

TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: 11 August 2022

J. Dong
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
V. Beeram
Juniper Networks
7 February 2022

Scalability Considerations for Network Resource Partition
draft-dong-teas-nrp-scalability-01

Abstract

The IETF Network Slice service aims to meet the connectivity demands of 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 traffic and meet the required SLOs and SLEs. One or multiple IETF Network Slice services can be mapped to one NRP.

As the demand for IETF Network Slice services increases, scalability would become an important factor for the large scale 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, there are concerns about the scalability of NRPs. This document describes the scalability considerations about NRPs in the network control plane and data plane, and some optimization mechanisms are proposed.

Status of This Memo

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

Internet-Draft

NRP Scalability Considerations

February 2022

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 11 August 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.	Network Resource Partition Scalability Requirements	3
3.	Network Resource Partition Scalability Considerations	5
3.1.	Control Plane Scalability	5
3.1.1.	Distributed Control Plane	5
3.1.2.	Centralized Control Plane	6
3.2.	Data Plane Scalability	6
3.3.	Gap Analysis of the Existing Mechanism	7
4.	Proposed Scalability Optimizations	8
4.1.	Control Plane Optimizations	8
4.1.1.	Distributed Control Plane Optimizations	8
4.1.2.	Centralized Control Plane Optimization	10
4.2.	Data Plane Optimizations	10
5.	Solution Evolution for Improved Scalability	12
6.	Security Considerations	13

7.	IANA Considerations	13
8.	Contributors	13
9.	Acknowledgments	13
10.	References	13
10.1.	Normative References	13

10.2.	Informative References	14
Authors'	Addresses	15

[1.](#) Introduction

The IETF Network Slice service aims to meet the connectivity demands of 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 Slice services, a concept called Network Resource Partition (NRP) is introduced, which refers to a set of network resources that are available in the underlay network to ensure the requested SLOs and SLEs of IETF Network Slices can be met.

[[I-D.ietf-teas-enhanced-vpn](#)] describes the layered framework and candidate technologies for delivering enhanced VPN (VPN+) services. VPN+ aims to meet the needs of customers or applications, including the applications that are associated with 5G, 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 built with a set of network resources. The VPN+ framework and technologies could be used for the realization of IETF Network Slice services. NRP could be used as the underlay network construct to support the VPN+ services.

As the demand for IETF Network Slice services increases, scalability would become an important factor for the large scale 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, there are concerns about the scalability of NRPs. This document describes the scalability considerations about NRPs in the network

control plane and data plane, and some optimization mechanisms are proposed.

2. Network Resource Partition Scalability Requirements

As described in [[I-D.ietf-teas-ietf-network-slices](#)], IETF Network Slices may be grouped together according to characteristics (including SLOs and SLEs). This grouping allows an operator to host a number of slices on a particular set of resources to reduce the amount of state information needed in the network. This can help to avoid the maintenance of per IETF Network Slice state in the underlay network.

With the development and evolution of 5G and other services, it is expected that an increasing number of IETF Network Slices will be deployed. The number of network slices required depends on how IETF Network Slices will be used, and the progress of network slicing for the vertical industrial services. The potential number of IETF Network Slice services and NRPs is analyzed by classifying the network slice deployment into three typical scenarios:

1. IETF Network Slices can be used by a network operator for different types of service. For example, in a converged multi-service network, different IETF Network Slices can be created to carry mobile transport services, fixed broadband services and enterprise services respectively: each type of service could be managed by a separate department or management team. Some service types, such as multicast services may also be deployed in a dedicated NRP. In this case, a separate NRP may need to be created for each service type. It is also possible that a network infrastructure operator provides IETF Network Slices to other network operators as a wholesale service, and a NRP may also be needed for each wholesale service customer. In this scenario, the number of NRPs in a network could be relatively small, such as in the order of 10 or so. This could be one of the typical cases in the beginning of IETF Network Slice deployment.
2. IETF Network Slices can be requested by customers of industrial verticals, where the assurance of SLOs and the fulfilment of SLEs are quite important. At the early stage of the vertical industrial services, a few top customers in some industries will

begin to use IETF Network Slices to provide performance assurance to their business, such as smart grid, manufacturing, public safety, on-line gaming, etc. The realization of such IETF Network Slices typically requires the provision of different NRPs for different industries, and some top customers can require dedicated NRPs for strict service performance guarantees. Considering the number of vertical industries, and the number of top customers in each industry, the number of NRPs needed may be in the order of 100.

3. With the evolution of 5G and cloud networks, IETF Network Slices could be widely used by various vertical industrial customers and enterprise customers who require guaranteed or predictable service performance. The total amount of IETF Network Slices may increase to thousands or more, although it is expected that the number of IETF Network Slices would still be less than the number of traditional VPN services in the network. 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 are 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, and the required NRPs may increase as well.

Thus, there may be large numbers of IETF Network Slices in some scenarios and the realization of IETF Network Slices needs to meet the scalability requirements. Mapping multiple IETF Network Slices to the same NRP presents a significant scaling benefit, but there can still be a requirement for a large number of NRPs which presents its own scalability challenges.

[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 hybrid of a centralized controller and a distributed control plane. The subsections that follow consider the scalability of these two approaches: the resultant scalability of the control plane will depend on how the hybrid is constructed.

[3.1.1.](#) Distributed Control Plane

It is necessary to create multiple NRPs for the delivery of IETF network slice services. Each NRP can be associated with a customized logical topology. The network resource attributes and the associated logical 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 in the following aspects:

- * The number of control protocol instances maintained on each node
- * The number of protocol sessions maintained on each link
- * The number of routes advertised by each node
- * The amount of attributes associated with each route

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

As the number of NRPs increases, it is expected that 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 separated control protocol instances (e.g., IGP instances) for different NRPs is considered higher than maintaining the information of separate 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 the information exchange of multiple NRPs. To meet the requirements of the increasing number of NRPs, it is suggested to choose a control plane mechanism which could improve the scalability while still provide the required

functionality.

[3.1.2.](#) Centralized Control Plane

By introducing a centralized network controller, the Software Defined Network (SDN) approach 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 centralized controller, thus the scalability of the controller also needs to be considered.

To provide global optimization for the Traffic Engineered (TE) paths in different NRPs, the controller needs to keep the topology and resource information of all the NRPs up-to-date. To achieve this, the controller may need to maintain a communication channel with each network node in the network. When there is significant change in the network, or multiple NRPs require global optimization concurrently, there may be a heavy processing burden at the controller, and a heavy load in the network surrounding the controller for the distribution of the updated network state and the TE paths.

[3.2.](#) Data Plane Scalability

To provide different IETF Network Slice services with the required SLOs and SLEs, it is necessary to allocate different subsets of network resources as different NRPs to avoid or reduce the unexpected interference from other services in the network. As the number of NRPs increases, it is required that the underlying network can provide fine-granular network resource partitioning, which means the amount of state about the partitioned network resources to be maintained on the network nodes will also increase.

In packet forwarding, IETF Network Slice service traffic needs to be processed according to the topology and 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 either directly or implicitly. Different approaches of encapsulating the NRP information in data packet can have different scalability implications.

One practical approach is to reuse some of the existing fields in the data packet to additionally identify the NRP the packet belongs to. For example, the destination IP addresses or the MPLS forwarding labels may be reused to further identify the NRP. This can avoid the cost of introducing new fields in the data packet, while since it introduces additional semantics to the existing fields, the processing of the existing fields in packet forwarding may need to be changed. Moreover, introducing resource semantics to existing identifiers in the packet (e.g., IP addresses, MPLS forwarding labels, etc.) may result in the increase of the amount of the existing IDs in proportion to the number of the NRPs, which may cause scalability problem in networks where a relatively large number of NRPs is needed.

An alternative approach is to introduce a new dedicated field in the data packet for identifying the NRP. This could avoid the impacts to the existing fields in the packet. And if this new field carries a globally-significant identifier of the NRP, it could be used together with the existing fields to determine the packet forwarding behavior. The potential issue with this approach is the difficulty in introducing a new field in some of the data plane technologies.

In addition, the introduction of NRP specific packet forwarding has an impact on the scalability of the forwarding entries on network nodes, as a network node may need to maintain separate forwarding entries for each NRP it participates in.

3.3. Gap Analysis of the Existing Mechanism

This section provides a gap analysis for an existing mechanism to perform NRP identification in the data plane and the related information distribution in the control plane.

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 distribute the resource attributes and the associated logical topology information in the control plane using mechanisms based on Multi-topology

[\[I-D.zhu-lsr-isis-sr-vtn-flexalgo\]](#). This mechanism is suitable for networks where a 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 needed 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. As 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, and the maximum number of Flexible Algorithms [\[I-D.ietf-lsr-flex-algo\]](#) is 128, which may not meet the required number of NRPs in some network scenarios.

[4. Proposed Scalability Optimizations](#)

[4.1. Control Plane Optimizations](#)

[4.1.1. Distributed Control Plane Optimizations](#)

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

The first optimization mechanism is to reduce the amount of control plane sessions used for the establishment and maintenance of the NRPs. For multiple NRPs which have the same connection relationship between two adjacent network nodes, it is proposed that one single control protocol session is used for each such group of NRPs. The information of different NRPs can be exchanged over the same session, with necessary identification information to distinguish the NRPs in the control message. This could reduce the overhead of maintaining a large number of separate control protocol sessions for each NRP, and could also reduce the amount of control plane messages flooded in the network.

The second optimization mechanism is to decouple the NRP information from the associated logical topology information in the control plane, so that the resource attributes and the topology attributes can be advertised and processed separately. In a network, it is possible that multiple NRPs associate with the same logical topology, and multiple NRPs may share the same set of network resources on a subset of network nodes and links. Then it is more efficient if only one copy of the topology information is advertised, and multiple NRPs sharing the same topology could refer to this topology information. More importantly, with this approach, the result of topology-based route computation could be shared by multiple NRPs, so that the overhead of per NRP route computation could be avoided. Similarly, information of a subset of network resources reserved on a particular network node or link could be advertised once and may be referred to by multiple NRPs which share the same set of resources.

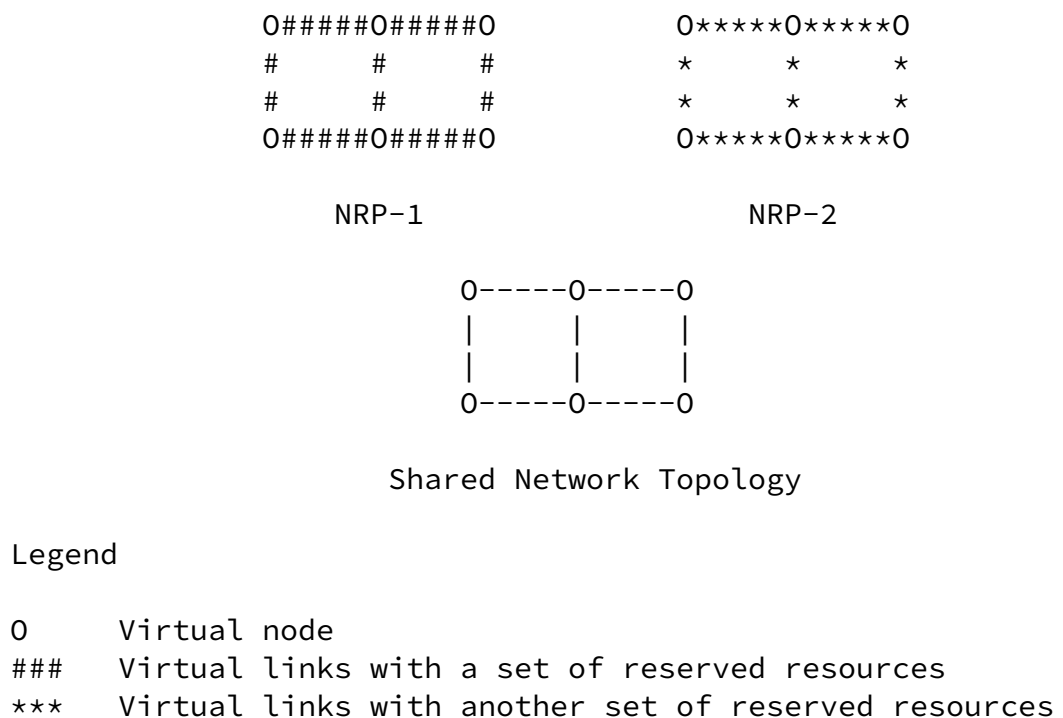


Figure 1. Topology Sharing between NRPs

Figure 1 gives an example of two NRPs which share the same logical topology. As shown in the figure, NRP-1 and NRP-2 are associated with the same topology, while the resource attributes of each NRP are different. In this case, only one copy of the network topology information needs to be advertised, and the topology-based route computation result can be shared by the two NRPs to generate the corresponding routing and forwarding tables.

Internet-Draft

NRP Scalability Considerations

February 2022

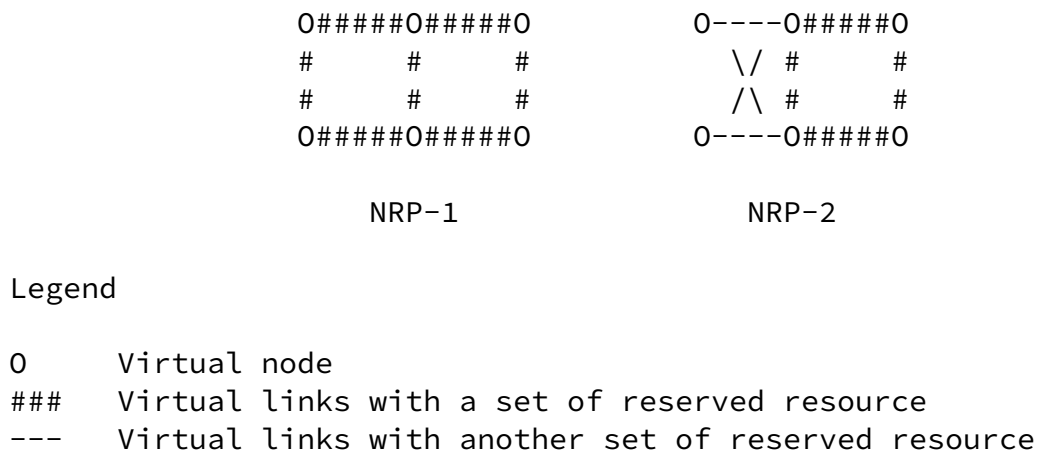


Figure 2. Resource Sharing between NRPs

Figure 2 gives another example of two NRPs which share the same set of network resources on some of the links. In this case, information about the resources allocated on each link only needs to be advertised once, then both NRP-1 and NRP-2 could refer to the reserved 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 mechanism to the distributed control plane rather than a replacement, so that the workload for NRP specific path computation in control plane could be shared by both the centralized controller and the network nodes, and the scalability of both systems could be improved. In addition, the centralized controller may be realized with multiple network entities, each is responsible for one subset or region of the network. This is the typical approach for scale out of the centralized control plane.

[4.2.](#) Data Plane Optimizations

To support more IETF Network Slice services while keeping the amount of data plane state at a reasonable scale, one typical approach is to classify a set of IETF Network Slice services which have similar

service characteristics and performance requirements into a group, and such groups of IETF Network Slice services are mapped to one NRP, which is allocated with an aggregated set of network resources and the union of the required logical topologies to meet the service requirement of the whole group of IETF Network Slice services. Different groups of IETF Network Slice services can be mapped to different NRPs with different set of network resources allocated from the underlay network. With appropriate grouping of IETF Network Slice services, a reasonable number of NRPs with network resources

reservation and aggregation could still meet the IETF Network Slice service requirements.

Another optimization in the data plane is to decouple the identifiers used for topology-based forwarding and the identifier used for the resource-specific processing introduced by NRP. One possible mechanism is to introduce a dedicated network-wide NRP Identifier (NRP-ID) in the packet header to uniquely identify the set of network resources allocated to a NRP on each involved network node and link for the processing and forwarding of the data packets mapped to the NRP. Then the existing identifiers in the packet header used for topology based forwarding (e.g., the 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 global NRP-ID field will be used together with other existing fields to determine the packet forwarding behavior, this may require network nodes to support a hierarchical forwarding table in data plane. Figure 3 shows the concept of using different data plane identifiers for topology-specific and resource-specific packet forwarding and processing respectively.

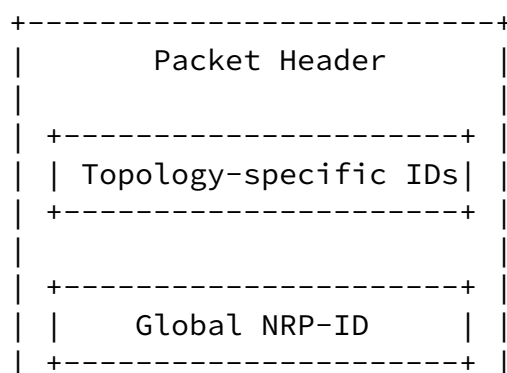


Figure 3. Decoupled Topology and Resource Identifiers in data packet

In an IPv6 [[RFC8200](#)] based network, this could be achieved by introducing a dedicated field in either the IPv6 fixed 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.dong-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](#)] based network, this may be achieved by introducing a dedicated NRP-ID either in the MPLS label stack or following the MPLS label stack. This way, the existing MPLS forwarding labels are used for topology-specific packet forwarding towards the destination node, 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 be 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 data plane are under discussion as part of the work in MPLS Open Design Team and is out of the scope of this document.

5. Solution Evolution for Improved Scalability

Based on the analysis in this document, the control plane 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.

At the first step, by introducing resource-awareness to segment routing SIDs [[I-D.ietf-spring-resource-aware-segments](#)], and using Multi-Topology or Flex-Algo as the control plane mechanism to define the logical topology, it could provide a solution for building a limited number of NRPs in the network, and can meet the requirements of 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 attribute from the resource attribute, so that multiple NRPs could share the same topology or resource attribute to reduce the control plane and data plane overhead. The data plane can still be based 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 [\[I-D.ietf-spring-sr-for-enhanced-vpn\]](#).

When 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, this could help to reduce the number of IP addresses or SR SIDs needed to support a large number of NRPs. This mechanism is called the NRP-ID based mechanism.

[6.](#) Security Considerations

This document describes the scalability considerations about the network control plane and data plane of NRPs in the realization of IETF Network Slice services, and proposes the 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.

[7.](#) IANA Considerations

This document makes no request of IANA.

[8.](#) Contributors

Zhibo Hu
Email: huzhibo@huawei.com

Hongjie Yang
Email: hongjie.yang@huawei.com

[9.](#) Acknowledgments

The authors would like to thank Adrian Farrel for the review and discussion of this document.

[10.](#) References

[10.1.](#) 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+) Services", Work in Progress, Internet-Draft, [draft-ietf-teas-enhanced-vpn-09](#), 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-09.txt>>.

Dong, et al.

Expires 11 August 2022

[Page 13]

Internet-Draft

NRP Scalability Considerations

February 2022

[I-D.ietf-teas-ietf-network-slices]
Farrel, A., Gray, E., 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-05](#), 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-05.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>>.

10.2. Informative References

- [I-D.dong-6man-enhanced-vpn-vtn-id]
Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network (VTN) Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, [draft-dong-6man-enhanced-vpn-vtn-id-06](#), 24 October 2021, <<https://www.ietf.org/archive/id/draft-dong-6man-enhanced-vpn-vtn-id-06.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-18](#), 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-18.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-02](#), 13 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-lsr-isis-sr-vtn-mt-02.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-03](https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-03), 12 July 2021,
<<https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-03.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-01](https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-01), 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-01.txt>>.

[I-D.zhu-lsr-isis-sr-vtn-flexalgo]

Zhu, Y., Dong, J., and Z. Hu, "Using Flex-Algo for Segment Routing based VTN", Work in Progress, Internet-Draft, [draft-zhu-lsr-isis-sr-vtn-flexalgo-03](https://www.ietf.org/archive/id/draft-zhu-lsr-isis-sr-vtn-flexalgo-03), 11 July 2021, <<https://www.ietf.org/archive/id/draft-zhu-lsr-isis-sr-vtn-flexalgo-03.txt>>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](https://www.rfc-editor.org/info/rfc4915), 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](https://www.rfc-editor.org/info/rfc5120), 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

Internet-Draft

NRP Scalability Considerations

February 2022

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

