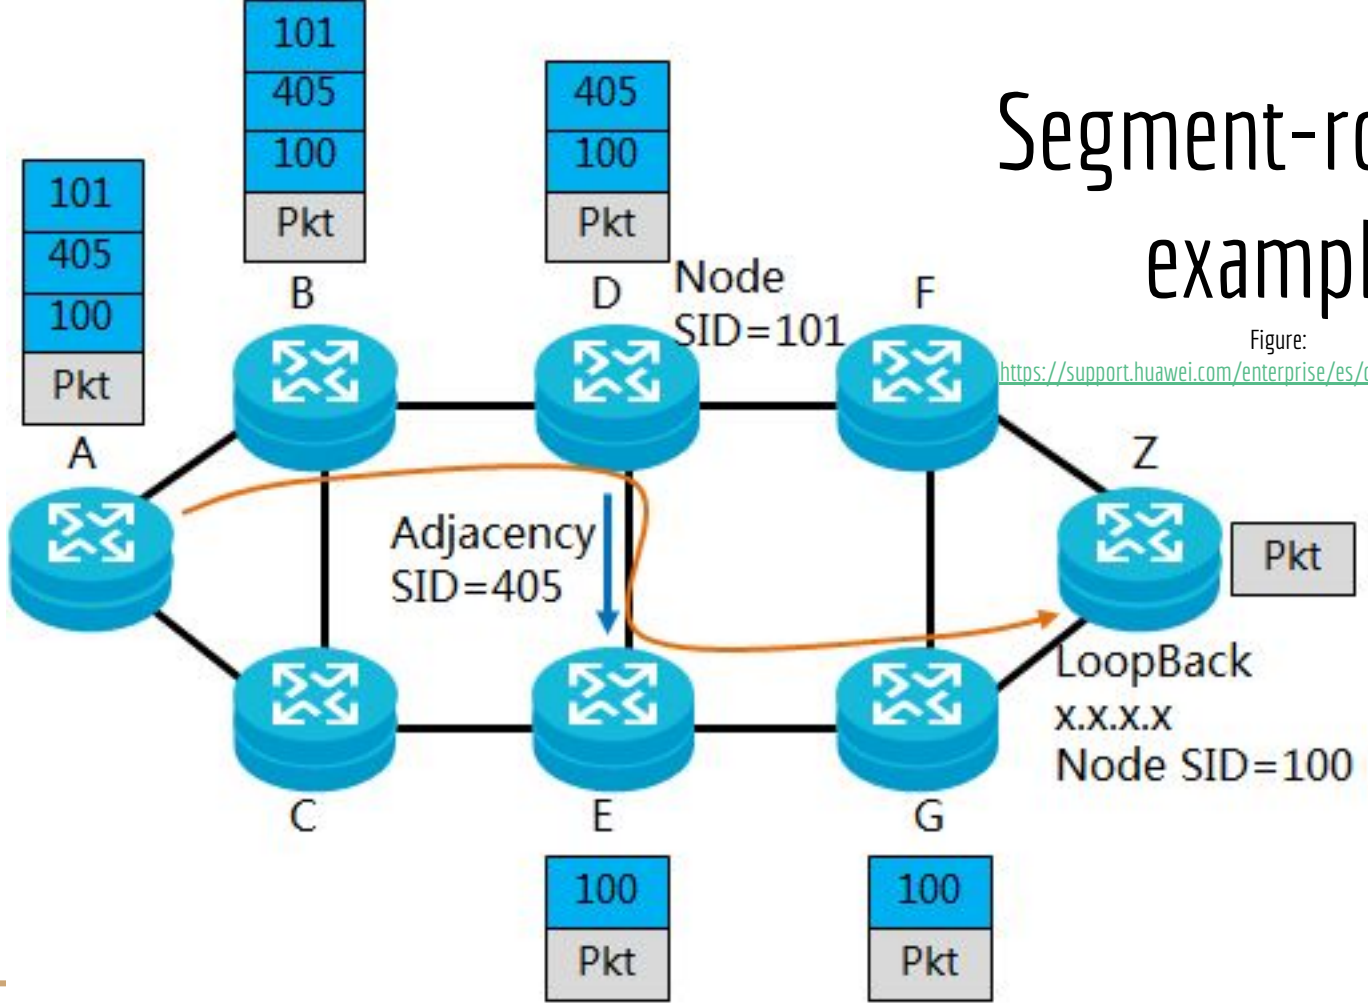

Figure: https://support.huawei.com/enterprise/es/doc/EDOC1100092117

SR-MPLS vs SRv6: Data Plane Protocol

SR-MPLS	SRv6
Encapsulation: MPLS labels are used to represent segments in the data plane.	**Encapsulation:** IPv6 is used as the data plane encapsulation.
Label Stack: The MPLS label stack is leveraged to encode the segment identifiers.	**Segment Identifier:** Segments are represented by IPv6 addresses.
Underlying Protocol: Uses MPLS as the underlying transport protocol.	**Routing Header:** The SRv6 architecture uses a new type of IPv6 routing header called the SRH (Segment Routing Header) to encode segment identifiers.

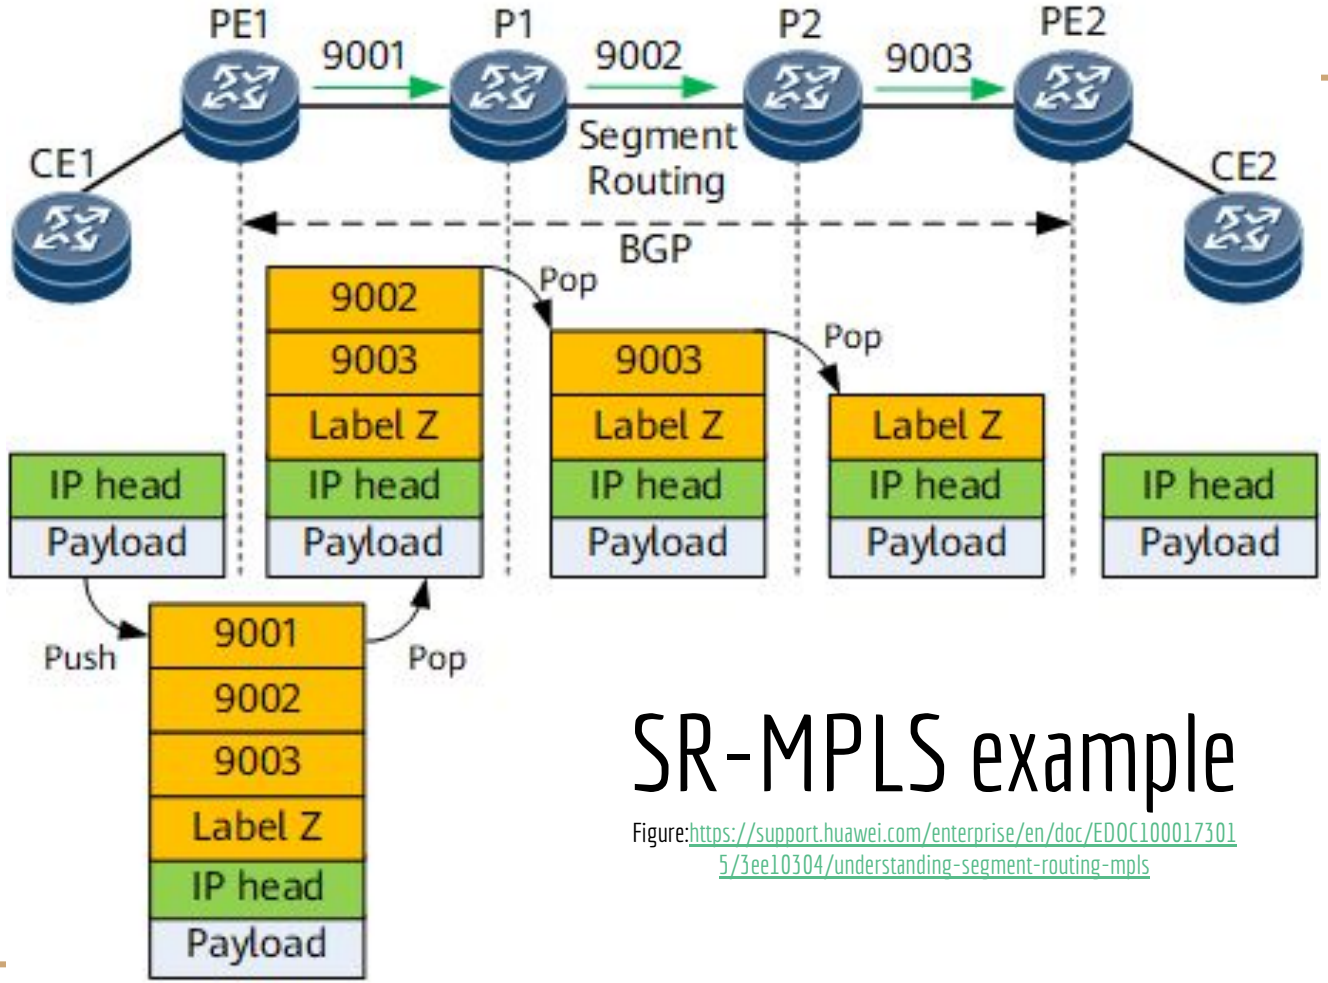

SR-MPLS example

Figure:https://support.huawei.com/enterprise/en/doc/EDOC1000173015/3ee10304/understanding-segment-routing-mpls

SR-MPLS vs SRv6: Network Layer Dependency

SR-MPLS	SRv6
Network Layer: Primarily operates at the MPLS layer, independent of the underlying IP layer.	**Network Layer:** Integrates directly with the IPv6 network layer.
Transition: Existing MPLS networks can incrementally deploy SR-MPLS without requiring changes to the IP infrastructure.	**IPv6 Dependency:** Requires native IPv6 support in the network.

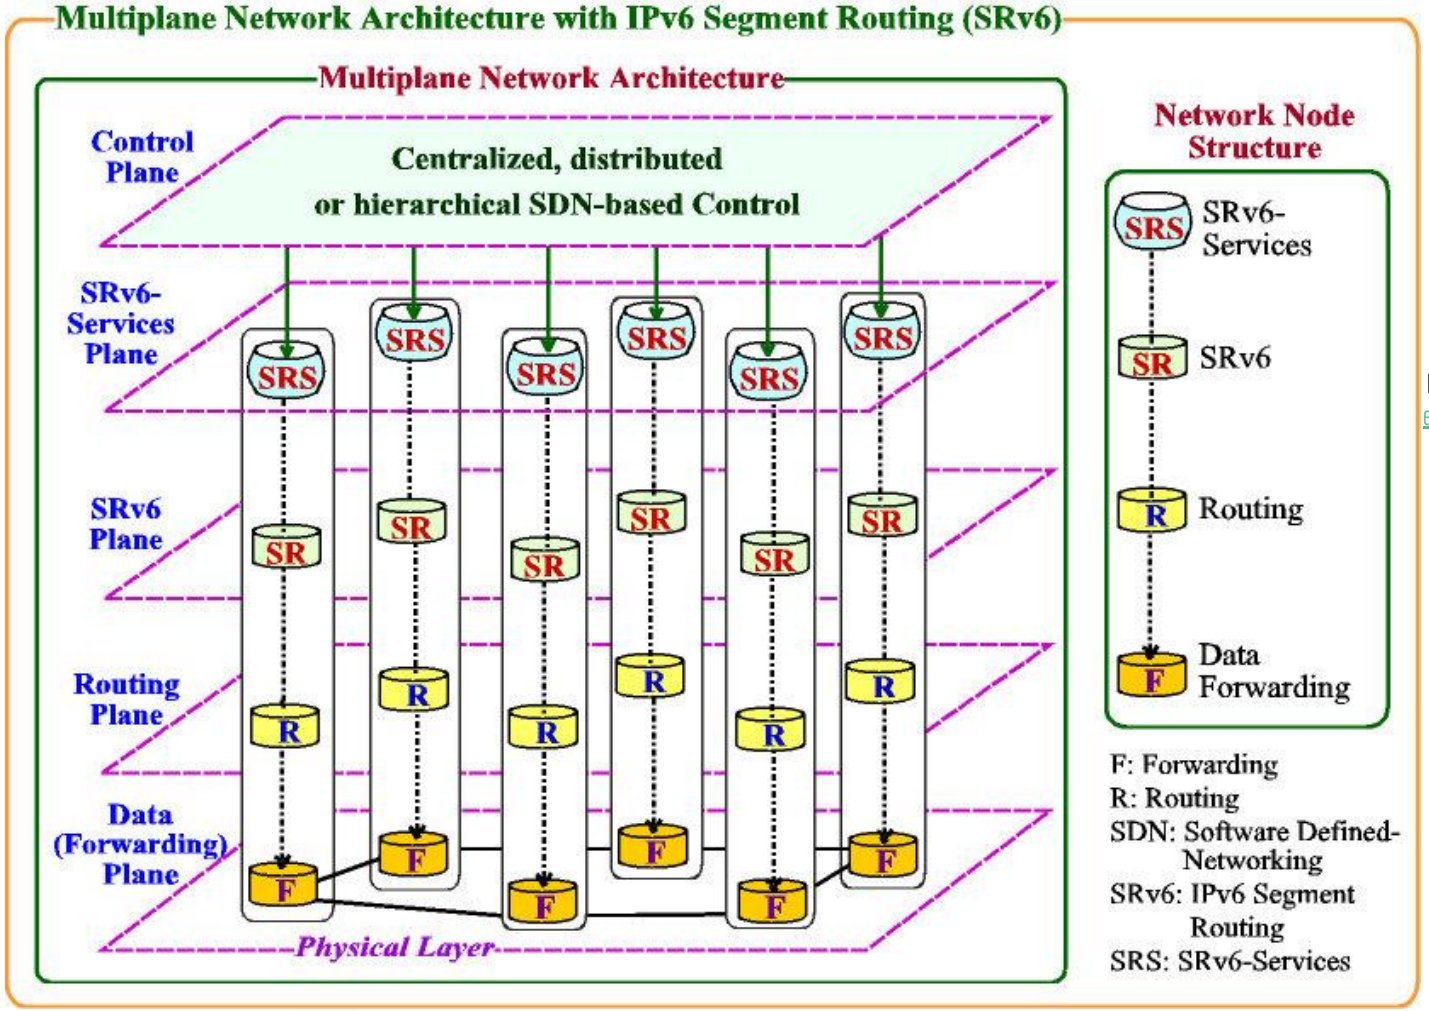

SRv6 example

Figure:https://www.researchgate.net/figure/General-Multiplane-Structure-of-IPv6-Networks-with-SRv6_fig2_336349238

SR-MPLS vs SRv6: Scalability and Flexibility

SR-MPLS	SRv6
Scalability: MPLS networks are known for their scalability, making SR-MPLS a suitable choice for large-scale networks.	**Scalability:** Inherits the scalability benefits of IPv6, potentially providing a large address space for segments.
Traffic Engineering: Provides flexible traffic engineering capabilities.	**Flexibility:** Supports flexible and extensible services with the use of IPv6 routing headers.

SR-MPLS vs SRv6: Segment Representation

SR-MPLS	SRv6
Label Stack: Segments are represented using MPLS labels in the label stack.	**IPv6 Addresses:** Segments are represented using IPv6 addresses.
Label Operations: Nodes in the network perform label operations to determine the next-hop based on the top label.	**Routing Header:** The SRH in the IPv6 header contains the list of segments to be processed.

SR-MPLS vs SRv6: Deployment Considerations

SR-MPLS	SRv6
Interoperability: Can be deployed in networks where MPLS is already in use.	**Greenfield Deployments:** Well-suited for new network deployments or networks with native IPv6 support.
Transition: Allows for a gradual migration of MPLS networks to SR-MPLS.	**IPv6 Infrastructure:** Requires the presence of IPv6 infrastructure in the network.

SRv6 concerns

IPv6 adoption and transition

IPv6 Dependency: SRv6 relies on IPv6, and its successful deployment depends on the widespread adoption of IPv6 across the network infrastructure. SR-MPLS, on the other hand, is agnostic to the underlying IP version.

SRv6 concerns

Header Manipulation

Security Implications of SRH: The Segment Routing Header (SRH) in SRv6 introduces a new header in the IPv6 packet, which may have security implications that are not present in SR-MPLS. Security considerations for this new header need to be carefully evaluated.

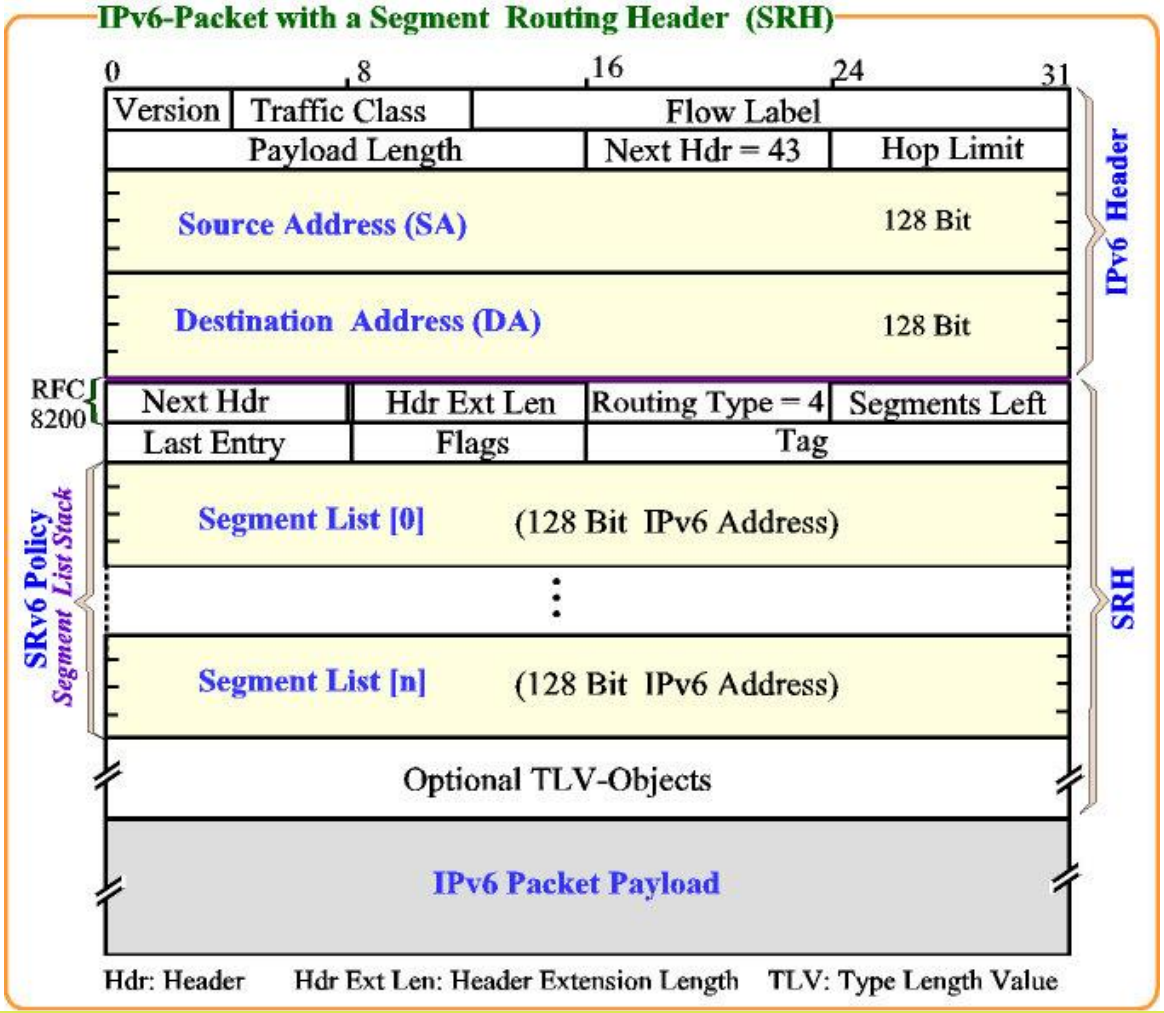

SRv6 Header

Figure: https://www.researchgate.net/figure/Structure-of-the-SRH-and-its-Encapsulation-in-the-IPv6-Packet_fig6_336349238

SRv6 concerns

Header Size and Fragmentation

Larger Headers in SRv6: The addition of the SRH in SRv6 headers increases the size of packets, potentially leading to fragmentation issues and susceptibility to fragmentation attacks. SR-MPLS does not introduce additional headers in the same way.

SRv6 concerns

Segment Identifier Spoofing

Spoofing Attacks: Malicious nodes might attempt to spoof or manipulate the Segment Identifiers (SIDs), leading to unauthorized redirection of traffic or bypassing security policies.

SRv6 concerns

Denial of Service (DoS) Attacks

Resource Exhaustion: Attackers may attempt to exhaust network resources by overwhelming nodes with a large number of SRv6 headers, leading to performance degradation or denial of service.

SRv6 concerns

Packet Length and Fragmentation

Fragmentation Attacks: The use of SRv6 headers increases the size of packets, potentially leading to fragmentation. Attackers may attempt to exploit vulnerabilities associated with fragmented packets.

SRv6 concerns

No default to fail closed

“A limited domain must default to fail closed, and require explicit configuration to open it. Think of MPLS: you must manually enable it on an interface (’set protocols mpls interface et-0/0/0’) before the border of my limited domain extends out the interface, and my neighbor has to enable it on their side too. We *can* both choose to do so, but if we don’t, the domain doesn’t cross the border.” (Warren Kumari)

SRv6 concerns

Any solutions?

“SRv6 as designed has evoked interest from various parties, though its deployment is being limited, amongst other things, by known security problems in its architecture. This document specifies a standard way to create a solution that closes some of the major security concerns, while retaining the tenants of the SRv6 protocol.”

SRv6 concerns

Any solutions?

To learn more, click on the following link:

https://datatracker.ietf.org/doc/draft-raviolli-intarea-trusted-domain-srv6/

...and SRm6?

Segment Routing Mapped To IPv6

SRm6 is a Segment Routing (SR) solution that supports a wide variety of use-cases while complying with IPv6 specifications. SRm6 is optimized for ASIC-based routers that operate at high data rates.

SRm6 vs SRv6

Instruction Encoding

SRm6 encodes topological instructions in 16 or 32-bit SIDs that appear in the CRH. It also encodes service instructions in IPv6 Destination Options.

SRv6 encodes all instructions in the low-order bits of the IPv6 Destination Address.

SRm6 vs SRv6

IPv6 Address Semantics

In SRm6 an IPv6 address always represents a network interface.

In SRv6, an IPv6 Destination Address can represent either of the following:

- A network interface
- An SRv6 SID, whose high-order bits are used for routing and low-order bits represent an instruction.

SRm6 vs SRv6

Routing Extension Header Size 1/2

SRv6 and SRm6 both encode path information in a Routing extension header. SRv6 uses the Segment Routing Header (SRH) [RFC8754], while SRm6 uses either the 16 or 32-bit version of the CRH.

SRm6 vs SRv6

Routing Extension Header Size 2/2

Large Routing headers are undesirable for the following reasons:

- Many ASIC-based forwarders copy all headers from buffer memory to on-chip memory. As header sizes increase, so does the cost of this copy.
- Because Path MTU Discovery (PMTUD) [RFC8201] is not entirely reliable, many IPv6 hosts refrain from sending packets larger than the IPv6 minimum link MTU (i.e., 1280 bytes). When packets are small, the overhead imposed by large Routing Headers is excessive.

SIMPLE DECISION TREE TO CHOOSE SR TECHNOLOGY

Figure: https://blogs.juniper.net/en-us/service-provider-transformation/choosing-the-right-segment-routing-flavor-for-business-agility

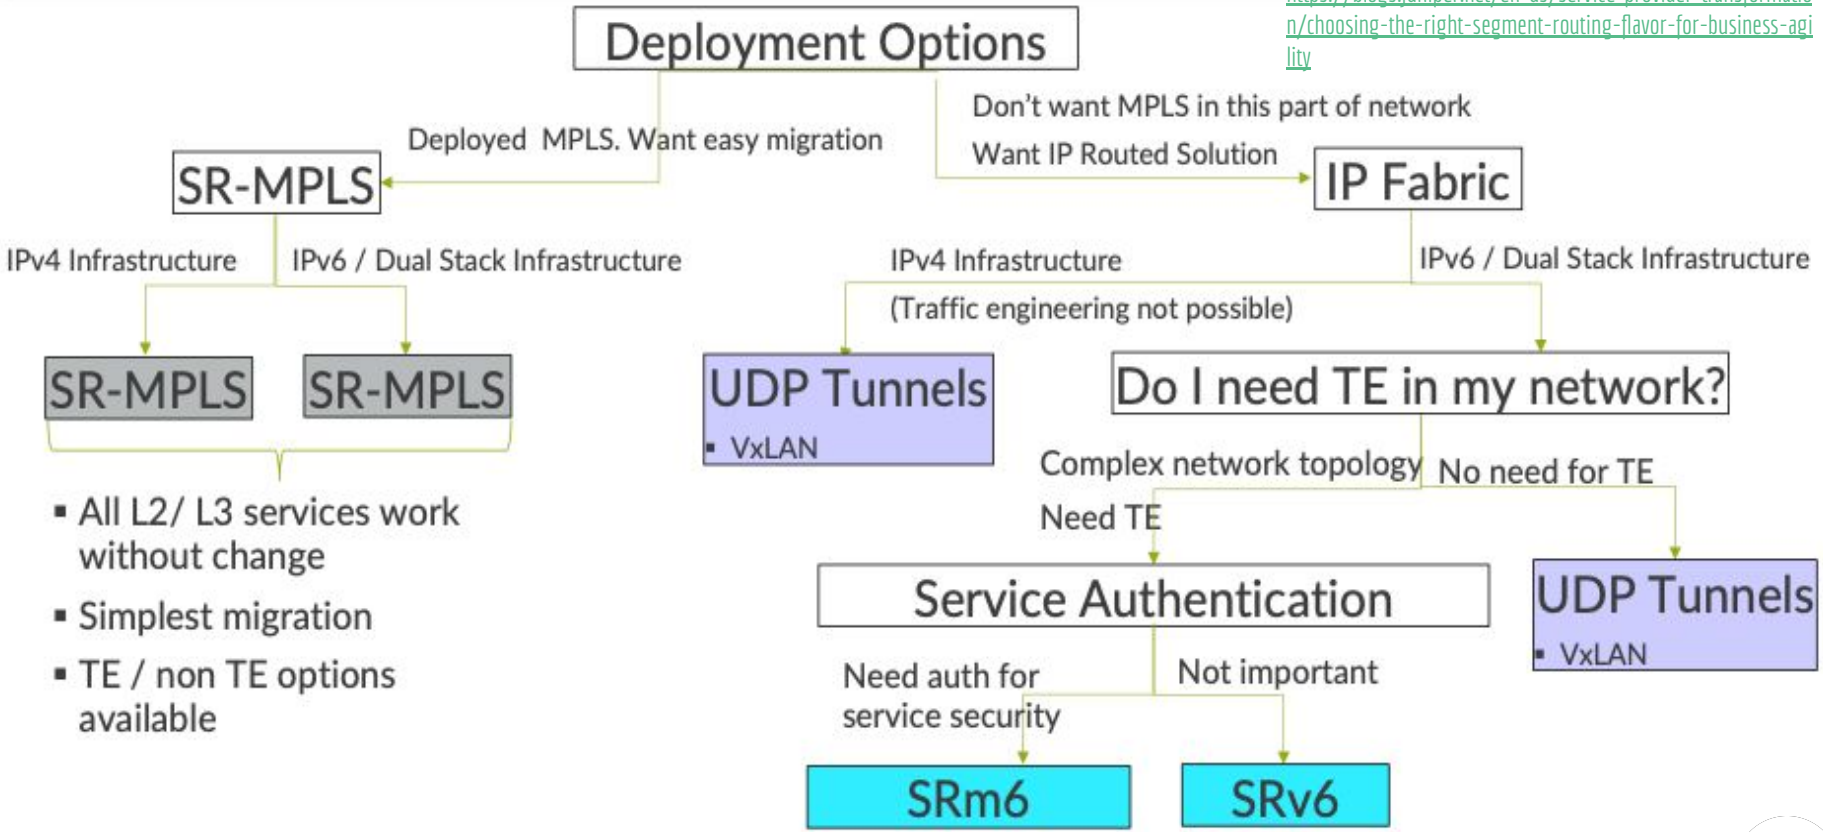

Unsolicited advice:
before adopting a new technology, it is helpful to ask yourself questions:

- what needs do I need to meet?
- what problems do I need to solve?
- do I risk vendor lock-in?
- do I have alternatives?

The answers to these questions can guide us to make a more thoughtful and informed choice.