

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 7 September 2022

J. Xie
Huawei Technologies
6 March 2022

Network Programming Interface for Provisioning of Underlay Services to
Overlay Networks Using SRv6
draft-xie-spring-srv6-npi-for-overlay-00

Abstract

This document describes a framework and a detailed suite of network programming interface (NPI) examples for provisioning of underlay services to overlay networks. It provides background by reviewing the growing pains that commonly faced today by enterprise for its flexible WAN sites connection. It assumes that WAN connection services are and will continue to be provided by multiple underlay networks operated by different administrative entities. Based on the pains and the assumptions, this document propose to use SRv6 binding SIDs (BSIDs) on a transport network (TN) edge router as NPI for service provisioning to be accessed remotely and securely by a customer router that constitutes to a higher overlay network, including the requirements of such service, the illustration of how it may work, and the possible applicability of the solution.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

Internet-Draft Network Programming Interface using SRv6

March 2022

This Internet-Draft will expire on 7 September 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	3
2.	Terms and Abbreviations	3
3.	Background and the Problem	4
3.1.	Internet and Transport Network for WAN Transit	4
3.2.	Seperation of Transport Network and Access Network	5
3.3.	The Problem	7
4.	Architecture and Design Consideration of NPI	7
4.1.	Concept and Architecture of NPI	7
4.2.	Transmission Association and Transmission Policy	9
4.3.	Data-plane Consideration	10
4.4.	Control-plane Consideration	11
5.	Requirements of NPI	11
5.1.	Scalable for Multiple Tenants	11
5.2.	Support for Common Service	12
5.2.1.	Slice Service	12
5.2.2.	SR-Policy Service	12
5.2.3.	Multicast Service	13
5.2.4.	Other Services	13
5.3.	Cost-Effective Encapsulation	13
5.4.	Secure for Remote Accessing	13
5.5.	Manageability	14
6.	Examples and Illustrations of NPI	14
6.1.	NPI.Slice: NPI for Slicing Service	15
6.2.	NPI.SR-Policy: NPI for SR-Policy Service	16
6.3.	NPI.Mcast: NPI for Multicast Services	16

6.4.	NPI.TI-LFA: NPI for TI-LFA Services	16
6.5.	NPI.DSCP: NPI for Diffserv Services	16
6.6.	NPI Syntax and Operation	16
7.	Applicability of NPI	17
8.	Security Considerations	18

9.	IANA Considerations	19
10.	Acknowledgements	19
11.	References	19
11.1.	Normative References	19
11.2.	Informative References	20
	Author's Address	21

[1.](#) Introduction

This document describes a framework and a detailed suite of network programming interface (NPI) examples for provisioning of underlay services to overlay networks. It provides background by reviewing the growing pains that commonly faced today by enterprise for its flexible WAN sites connection. It assumes that WAN connection services are and will continue to be provided by multiple underlay networks operated by different administrative entities. Based on the pains and the assumptions, this document propose to use SRv6 binding SIDs (BSIDs) on a transport network (TN) edge router as NPI for service provisioning to be accessed remotely and securely by a customer router that constitutes to a higher overlay network like SD-WAN or CDN, including the requirements of such service, the illustration of how it may work, and the possible applicability of the solution.

[2.](#) Terms and Abbreviations

The following abbreviations are used in this document.

- * AN: Access Network
- * TN: Transport Network
- * VN: Virtual Network
- * VTN: Virtuan Transport Network

- * TSN: Time-Sensitive Networking
- * MAN: Metro Area Network
- * WAN: Wide Area Network
- * SD-WAN: Software-defined Wide Area Network
- * CDN: Content Distribution Network
- * VPN: Virtual Private Network

Xie

Expires 7 September 2022

[Page 3]

Internet-Draft Network Programming Interface using SRv6

March 2022

- * VRF: Virtual Routing and Forwarding
- * PE: Provider Edge
- * CPE: Customer Premises Equipment
- * SR: Segment Routing
- * MPLS: Multi-Protocol Label Switching
- * LDP: Label Distribution Protocol
- * BGP: Border Gateway Protocol
- * SRv6: Segment Routing over IPv6
- * BSID: Binding Segment Identifier
- * NBI: NorthBound Interface
- * NSP: Network Service Provider
- * ISP: Internet Service Provider
- * IXP: Internet eXchange Provider
- * SRH: Segment Routing Header
- * TA: Transmission Association

- * TAB: Transmission Association Database
- * TP: Transmission Policy
- * TPB: Transmission Policy Database
- * NPI: Network Programming Interface

3. Background and the Problem

3.1. Internet and Transport Network for WAN Transit

The following figure demonstrates a network architecture that an enterprise commonly uses for its WAN sites connection. On one site, a network service provider (NSP) may provide Internet access and Transport Network (TN) access to the same customer. On the other site, NSP2 provides the Internet access (act as Internet Service Provider or ISP), and NSP1 provides the TN access (act as NSP).




Figure 1: Multiple Transit Network

Inside the TN, there are many services provided, like Slice [[I-D.ietf-teas-ietf-network-slices](#)], SR-Policy [[I-D.ietf-spring-segment-routing-policy](#)], Multicast VPN [[RFC6513](#)],

TI-LFA [[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)], or simple Diffserv using DSCP. They can each be deployed with different implementations but provide the same or similar functions to its user instance.

Take Slice as an example, the network slices can be realized in different technologies, including physical connections, MPLS, time-sensitive networking (TSN), Flex-E, virtual transport network (VTN), or virtual network (VN) for independent transportation on the same infrastructure. This document uses the VTN to represent the TN (default slice) or a (non-default) Network Slice in the TN.

[3.2.](#) Separation of Transport Network and Access Network

The following figure demonstrates a network architecture that further separates Transport network (TN) and Access Network (AN). In this figure, the network operator of AN, TN and Internet can be different from each other. The physical connection between AN and Internet/TN is pre-built by the two operators. The connection between AN and CPE is provisioned on-demand. In some scenarios, the AN can be an Internet exchange provider (IXP) independent of ISP and NSP. In some other scenarios, the AN can be an ISP that runs Internet backbone as well.

Once a CPE is connected to AN, a customer wants to select its WAN transit service between site of CPE1 to site of CPE2 in a flexible way. For example, CPE1 may want to use the following policy for its different flows:

- * For flows to CPE2 with characteristics X, use TN (default slice) to transmit.
- * For flows to CPE2 with characteristics Y, use TN (non-default) slice Y to transmit.
- * For other flows to CPE2, use Internet to transmit.

In some circumstances, an enterprise may want even more flexible way as its flow policy in a timely manner:

- * Test the quality of some type of TN service before it use this TN service in a short time.
- * Use some type of TN service for flows with characteristics Z for some time (one month for example).
- * Use another type of TN service for flows with characteristics Z for some other time (another one month for example).




Figure 2: Seperation of AN and TN

3.3. The Problem

In above figure, PE2 and CPE2 are connected through an AN. Assuming that AN is an Layer-3 network serving in a local or metro area network (MAN). This is different from the previous case. CPE and TN are directly-connected in previous case and non-directly connected in

this case.

Conventional approach provided by VPN technology [[RFC4364](#)] depends heavily on a physical or logical port-based "VRF attachment circuit".

A physical "VRF attachment circuit" means that a directly-connected CPE1 (as described in previous section) is deployed, which needs heavy manual work and costs a fairly long time.

A logical "VRF attachment circuit" means that a set of joint segments, first a physical CPE-to-VRF attachment circuit from CPE1 to AN (west), then a pseudo-circuit like PW (Pseudo-Wire) service from AN (west) to AN (east), and last a physical VRF-to-VRF attachment circuit between AN (east) and PE1 (Each will treat the other as if it were a CPE router). In this approach, adding or removing port-based VRF attachment circuit(s) between AN and TN is necessary when a customer requires to use the TN service in WAN area, and this usually means collaboration and cost between AN operator and TN operator. It also means that, a customer who wants to use various transport services of TN has to depend additionally on AN. This approach is practical when AN operator and TN operator are the same NSP, but may be difficult for deployment when AN and TN are different NSPs. In a cloud environment, the CPE may be a virtual CPE (vCPE) located in Cloud DC and there is no direct connection between the vCPE to the PE, and the challenge is also described in [[I-D.ietf-rtgwg-net2cloud-problem-statement](#)].

The general problem is that, the various services in a TN is not accessible to a CPE not having a direct connection to the PE of TN flexibly.

[4.](#) Architecture and Design Consideration of NPI

[4.1.](#) Concept and Architecture of NPI

This section defines general concept and architecture of the solution.

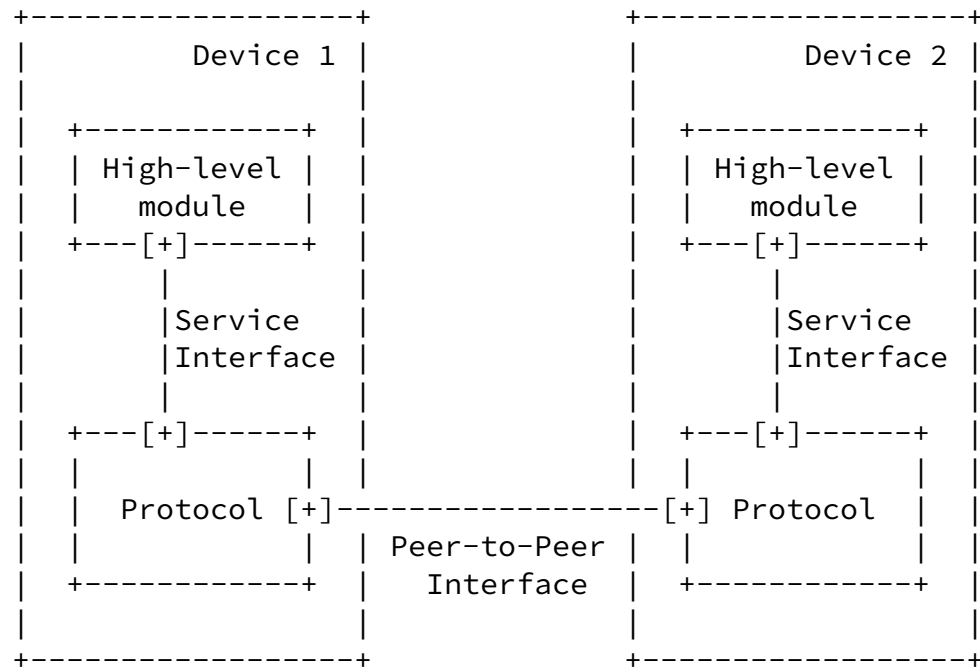


Figure 4: Layering Mode of Protocol

In the figure of NPI Reference Mode, the TN as a whole provides the operations that a device in a higher layer network (the overlay network) can invoke, through a message or packet between remote peers (CPE1 and PE1 in the diagram). The layering and peering relation between underlay network and overlay network is comparable to the "Service Interface" and "Peer-to-peer Interface" respectively. The "Network Programming Interface" provided to a higher overlay network is comparable to the "Application Programming Interface" as well. NPI provides a syntax and message specification by which the transport service can be invoked by overlay network, and the implementation of the NPI is responsible for mapping the tangible set of operations and objects defined by the NPI onto the abstract set of service defined by an underlay network. The flexible use of the NPI by overlay network is the so-called "programming" or "software-define".

[4.2.](#) Transmission Association and Transmission Policy

This section defines management framework for implementations and invokers of NPI between PE and CPE. The concept of a "Transmission Association" (TA) and "Transmission Policy" (TP) is fundamental to NPI.

Transmission Association (TA) is an object that defines the transmission characteristics that can be offered by TN operator to a

customer. In the NPI architecture, VTN is used to represent an example of the Transmission characteristics, and VRF is used to

represent a customer, and SRv6 BSID is used as the TA object. TA is defined in PE of a TN, and all the TAs for all customers defined in a PE constitutes the transmission association database (TAD).

Transmission Policy (TP) is a mapping from a set of flows of a customer to a specific TA object. In the NPI architecture, TP is defined in CPE and all the TPs defined in a CPE constructs the Transmission policy database (TPD).

Implementations of NPI should include implementation on PE and implementation on CPE separately. Implementation on PE (of underlay TN) should implement TA and TAD, and the TAD should support multiple customers, meaning that VRF routing and forwarding should be built for each customer. Implementation on CPE (of overlay) should implement TP and TPD. The TPD on a CPE should be able to update according to the change of TAs available on it.

[4.3.](#) Data-plane Consideration

The problem to be solved is how to access of Underlay network service by a non-directly connected (or remote) CPE belonging to an overlay network. There is no physical connection (and its secure mechanism) between PE1 and CPE1, nor logical connection like E-Line (and its secure mechanism) between PE1 and CPE1.

SRv6 Binding SID (BSID) provides a potential solution for this problem for its "binding/mapping" and "reachability" capabilities intrinsically.

SRv6 Binding SIDs (BSIDs) are configured on PE1, advertised to AN, and accessed by CPE1 using IPv6 routing capability. When the CPE1 sends a packet (customer packet, c-packet) using SRv6 with the active segment being the BSID, PE1 transits the packet to PE2 using the transport service (e.g., network slice) that is bound to the BSID. In a common case, it means a proper encapsulation of the c-packet by PE1 to use the VTN bound to the BSID.

Multiple BSIDs can be configured on PE1 and each BSID is mapped to a different services. Whenever a new VTN service is desired by the

customer, the network operator of TN only need to expose an additional BSID to the customer.

As an example, the CPE encapsulates a received packet with an outer IPv6 header where SRv6 BSID is in the destination address field to indicate the VTN it uses, followed by an SRH to contain egress CPE's SID as final destination.

[4.4.](#) Control-plane Consideration

PE needs to advertise the TA (SRv6 BSID) and its transmission characteristics to CPE. Thus the CPE could know the available TAs provided by a PE, and build or update its TPD. On the other hand, the PE need to know the IPv6 prefixes relevant to the transmission for a customer, and build or update the routing and forwarding information for the VRF of the customer. For both case, the control-plane Interface could use routing protocol like BGP, or some other means like NorthBound Interface, and the detailed control-plane is outside the scope of this document.

[5.](#) Requirements of NPI

[5.1.](#) Scalable for Multiple Tenants

A typical example, the transport network (TN) provides transmission association (TA) as NPI for a customer to access its VTNs. Each TA is mapped from an SRv6 BSID to a VRF (representing the customer) and a VTN (representing the transmission characteristics). Multiple SRv6 BSIDs can serve different customers. For example, these SRv6 SIDs can be configured under an SRv6 Locator for ease of management. The SRv6 locator can be advertised to AN as an aggregated prefix to make it routable. Locally, the SRv6 locator can be bound to a VRF and thus separate it from other VRF or default-VRF (public VRF).

Following is an example on PE1:

- * PE1 connect to AN using an interface that is bound to VRF-z.
- * SRv6 BSID-11/12/13 are bound to customer-1 (VRF-x) VTN-1/2/3 respectively.

- * SRv6 BSID-21/22/23 are bound to customer-2 (VRF-y) VTN-1/2/3 respectively.

In this example, BSID-11/12/13 and BSID-21/22/23 are allocated from an SRv6 locator, and the SRv6 locator is advertised to AN and is routable through AN. Locally the SRv6 locator and its SIDs are setup in VRF-z context. When an PE1 receives an IPv6 packet from the interface connected to AN, it performs longest-prefix-match lookup on the packet's destination address in VRF-z table. The lookup will return one of the following:

- * A FIB entry that represents (BSID-11, VRF-x, VTN-1)
- * A FIB entry that represents (BSID-12, VRF-x, VTN-2)

- * A FIB entry that represents (BSID-13, VRF-x, VTN-3)
- * A FIB entry that represents (BSID-21, VRF-y, VTN-1)
- * A FIB entry that represents (BSID-22, VRF-y, VTN-2)
- * A FIB entry that represents (BSID-23, VRF-y, VTN-3)

PE1 then performs a longest-prefix-match lookup on the next SID in SRH (an SID of CPE2), in VTN-x or VTN-y table, to determine the next hop (tunnel endpoint on TN). In the example, the next hop is PE2. PE1 further performs the transmission of the packet from PE1 to PE2 over VTN-x or VTN-y, and this usually includes the necessary encapsulation.

In this way, the NPI architecture can support multiple tenants and multiple services in a scalable way.

[5.2.](#) Support for Common Service

[5.2.1.](#) Slice Service

Slice service is an emerging service that NSPs are deploying for the 5G and other wider use cases. The NPI should support Slice service

to be accessible for overlay networks to invoke.

The SRv6 Binding SID for NPI can be independent to any VTN implementation inside the TN. For example, the VTN can use MPLS encapsulation, SR-MPLS or SRv6 encapsulation, or SRv6/SR-MPLS encapsulation with any kind of additional info like SR/SRv6 policy, SR/SRv6 slicing, Flex-E, etc.

[5.2.2.](#) SR-Policy Service

SR-Policy service is widely implemented and deployed in NSPs. The NPI should support SR-Policy service to be accessible for overlay networks to invoke.

The SRv6 BSID for NPI can be independent to any SR-Policy implementation inside the TN. For example, the SR-Policy can use either SR-MPLS encapsulation SRv6 encapsulation for its SR-Policy inside the TN, but the NPI for the overlay network to invoke is through SRv6 BSID NPI.

[5.2.3.](#) Multicast Service

It is rare for Internet (ISP) to provide multicast service, but is common for a TN to support multicast service (e.g., MVPN service). The SRv6 BSID service should support, or be able to extend to support multicast service provided by TN. The multicast service provided by TN can use any data plane or control plane. For example, the data plane can be MPLS, GRE or IP, and the control plane can be NG-MVPN [[RFC6513](#)].

[5.2.4.](#) Other Services

Some other services like TI-LFA, or simple DSCP/QoS can also be considered as a service that could be provided by TN to Overlay networks, determined by market and commercial design between TN operators and Overlay network. In such case, the NPI should be defined for Overlay network to invoke.

[5.3.](#) Cost-Effective Encapsulation

The SRv6 Binding SID for NPI is used for customer CPE to identify the transport service it desired for a packet, and is in the destination address of the packet as active segment. The real destination address of the packet in the overlay is the last segment in the SRH. When the PE1 receives a packet with the destination address being an SRv6 BSID mapping to a VTN, PE1 gets the BSID semantic from the packet and transit the packet through the VTN that is bound to the SRv6 BSID. Since the real destination address and customer (VRF) of the packet has been obtained, the SRH should be delete before it is encapsulated and transmitted through the TN, in case there is no additional SID(s) other than the BSID and the real destination. A Delete On-demand (DoD) flavor may be used for this purpose.

[5.4.](#) Secure for Remote Accessing

Traditionally, the security of accessing an TN service is guaranteed by a strictly managed (configured) object between CPE and PE named "VRF attachment circuit" [[RFC4364](#)]. The concept of "VRF accachment circuit" makes the accessing of Provider Edge router a "local" behavior with strict and static configuration (e.g., physical wire, or/and E-Line). Such kind of strict amd static configuration per-site per-customer is the pain for a flexiable and instant accessing of TN WAN transport service, and is the problem to be solved.

In the NPI concept, the "VRF attachement circuit" is replaced by the SRv6 BSID which is to be accessed remotely and securely by a customer router. It needs to provide the same security level as described as in [section 4.5 of \[RFC3809\]](#). Particularly, it needs to ensure that

the SRv6 SID allocated for a customer is accessed exactly by the CPE belonging to the customer and no spoofed CPE could access the SRv6 SID. A candidate solution is to use SRH HMAC defined in [[RFC8754](#)]. This solution ensures the SRv6 SID accessiable only by the serving customer holding correct HMAC key(s) in a per-customer granularity.

[5.5.](#) Manageability

Use of anycast SRv6 SID in every PE of multiple TN sites can provide the same access identifier for a customer. This means that, the SRv6 SID for NPI is allocated and managed per-customer and per-VTN.

Anycast SRv6 SID as a NPI can provides support when vCPE migrate from one Data center to another center when vCPE uses a hard-coded "IPv6 address" to invoke the NPI. Alternatively, the DNS system can be aided for the use of the NPI and provids the migration feature for vCPE.

6. Examples and Illustrations of NPI

Following figure is a reference to illustrates the NPI examples.

```

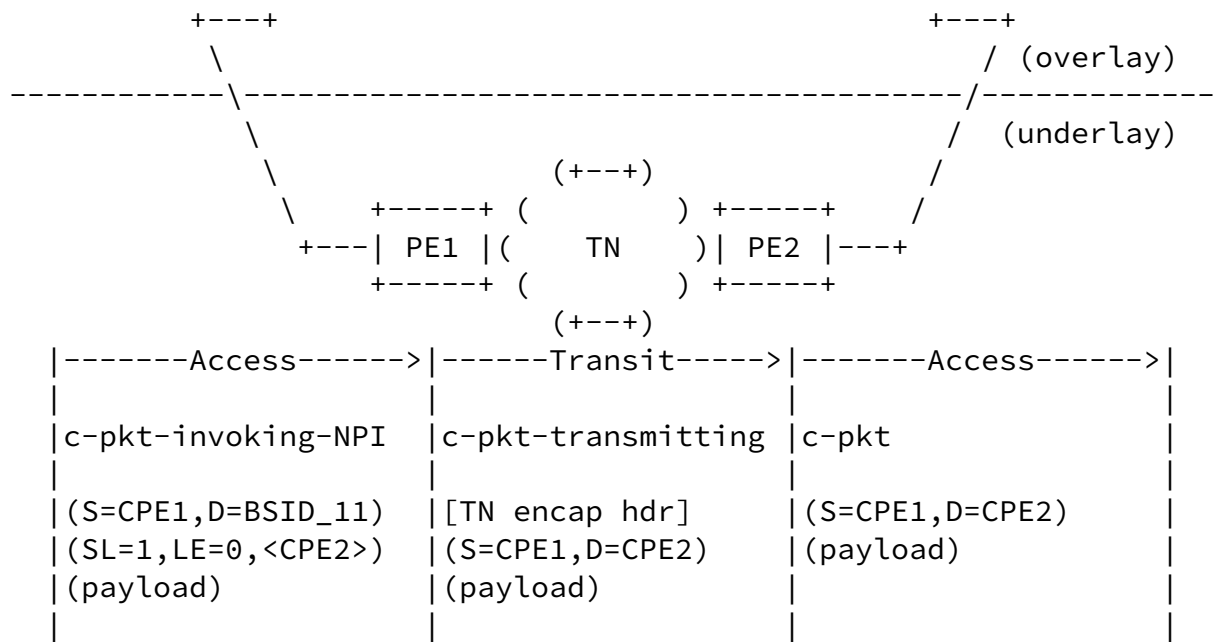
      +----+
+----+ (    )
|CPE1|---( AN  )
+----+ (    )

```

```

      +----+
(    ) +----+
( AN  )---|CPE2|
(    ) +----+

```

c-pkt-invoking-NPI: Customer Packet invoking NPI (using SRv6 BSID)
c-pkt-transmitting: Customer Packet Transmitting through TN
c-pkt: Customer Packet

Figure 5: Reference Figure of NPI Examples

6.1. NPI.Slice: NPI for Slicing Service

PE1 receives a packet, and the destination address is an SRv6 BSID bound to VTN-x, and the only left SID in SRH (with HMAC) is the final destination of the overlay network - CPE2. PE1 deletes the SRH on-demand, and transmits the packet using VTN-x. The "TN encap hdr" in the reference mode is a kind of encapsulation that is corresponding to VTN-x.

The TN encapsulation header corresponding to VTN-x can be based on MPLS, SR-MPLS, SRv6 or any other kind of encapsulation to implement the VTN-x transmitting.

PE1's interface connecting AN is bound to a VRF (VRF-z), and the SRv6 locator containing BSID_11 is also bound to the same VRF.

For ease of use by CPE2 of the same customer, PE2 has the same SRv6 BSID, we call anycast BSID, for the same purpose of using VTN-x to transmit a packet from CPE2.

[6.2.](#) NPI.SR-Policy: NPI for SR-Policy Service

PE1 receives a packet, and the destination address is an SRv6 BSID bound to an SR-Policy, and the only left SID in SRH (with HMAC) is the final destination of the overlay network - CPE2. PE1 deletes the SRH of the received packet, and re-encapsulates the packet with the SR-Policy header and transmitted through the TN to PE2. The SR-Policy can be SR-MPLS or SRv6 data plane. Comparing to the End.BM and End.B6.Encaps/End.B6.Encaps.Red defined in [[RFC8986](#)], a delete on-demand (DoD) flavor is needed to delete the SRH in the received packet.

[6.3.](#) NPI.Mcast: NPI for Multicast Services

PE1 receives a packet, and the destination address is an SRv6 BSID bound to a multicast group, and there is SRH (with HMAC) with no left SID (SL=0 and LE=0). PE1 deletes the SRH of the received packet, get the multicast group that is bound to the BSID, and replace the destination address with the multicast group address. Further, PE1 re-encapsulates the packet with proper P2MP header like MPLS P2MP label-stack or mGRE header. The packet then is transmitted through the TN to the PEs that receives multicast join of the multicast group.

On the receiver site, PE2 get a multicast packet by decapsulating the received packet. Further, PE2 transforms the destination address to unicast address of CPE2 and send the packet to CPE2. There could be some other alternatives for receiver site and the detail is outside the scope of this document.

[6.4.](#) NPI.TI-LFA: NPI for TI-LFA Services

PE1 receives a packet, and the destination address is an SRv6 BSID bound to null, but the TN has enabled TI-LFA as default service to provide 50ms failover guarantee. PE1 simply transfer the packet through the TN to PE2 with the TI-LFA guaranteed 50ms failover.

[6.5.](#) NPI.DSCP: NPI for Diffserv Services

PE1 receives a packet, and the destination address is an SRv6 BSID bound to a DSCP value representing a diffserv service in TN. PE1 simply transfer the packet through the TN to PE2 with the proper DSCP/TOS/Traffic-Class value used for a specified queue.

[6.6.](#) NPI Syntax and Operation

The following is the NPI Syntax and Operation of the NPI examples.

NPI Example	Message Syntax	NPI Operation
NPI.Slice	(DA=BSID) (SL=1,LE=0,FDA)	Transport packet in a Slice bound to BSID to FDA (Final Destination Address)
NPI.SRplcy	(DA=BSID) (SL=1,LE=0,FDA)	Transport packet in an SRpolicy bound to BSID to FDA
NPI.Mcast	(DA=BSID) (SL=0,LE=0,HMAC)	Transport packet to sites joined in a multicast group that is bound to BSID
NPI.TILFA	(DA=BSID) (SL=1,LE=0,FDA)	Transport packet to a site with 50ms protection guaranteed by TI-LFA
NPI.DSCP	(DA=BSID) (SL=1,LE=0,FDA)	Transport packet to FDA using DSCP bound to BSID through the TN

Figure 6: NPI Syntax and Operation

The list is not exhaustive. In practice, any service provided by a TN can have a NPI using SRv6 SID for accessing remotely and securely in this framework.

7. Applicability of NPI

1. The solution provides a new approach to make various service accessible to customers that need flexible selection between Internet/VTNs/Slices.

It provides a self-service mode for each customer to select service by using the SRv6 BSID in the packet as its "transmitting request". The TN edge router does not need to configure or reconfigure its "steering policy" on a per-flow basis. Each customer can have very flexible policy to determine which flows to use which service in which time slot. The TN is not aware of the customer's policy (TP) and its timely change, but only behaves according to the active SRv6 BSID of a packet (TA).

By contrast, a northbound interface (NBI) is likely to be necessary in conventional approach. A customer expresses requirements for a particular service (which flows use which service) and the TN need to dynamically change its configuration when the requirements changes.

Note that customers accessing the VTNs don't need to have a "directly-connected" line to the Transport network. Any customer connected to AN (usually part of Internet Service Provider, ISP) can access high-quality TN transmission service in a software-define way.

2. This solution provides a separation of AN and TN with very low dependency of each other.

It encourages each operator to serve customer's requirement by using its network advantage. AN operator (local operator) can use its small-area but high-bandwidth network to get more customer connected, and TN operator (WAN operator) can use its large-area network with low-latency and/or ultra-reliable capabilities to get more customer to use.

3. SRv6 BSID NPI can serve as a solution for inter-connecting distributed Cloud/Edge-cloud/NFV in WAN scope as well as On-premise enterprise site.

Using the ubiquitous AN and Internet connection, distributed Cloud/Edge-cloud/NFV in WAN scope could easily be connected for its applications initially. Using the VTNs provided by TN operator, advanced requirements such as ultra-reliable and/or low-latency communication in WAN scope could be met and thus enable the applications to promote their performance and user experience incrementally.

4. This solution provides an alternative evolution path for MPLS network.

The SRv6 BSID can be deployed on edge routers of an MPLS (conventional LDP/RSVP-TE MPLS or SR MPLS) network, and the MPLS data plane can be used as the encapsulation for many value-added services

like VTN, SR-Policy, Multicast, etc.

Technologies including physical connections, time-sensitive networking (TSN), Flex-E, etc can also be used in TN, and the solution can combine with these technologies, and make these technologies accessible through the SRv6 BSID Service.

[8.](#) Security Considerations

TBD.

Xie

Expires 7 September 2022

[Page 18]

Internet-Draft Network Programming Interface using SRv6

March 2022

[9.](#) IANA Considerations

Allocation is expected from IANA for implementation of the NPIs from the SRv6 Endpoint Behaviors Registry.

Value	Endpoint Behavior	Reference
TBD	NPI.Slice with DoD Flavor	This draft
TBD	NPI.SRplcy with DoD Flavor	This draft
TBD	NPI.Mcast with DoD Flavor	This draft
TBD	NPI.TILFA with DoD Flavor	This draft
TBD	NPI.DSCP with DoD Flavor	This draft

[10.](#) Acknowledgements

TBD.

[11.](#) References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](#), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", [RFC 8986](#), DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

Xie

Expires 7 September 2022

[Page 19]

Internet-Draft Network Programming Interface using SRv6

March 2022

11.2. Informative References

- [I-D.ietf-rtgwg-net2cloud-problem-statement]
Dunbar, L., Consulting, M., Jacquenet, C., and M. Toy, "Dynamic Networks to Hybrid Cloud DCs Problem Statement", Work in Progress, Internet-Draft, [draft-ietf-rtgwg-net2cloud-problem-statement-11](#), 26 July 2020, <<https://www.ietf.org/archive/id/draft-ietf-rtgwg-net2cloud-problem-statement-11.txt>>.
- [I-D.ietf-rtgwg-segment-routing-ti-lfa]
Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, [draft-ietf-rtgwg-segment-routing-ti-lfa-08](#), 21 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-rtgwg-segment-routing-ti-lfa-08.txt>>.
- [I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, [draft-ietf-spring-segment-routing-policy-18](https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-18), 17 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-18.txt>>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, [draft-ietf-teas-ietf-network-slices-08](https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-08), 6 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-08.txt>>.

[RFC3809] Nagarajan, A., Ed., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", [RFC 3809](https://www.rfc-editor.org/info/rfc3809), DOI 10.17487/RFC3809, June 2004, <<https://www.rfc-editor.org/info/rfc3809>>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](https://www.rfc-editor.org/info/rfc4364), DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

[RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", [RFC 6513](https://www.rfc-editor.org/info/rfc6513), DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.

Xie

Expires 7 September 2022

[Page 20]

Internet-Draft Network Programming Interface using SRv6

March 2022

Author's Address

Jingrong Xie
Huawei Technologies
Email: xiejingrong@huawei.com

