Workgroup: TEAS Working Group
Internet-Draft:
draft-ietf-teas-nrp-scalability-01
Published: 24 October 2022
Intended Status: Informational
Expires: 27 April 2023
Authors: J. Dong
Huawei Technologies
Z. Li
Huawei Technologies
L. Gong
China Mobile
G. Yang
China Telecom
J. Guichard
Futurewei Technologies
G. Mishra
Verizon Inc.
F. Qin
China Mobile
T. Saad
Juniper Networks
V. Beeram
Juniper Networks

Scalability Considerations for Network Resource Partition

Abstract

The IETF Network Slice aims to offer a connectivity service to a network slice customer with specific Service Level Objectives (SLOs) and Service Level Expectations (SLEs) over a common underlay network. A Network Resource Partition (NRP) is a set of network resources that are allocated from the underlay network to carry a specific set of network slice service traffic and meet specific SLOs and SLEs.

As the demand for IETF Network Slice increases, scalability would become an important factor for the deployment of IETF Network Slices. Although the scalability of IETF Network Slices can be improved by mapping a group of IETF Network Slices to one NRP, that design may not be suitable or possible for all deployments, thus there are concerns about the scalability of NRPs.

This document discusses some scalability considerations about NRPs in the network control and data plane. It also investigates a set of optimization mechanisms.

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."

This Internet-Draft will expire on 27 April 2023.

Table of Contents

1. Introduction

The IETF Network Slice aims to offer a connectivity service to a network slice customer with specific Service Level Objectives (SLOs)

and Service Level Expectations (SLEs) over a common underlay network. [I-D.ietf-teas-ietf-network-slices] defines the terminologies and the characteristics of IETF Network Slices. It also discusses the general framework, the components and interfaces for requesting and operating IETF Network Slices. For the realization of IETF Network Slices, the concept called Network Resource Partition (NRP) is introduced by [I-D.ietf-teas-ietf-network-slices]. An NRP is a collection of network resources in the underlay network, which can be used to ensure the requested SLOs and SLEs of IETF Network Slices services are met.

[I-D.ietf-teas-enhanced-vpn] describes a layered architecture and the candidate technologies in different layers for delivering enhanced VPN (VPN+) services. VPN+ aims to meet the needs of customers or applications which require connectivity services with advanced characteristics, such as the assurance of Service Level Objectives (SLOs) and specific Service Level Expectations (SLEs). VPN+ services can be delivered by mapping one or a group of overlay VPNs to a virtual underlay network which is allocated with a set of network resources. The VPN+ architecture and technologies could be used for the realization of IETF Network Slices. And in the context of network slicing, NRP could be used to instantiate the virtual underlay network construct in VPN+.

As the demand for IETF Network Slice services increases, scalability (the number of network slices a network can support) would become an important factor for the deployment of IETF Network Slices in specific networks. Although the scalability of IETF Network Slices can be improved by mapping a group of IETF Network Slices to one NRP, that design may not be suitable or possible for all deployments, thus there are concerns about the scalability of NRPs.

This document discusses some scalability considerations about NRPs in the network control and data plane. It also investigates a set of optimization mechanisms.

2. Network Resource Partition Scalability Requirements

As described in [I-D.ietf-teas-ietf-network-slices], the connectivity constructs of IETF Network Slices may be grouped together according to their characteristics (including SLOs and SLEs) and mapped to a given NRP. The grouping and mapping of IETF Network Slices are policy-based and under the control of operator. For example, a policy can be considered by an operator to host a large number of IETF Network Slices into a relatively small number of NRPs to reduce the amount of state information to be maintained in the network. This can help to avoid the maintenance of per IETF

Network Slice state in the underlay network, which is equivalent to the one-to-one mapping between IETF Network Slices and NRPs.

With the introduction of various services which require enhanced connectivity, it is expected that the number of IETF Network Slices will increase. The potential number of IETF Network Slices and the underlying NRPs are estimated by classifying the network slice deployment into three typical scenarios:

1. IETF Network Slices can be used by a network operator to deliver different types of services. For example, in a multi-service network, different IETF Network Slices can be created to carry, e.g., mobile transport services, fixed broadband services, and enterprise services respectively: each type of service could be managed by a separate team. Some other type of services, such as multicast services, may also be deployed in a separate virtual underlay network. Then a separate NRP may be created for each service type. It is also possible that a network infrastructure operator provides IETF Network Slice services to other network operators as wholesale services, and an NRP may also be needed for each wholesale service operator. In this scenario, the number of NRPs in a network could be relatively small, such as in the order of 10 or so.

2. IETF Network Slice services can be requested by customers of industrial verticals, where the assurance of SLOs and the fulfilment of SLEs are contractually defined between the customer and the slice service provider, possibly including financial penalties in case the service provider fails to honor the contract (SLO or SLE). At the early stage of the vertical industrial deployment, a few customers in some industries will start using IETF Network Slices to address the connectivity requirements and performance assurance raised by their business, such as smart grid, manufacturing, public safety, on-line gaming, etc. The realization of such IETF Network Slices may require the provision of different NRPs for different industries, and some customers may require dedicated NRPs for strict service performance guarantees. Considering the number of vertical industries and the number of customers in each industry, the number of NRPs needed may be in the order of 100.

3. With the advent of 5G and cloud networks, IETF Network Slices services could be widely used by customer of various vertical industries and enterprises who require guaranteed or predictable network service performance. The amount of IETF Network Slices may increase to the order of thousands or more. Accordingly, the number of NRPs needed may be in the order of 1000.

In [TS23501], the 3GPP defines a 32-bit identifier for a 5G network slice with an 8-bit Slice/Service Type (SST) and a 24-bit Slice Differentiator (SD). This allows mobile networks (the RAN and mobile core networks) to potentially support a large number of 5G network slices. It is likely that multiple 5G network slices may be mapped to the same IETF Network Slice, but in some cases (for example, for specific SST or SD) the mapping may be closer to one-to-one. This may require increasing number of IETF Network Slices, the number of required NRPs may increase as well.

Thus the question of scalable IETF network slice services arises. Mapping multiple IETF Network Slices to the same NRP presents a significant scaling benefit, while a large number of NRPs may also be required, which raises scalability challenges too.

3. Network Resource Partition Scalability Considerations

This section analyses the scalability of NRPs in the control plane and data plane to understand the possible gaps in meeting the scalability requirements of IETF Network Slices.

3.1. Control Plane Scalability

The control plane for establishing and managing NRPs could be based on the combination of a centralized controller and a distributed control plane. The following subsections consider the scalability property of both the distributed and centralized control plane in such design.

3.1.1. Distributed Control Plane

In some networks, multiple NRPs may need to be created for the delivery of IETF Network Slice services. Each NRP is associated with a logical topology. The network resource attributes and the associated topology information of each NRP may need to be exchanged among the network nodes. The scalability of the distributed control plane used for the distribution of NRP information needs to be considered from the following aspects:

* The number of control protocol instances maintained on each node

* The number of control protocol sessions maintained on each link

* The number of control messages advertised by each node

* The amount of attributes associated with each message

* The number of computations (e.g., SPF computation) executed by each node

As the number of NRPs increases, it is expected that at least in some of the above aspects, the overhead in the control plane may increase in proportion to the number of the NRPs. For example, the overhead of maintaining separate control protocol instances (e.g., IGP instances) for each NRP is considered higher than maintaining the information of multiple NRPs in the same control protocol instance with appropriate separation, and the overhead of maintaining separate protocol sessions for different NRPs is considered higher than using a shared protocol session for exchanging the information of multiple NRPs. To meet the scalability and performance requirements with increasing number of NRPs, it is suggested to select the control plane mechanisms which have better scalability while can still provide the required functionality, isolation and security for the NRPs.

3.1.2. Centralized Control Plane

The use of centralized network controllers may help to reduce the amount of computation overhead in the distributed control plane, while it may also transfer some of the scalability concerns from network nodes to the network controllers, thus the scalability of the controller also needs to be considered.

A centralized controller can have a global view of the network, and is usually used for Traffic Engineering (TE) path computation with various constraints, or the global optimization of TE paths in the network. To provide TE paths computation and optimization for multiple NRPs, the controller needs to keep the topology and resource information of all the NRPs up-to-date. And for some events such as link or node failures, the resulting updates to the NRPs may need to be distributed to the controller in real time, and may affect the planning and operation of some NRPs. When there is a significant change in the network which impacts multiple NRPs, or multiple NRPs require global optimization concurrently, there may be a heavy processing burden at the controllers, and a large amount of signaling traffic to be exchanged between the controller and corresponding NRP components. These need to be taken into consideration from a scalability and performance standpoints.

3.2. Data Plane Scalability

To provide different IETF Network Slice services with the required SLOs and SLEs, it is important to allocate as many different subsets of network resources as there are different NRPs to avoid or reduce the risk of interference both between different IETF network slice services and between slice services and other services in the network. With both the use cases and the number of NRPs increases, it is required that the underlay network can provide a finer granularity of network resource partitioning for more IETF network

slice services, which means the amount of state about the partitioned network resources to be maintained on the network nodes is likely to increase.

IETF Network Slice service traffic needs to be processed and forwarded by network nodes according to a forwarding policy that is associated with the topology and the resource attributes of the NRP it is mapped to, this means that some fields in the data packet need to be used to identify the NRP and its associated topology and resources either directly or implicitly. Different approaches for encapsulating the NRP information in data packets may have different scalability implications.

One practical approach is to reuse some of the existing fields in the data packet to additionally indicate the NRP the packet belongs to. For example, destination IP address or MPLS forwarding label may be reused to identify the NRP. This avoids the complexity of introducing new fields in the data packet, while the additional semantics introduced to the existing fields may require additional processing. Moreover, introducing NRP-specific semantics to existing identifiers in the packet may result in the amount of the existing identifiers increasing in proportion to the number of the NRPs. For example, if IP address is reused to further identify an NRP, for a node which pariticipate in M NRPs, the amount of IP addresses needed for reaching this node in different NRPs would increase from 1 to M. This may cause scalability problems in networks where a relatively large number of NRPs is in operation.

An alternative approach is to introduce a new dedicated field in the data packet for identifying an NRP. And if this new field carries a network wide unique NRP identifier (NRP ID), it could be used together with the existing fields to determine the packet forwarding behavior. The potential issue with this approach lies in the difficulty of introducing a new field in some of data plane technologies.

In addition, the introduction of NRP-specific packet forwarding impacts the number of the forwarding entries maintained by the network nodes.

3.3. Gap Analysis of Existing Mechanisms

This section provides gap analysis of existing mechanisms which may be used to provide NRP identification in the data plane and the distribution of NRP related information using control plane protocols.

One existing mechanism of building NRPs is to use resource-aware Segment Identifiers (either SR-MPLS or SRv6)

[I-D.ietf-spring-resource-aware-segments] to identify the allocated network resources in the data plane based on the mechanisms described in [I-D.ietf-spring-sr-for-enhanced-vpn], and then distribute the resource attributes and the associated logical topology information in the control plane using mechanisms based on Multi-topology [I-D.ietf-lsr-isis-sr-vtn-mt] or Flex-Algo [I-D.zhu-lsr-isis-sr-vtn-flexalgo]. This mechanism is suitable for networks where a relatively small number of NRPs are needed. As the number of NRPs increases, there may be several scalability challenges with this approach:

1. The number of SR SIDs will increase in proportion to the number of NRPs in the network, which will bring challenges both to the distribution of SR SIDs and the related information in the control plane, and to the installation of forwarding entries for resource-aware SIDs in the data plane.

2. If each NRP is associated with an independent logical topology or algorithm, the number of route computations (e.g., SPF computations) will increase in proportion to the number of NRPs in the network, which may introduce significant overhead to the control plane of network nodes.

3. The maximum number of logical topologies supported by OSPF [RFC4915] is 128, the maximum number of logical topologies supported by IS-IS [RFC5120] is 4096, and the maximum number of Flexible Algorithms [I-D.ietf-lsr-flex-algo] is 128. Some of these technologies may not meet the required number of NRPs in some network scenarios.

4. Proposed Scalability Optimizations

To support more IETF Network Slice services while keeping the amount of network state at a reasonable scale, one basic approach is to classify a set of IETF Network Slice services (e.g., services which have similar service characteristics and performance requirements) into a group, and such group of IETF Network Slice services are mapped to one NRP, which is allocated with an aggregated set of network resources and the combination of the required logical topologies to meet the service requirements of the whole group of IETF Network Slice services. Different groups of IETF Network Slice services may be mapped to different NRPs, each of which is allocated with different set of network resources from the underlay network. According to operator's deployment policy, appropriate grouping of IETF Network Slice services and mapping them to a set of NRPs with proper network resource allocation could still meet the IETF Network Slice service requirements. However, in some network scenarios, such aggregation mechanism may not be applicable. The following sub-

sections proposes further optimization in control plane and data plane respectively.

4.1. Control Plane Optimization

4.1.1. Distributed Control Plane Optimization

Several optimization mechanisms can be considered to reduce the distributed control plane overhead and improve its scalability.

The first control plane optimization consists in reducing the amount of control plane sessions used for the establishment and maintenance of the NRPs. When multiple NRPs have the same connection relationship between two adjacent network nodes, it is proposed that one single control protocol session is used for these NRPs. The information specific to the different NRPs can be exchanged over the same control protocol session, with necessary identification information to distinguish the information of different NRPs in the control message. This could reduce the overhead of node in creating and maintaining a separate control protocol session for each NRP, and could also reduce the amount of control plane messages.

The second control plane optimization is to decouple the resource information of the NRP from the associated logical topology information, so that the resource attributes and the topology attributes of the NRP can be advertised and processed separately. In a network, it is possible that multiple NRPs are associated with the same logical topology, or multiple NRPs may share the same set of network resources hosted by a specific set of network nodes and links. With topology sharing, it is more efficient to advertise only one copy of the topology information, and multiple NRPs deployed over the very same topology could exploit such topology information. More importantly, with this approach, the result of topology-based route computation could also be shared by multiple NRPs, so that the overhead of per NRP route computation could be avoided. Similarly, for the resource sharing case, information about a set of network resources allocated on a particular network node or link could be advertised in the control plane only once and then be referred to by multiple NRPs which share that set of resource.

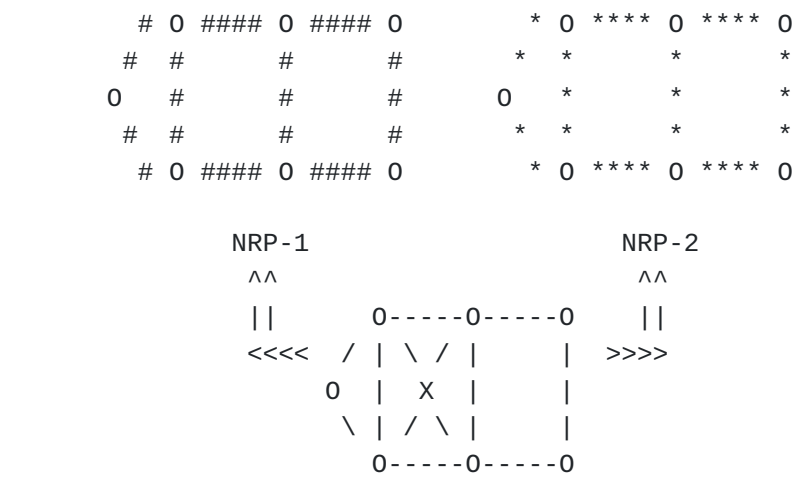

Legend

O	Virtual node
###	Virtual links with a set of reserved resources
***	Virtual links with another set of reserved resources

Figure 1. Topology Sharing between NRPs

Figure 1 gives an example of two NRPs that share the same logical topology. NRP-1 and NRP-2 are associated with the same logical topology, while the resource attributes of each NRP are different. In this case, the information of the shared network topology can be advertised using either MT or Flex-Algo, then the two NRPs can be associated with the same MT or Flex-Algo, and the outcomes of topology-based route computation can be shared by the two NRPs for further generating the corresponding NRP-specific routing and forwarding entries.

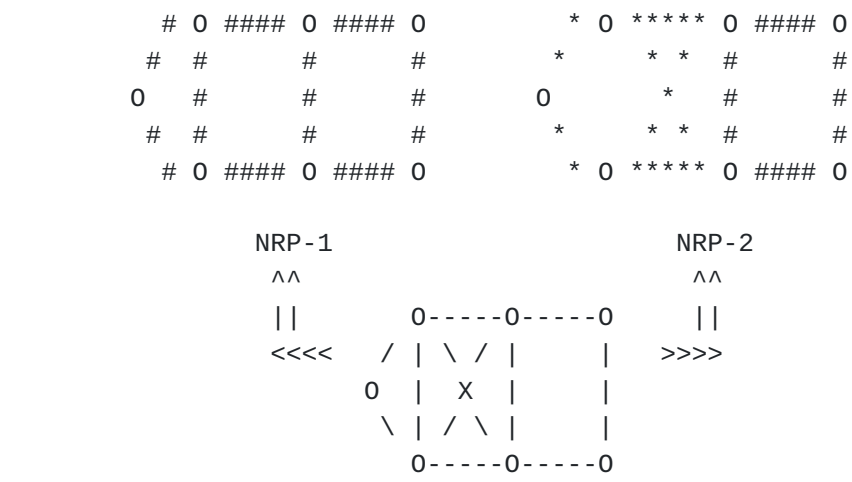

Underlay Network Topology

Legend

O	Virtual node
###	Virtual links with a set of reserved resource
***	Virtual links with another set of reserved resource

Figure 2. Resource Sharing between NRPs

Figure 2 gives another example of two NRPs which have different logical topologies, while they share the same set of network resources on a subset of the links. In this case, the information about the shared resources allocated on the those links needs to be advertised only once, then both NRP-1 and NRP-2 can refer to the common set of allocated link resource for constraint based path computation.

4.1.2. Centralized Control Plane Optimization

For the optimization of the centralized control plane, it is suggested that the centralized controller is used as a complementary computational facility to the distributed control plane rather than a replacement, so that the workload for NRP-specific path computation can be shared by both the centralized controller and the network nodes. In addition, the centralized controller may be realized with multiple network entities, each of which is responsible for one subset or region of the network. This is the typical approach for scale out of the centralized controller.

4.2. Data Plane Optimization

One optimization in the data plane consists in decoupling the identifiers used for topology-based forwarding from the identifier used for the NRP-inferred resource-specific processing. One possible mechanism is to introduce a dedicated network-wide NRP Identifier

(NRP ID) in the packet header to uniquely identify the set of local network resources allocated to an NRP on each participating network node and link for the processing of packets. Then the existing identifiers in the packet header used for topology based forwarding (e.g., destination IP address, MPLS forwarding labels) are kept unchanged. The benefit is the amount of the existing topology-specific identifiers will not be impacted by the increasing number of NRPs. Since this new NRP ID field will be used together with other existing fields of the packet to determine the packet forwarding behavior, this may require network nodes to maintain a hierarchical forwarding table in data plane. Figure 3 shows the concept of using separate data plane identifiers for topology-specific and resource-specific packet forwarding and processing purposes.

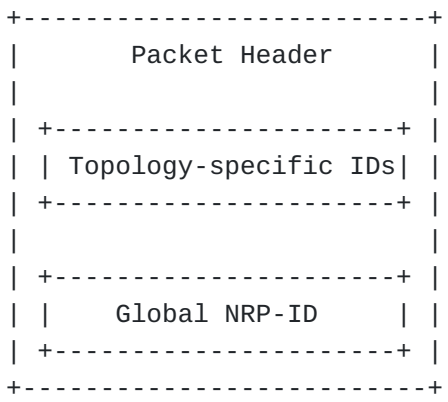

Figure 3. Decoupled Topology and Resource Identifiers in data packet

In an IPv6 [RFC8200] network, this could be achieved by introducing a dedicated field in either the IPv6 base header or the extension headers to carry the NRP ID for the resource-specific forwarding, while keeping the destination IP address field used for routing towards the destination prefix in the corresponding topology. Note that the NRP ID needs to be parsed by every node along the path which is capable of NRP-aware forwarding.
[I-D.ietf-6man-enhanced-vpn-vtn-id] introduces the mechanism of carrying the VTN resource ID (which is equivalent to NRP ID in the context of network slicing) in IPv6 Hop-by-Hop extension header.

In an MPLS [RFC3032] network, this may be achieved by inserting a dedicated NRP ID either in the MPLS label stack or a specific field that follows the MPLS label stack. Thus the existing MPLS forwarding labels are used for topology-specific packet forwarding purposes, and the NRP ID is used to determine the set of network resources for packet processing. This requires that both the forwarding label and the NRP ID are parsed by nodes along the forwarding path of the packet, and the forwarding behavior may depend on the position of the NRP ID in the packet. The detailed extensions to MPLS is

currently under discussion as part of the work conducted by the MPLS Open Design Team, and is out of the scope of this document.

5. Solution Evolution Perspectives

Based on the analysis provided by this document, the control and data plane for NRP need to evolve to support the increasing number of IETF Network Slice services and the increasing number of NRPs in the network. This section describes the foreseeable solution evolution taking the SR-based NRP solutions as an example, while the analysis and optimization in this document are generic and not specific to SR.

First, by introducing resource-awareness with specific SR SIDs [I-D.ietf-spring-resource-aware-segments], and using Multi-Topology or Flex-Algo mechanisms to define the logical topology of the NRP, providing a limited number of NRPs in the network is possible, and can meet the requirements for a relatively small number of IETF Network Slice services. This mechanism is called the "basic SR-based NRP".

As the required number of IETF Network Slice services increases, more NRPs may be needed, then the control plane scalability could be improved by decoupling the topology attributes from the resource attributes, so that multiple NRPs could share the same topology or resource attributes to reduce the overhead. The data plane can still rely on the resource-aware SIDs. This mechanism is called the "scalable SR-based NRP". Both the basic and the scalable SR-based NRP mechanisms are described in [I-D.ietf-spring-sr-for-enhanced-vpn].

Whenever the data plane scalability becomes a concern, a dedicated NRP ID can be introduced in the data packet to decouple the resource-specific identifiers from the topology-specific identifiers in the data plane, so as to reduce the number of IP addresses or SR SIDs needed in supporting a large number of NRPs. This is called the NRP-ID-based mechanism.

6. Operational Considerations

The instantiation of NRPs require NRP-specific configurations of the participating network nodes and links. There can also be cases where the topology or the set of network resources allocated to an existing NRP needs to be modified. Of course, the amount of configurations for NRP instantiation and modification will increase with the number of NRPs.

For the management and operation of NRPs and the optimization of paths within the NRPs, the status of NRPs needs to be monitored and

reported to the network controller. The increasing number of NRPs would require additional NRP status information to be monitored.

7. Security Considerations

This document discusses scalability considerations about the network control plane and data plane of NRPs in the realization of IETF Network Slice services, and investigates some mechanisms for scalability optimization. As the number of NRPs supported in the data plane and control plane of the network can be limited, this may be exploited as an attack vector by requesting a large number of network slices, which then result in the creation of a large number of NRPs.

One protection against this is to improve the scalability of the system to support more NRPs. Another possible solution is to make the network slice controller aware of the scaling constraints of the system and dampen the arrival rate of new network slices and NRPs request, and raise alarms when the thresholds are crossed.

The security considerations in [I-D.ietf-teas-ietf-network-slices] and [I-D.ietf-teas-enhanced-vpn] also apply to this document.

8. IANA Considerations

This document makes no request of IANA.

9. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

Hongjie Yang
Email: hongjie.yang@huawei.com

10. Acknowledgments

The authors would like to thank Adrian Farrel, Dhruv Dhody, Donald Eastlake, Kenichi Ogaki, Mohamed Boucadair, Christian Jacquenet and Kiran Makhijani for their review and valuable comments to this document.

11. References

11.1. Normative References

[I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+)", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-11, 19 September 2022, <https://

www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-11.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-15, 21 October 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-15.txt>.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <https://www.rfc-editor.org/info/rfc3032>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

11.2. Informative References

[I-D.ietf-6man-enhanced-vpn-vtn-id] Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network (VTN) Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-01, 11 July 2022, <https://www.ietf.org/archive/id/draft-ietf-6man-enhanced-vpn-vtn-id-01.txt>.

[I-D.ietf-lsr-flex-algo] Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-26, 17 October 2022, <https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-26.txt>.

[I-D.ietf-lsr-isis-sr-vtn-mt] Xie, C., Ma, C., Dong, J., and Z. Li, "Using IS-IS Multi-Topology (MT) for Segment Routing based Virtual Transport Network", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-sr-vtn-mt-03, 10 July 2022, <https://www.ietf.org/archive/id/draft-ietf-lsr-isis-sr-vtn-mt-03.txt>.

[I-D.ietf-spring-resource-aware-segments]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-06, 11 October 2022, <https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-06.txt>.

[I-D.ietf-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Segment Routing based Virtual Transport Network (VTN) for Enhanced VPN", Work in Progress, Internet-Draft, draft-ietf-spring-sr-for-enhanced-vpn-04, 11 October 2022, <https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-04.txt>.

[I-D.zhu-lsr-isis-sr-vtn-flexalgo] Zhu, Y., Dong, J., and Z. Hu, "Using Flex-Algo for Segment Routing (SR) based Virtual Transport Network (VTN)", Work in Progress, Internet-Draft, draft-zhu-lsr-isis-sr-vtn-flexalgo-05, 11 July 2022, <https://www.ietf.org/archive/id/draft-zhu-lsr-isis-sr-vtn-flexalgo-05.txt>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[TS23501] "3GPP TS23.501", 2016, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China

Email: jie.dong@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China

Email: lizhenbin@huawei.com

Liyan Gong
China Mobile
No. 32 Xuanwumenxi Ave., Xicheng District
Beijing
China

Email: gongliyan@chinamobile.com

Guangming Yang
China Telecom
No.109 West Zhongshan Ave., Tianhe District
Guangzhou
China

Email: yangguangm@chinatelecom.cn

James N Guichard
Futurewei Technologies
2330 Central Express Way
Santa Clara,
United States of America

Email: james.n.guichard@futurewei.com

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com

Fengwei Qin
China Mobile
No. 32 Xuanwumenxi Ave., Xicheng District
Beijing
China

Email: qinfengwei@chinamobile.com

Tarek Saad
Juniper Networks

Email: tsaad@juniper.net

Vishnu Pavan Beeram
Juniper Networks

Email: vbeeram@juniper.net