TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: 28 April 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

25 October 2021

Scalability Considerations for Enhanced VPN (VPN+) draft-dong-teas-enhanced-vpn-vtn-scalability-04

Abstract

Enhanced VPN (VPN+) aims to meet the needs of some customers or applications, including the customers and applications that are associated with 5G, which requires connectivity services with advanced characteristics, such as the assurance of some Service Level Objectives (SLOs) and specific Service Level Expectations (SLEs). VPN+ could be used for network slice realization both in the context of 5G and in more generic scenarios, such as enterprise services which have requirement on the performance assurance. With the demand for VPN+ services increases, scalability would become an important factor for the large scale deployment of VPN+. This document describes the scalability considerations about the network control plane and data plane in enabling VPN+ services, some optimization mechanisms are also proposed.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Virtual Private Networks (VPNs) have served the industry well as a means of providing different customers with logically separated connectivity services over a common network infrastructure. The common or base network that is used to provide the VPNs is often referred to as the underlay, and the VPNs are often called the overlay. The underlay network is responsible for establishing the network connectivity and managing the network resources to meet specific service requirement. The overlay network is used to distribute the membership and reachability information of the customers, and provide logical separation in terms of service delivery between different customers in the shared network.

Enhanced VPN (VPN+) aims to meet the needs of some customers or applications, including the applications that are associated with 5G, which requires connectivity services with advanced characteristics, such as the assurance of Service Level Objectives (SLOs) and specific Service Level Expectations (SLEs).
[I-D.ietf-teas-ietf-network-slices] defines the terminologies and the general framework of IETF network slices. VPN+ could be used for IETF network slice realization both in the context of 5G and in more generic scenarios, such as enterprise services which have requirement on the performance assurance.

[I-D.ietf-teas-enhanced-vpn] describes the framework for delivering VPN+ services. To meet the requirement of some VPN+ services, a Virtual Transport Networks (VTNs) need to be created, which has a subset of network resources allocated from the physical network and is associated with a logical network topology to meet the requirements of one or a group of VPN+ services. VPN+ services can be delivered by mapping one or a group of overlay VPNs to the appropriate VTNs as the virtual underlay.

Section 6 of [I-D.ietf-teas-enhanced-vpn] provides some general analysis of the scalability of VPN+. This document gives further analysis of the scalability considerations when a large number of VPN+ services needs to be provided. Since the scalability of the overlay is usually not the major bottleneck, this document mainly focuses on the scalability of the VTNs in the underlay .

2. VPN+ Scalability Requirements

As described in [I-D.ietf-teas-enhanced-vpn], VPN+ services may require additional state to be introduced into the network to take advantage of the enhanced functionality. This may introduces some concerns about the network scalability. This section gives some analysis of the number of VPN+ services and the VTNs that might be needed in different network scenarios.

Since the typical use case of VPN+ is to deliver IETF network slice [I-D.ietf-teas-ietf-network-slices] for customers and services in 5G and other scenarios, the number of IETF network slices required could reflect the number of VPN+ needed in the 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 VPN+ services and VTNs 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 services. For example, in a converged multi-service network, different IETF network slices can be created to carry mobile transport service, fixed broadband service and enterprise services respectively, each type of service could be managed by a separate department or management team. Some service types, such as multicast service may also be deployed in a dedicated network slice. In this case, a separate VTN 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 VTN may also be needed for each wholesale service customer. In this scenario, the number of VTNs 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 in vertical industries, 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 to provide different VTNs for different industries, and some top customers can require dedicated VTNs for strict service performance guarantee.

Considering the number of vertical industries, and the number of top customers in each industry, the number of VTNs 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 VTNs needed may be in the order of 1000.

As defined by 3GPP [TS23501], a 5G network slice is identified using the Single Network Slice Selection Assistance Information (S-NSSAI), which is a 32-bit identifier comprised of 8-bit Slice/Service Type (SST) and 24-bit Slice Differentiator (SD). This allows the mobile networks (the RAN and mobile core networks) to support a large number of 5G network slices. Although it is likely that multiple 5G network slices are mapped to the same IETF network slice, in some cases the number of IETF network slices may still be comparable to the number of 5G network slices.

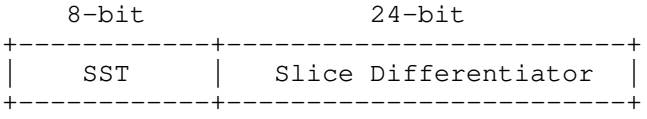

Figure 1. Format of S-NSSAI in 3GPP

Thus solution of VPN+ and VTN needs to meet the scalability requirement of IETF network slices in different scenarios. The increased number of VPN+ services will introduce additional complexity and overhead both to the control plane and the data plane, especially in the aspects related to the underlay VTNs. Although in many cases multiple VPN+ services can be mapped to the same VTN as the underlay, there still can be scalability challenges with the increased number of VTNs.

3. VTN Scalability Considerations

In this section, the scalability of VTN in the control plane and data plane is analyzed to understand the possible gaps in meeting the scalability requirement of VPN+ and VTN.

3.1. Control Plane Scalability

As described in [I-D.ietf-teas-enhanced-vpn], the control plane of VPN+ could be based on the hybrid of a centralized controller and the distributed control plane.

3.1.1. Distributed Control Plane

At part of the delivery of VPN+ services, it is necessary to create multiple VTNs, each of which is allocated with a set of dedicated or shared network resources, and is associated with a customized logical topology. The topological and resource attributes and the state information of each VTN may need to be exchanged among the network nodes. The scalability of the distributed control plane used for the distribution of VTN 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 computation (i.e. SPF computation) executed by each node

As the number of VTNs increases, it is expected that in some of the above aspects, the overhead in the control plane may increase dramatically. For example, the overhead of maintaining separated control protocol instances (e.g. IGP instances) for different VTNs is considered higher than maintaining the information of separated VTNs in the same control protocol instance with appropriate separation, and the overhead of maintaining separate protocol sessions for different VTNs is considered higher than using a shared protocol session for the information exchange of multiple VTNs. To meet the requirement of the increasing number of VTNs, It is suggested to choose the control plane mechanisms which could improve the scalability while still provide the required functionality.

3.1.2. Centralized Control Plane

By introducing the centralized network controller, the SDN approach can reduce the amount of control plane 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 VTNs, the controller needs to keep the topology and resource information of all the VTNs 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 VTNs requires 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 VPN+ services with the required SLOs and SLEs, it is necessary to allocate different subsets of network resources to different VTNs to avoid or reduce unexpected interruption. As the number of VTNs 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, VPN+ service traffic needs to be processed separately according to the topology and resource attributes of the VTN it mapped to, this means that some fields in the data packet needs to be used to identify the VTN topology and resources either directly or implicitly. Different approaches of encapsulating the VTN 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 VTN the packet belongs to. For example, the destination IP addresses or the MPLS forwarding labels may be reused to further identify a VTN. 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 VTN 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 VTNs, which may cause scalability problem in networks where a relatively large number of VTNs is needed.

An alternative approach is to introduce a new dedicated field in the data packet for VTN identification. This could avoid the impacts to the existing fields in the packet. And if this new field carries a global-significant VTN identifier, it could be used together with the existing fields to determine the VTN-specific packet forwarding. 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 per VTN packet forwarding has impact on the scalability of the forwarding entries on network nodes, as a network node may need to maintain separate forwarding entries for each VTN it participates in.

3.3. Gap Analysis of Existing Mechanisms

One candidate mechanism to build VTN is to use VTN-specific Segment Routing (either SR-MPLS or SRv6) Identifiers in the data plane as described in [I-D.ietf-spring-sr-for-enhanced-vpn], and define and distribute the associated topology and resource attribute of each VTN based on either Multi-topology [I-D.ietf-lsr-isis-sr-vtn-mt], Flex-Algo [I-D.zhu-lsr-isis-sr-vtn-flexalgo] or the combination of these mechanisms in the control plane. This mechanism is suitable for networks where a small number of VTNs is needed. As the number of VTNs 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 VTNs in the network, which will bring challenges both to the distribution of SIDs and the related information in the control plane, and to the installation of forwarding entries for VTN-specific SIDs in the data plane.

2. The number of route computation (e.g. SPF computation) will increase in proportion to the number of VTNs 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 is 128, and the maximum number of Flex-Algo is 128, which may not meet the required number of VTNs in some network scenarios.

4. Proposed Scalability Optimizations

4.1. Control Plane Optimizations

For the distributed control plane, several optimizations can be considered to reduce the control plane overhead and improve the control plane scalability.

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

The second optimization mechanism is to decompose the attributes of a VTN into different groups, so that different types of VTN attribute can be advertised and processed separately in control plane. There are two basic types of attributes associated with a VTN: the topology attribute and the network resource attribute. In a network, it is possible that multiple VTNs share the same topology, and multiple VTNs may share the same set of network resources on particular network nodes and links. Then it is more efficient if only one copy of the topology information is advertised, and multiple VTNs 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 VTNs, so that the overhead of per-VTN route computation could also be reduced . Similarly, information of a subset of network resources reserved on a particular network node or link could be advertised once and be referred to by multiple VTNs which share the same set of resources. This methodology could also apply to other attributes of VTN which may be introduced later and can be processed independently.

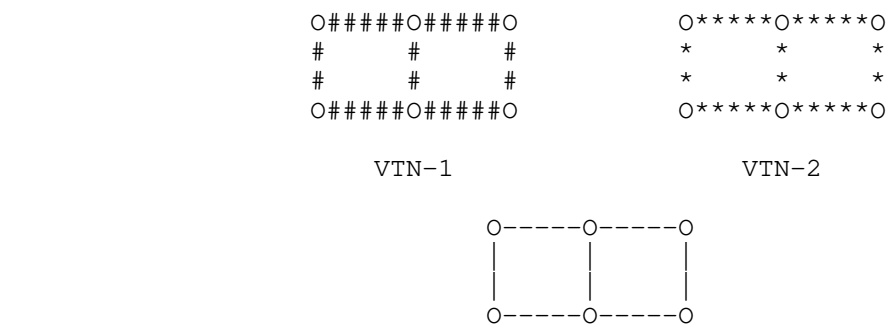

Figure 2. Topology Sharing between VTNs

Figure 1: FIG-2

Figure 2 gives an example of two VTNs which share the same logical topology. As shown in the figure, VTN-1 and VTN-2 are associated with the same topology, while the resource attributes of each VTN 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 VTNs to generate the corresponding routing and forwarding tables.

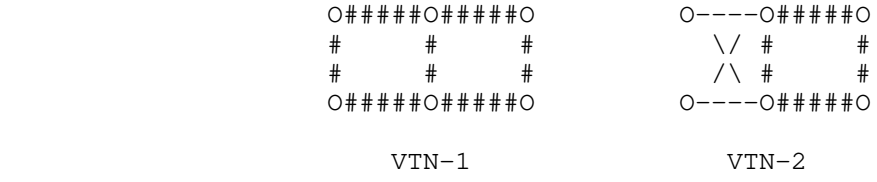

Figure 3. Resource Sharing between VTNs

Figure 3 gives another example of two VTNs 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 VTN-1 and VTN-2 could refer to the reserved link resource for constraint based path computation.

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 VTN 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.

4.2. Data Plane Optimizations

To support more VPN+ services while keeping the amount of data plane state at a reasonable scale, one typical approach is to classify a set of VPN+ services which have similar service characteristics and performance requirements into a group, and such group of VPN+ services are mapped to one VTN, 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 VPN+ services. Different groups of VPN+ services can be mapped to different VTNs with different set of network resources allocated. With appropriate grouping of VPN+ services, a reasonable number of VTNs with network resources reservation and aggregation could still meet the 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 VTN. One possible mechanism is to introduce a dedicated VTN Resource identifier in the packet header to uniquely identify the set of local network resources allocated to a VTN on each network node for the processing and forwarding of the received packets. 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 VTNs. Since this new VTN Resource ID field will be used together with other existing fields to determine the VTN-specific packet forwarding, this may require network nodes to support a hierarchical forwarding table in data plane. Figure 4 shows the concept of using different data plane identifiers for topology-specific and resource-specific packet forwarding and processing in a VTN respectively.

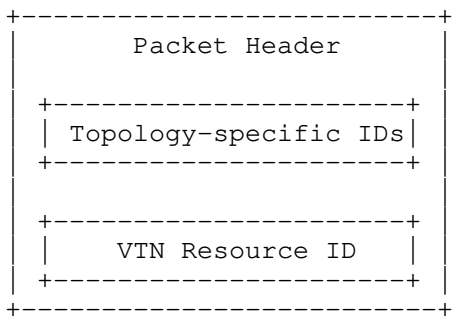

Figure 4. Decoupled Data Plane Topology and Resource Identifiers

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 VTN resource identifier 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 VTN resource ID needs to be parsed by every node along the path which is capable of VTN-specific forwarding. [I-D.dong-6man-enhanced-vpn-vtn-id] introduces the mechanism of carrying the VTN resource ID in IPv6 Hop-by-Hop extension header.

In an MPLS [RFC3032] based network, this may be achieved by introducing a dedicated VTN resource ID either in the MPLS label stack or following the MPLS label stack. This way, the existing MPLS forwarding labels could be used for topology-specific packet forwarding towards the destination node, and the VTN resource ID is used to determine the set of network resources for packet processing. This requires that both the forwarding label and the VTN Resource ID be parsed by nodes along the forwarding path of the packet, and the forwarding behavior may depend on the position of the VTN resource ID in the packet. The detailed extensions in MPLS data plane are 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 VPN+ and VTN needs to evolve to support the increasing number of VPN+ services and the increasing number of VTNs 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, it could provide a solution for building a limited number of VTNs in the network to meet the requirement of a relatively small number of VPN+ services in the network. This mechanism is considered as the basic SR VTN.

As the required number of VPN+ services increases, more VTNs may be needed, then the control plane scalability could be improved by decoupling the topology attribute from the resource attribute and other attributes of VTN, so that multiple VTNs could share the same topology or resource attribute to reduce the control plane and data plane overhead. This mechanism is considered as the scalable SR VTN. Both the basic and the scalable SR VTN mechanisms are described in [I-D.ietf-spring-sr-for-enhanced-vpn].

If the data plane scalability becomes a concern, a dedicated VTN resource ID can be introduced in the data packet to decouple the topology-specific identifiers from the VTN resource identifiers in the data plane, this could help to reduce the number of SR SIDs needed to support a large number of VTNs. This mechanism is considered as the Resource-Independent (RI) VTN.

6. Security Considerations

This document describes the scalability considerations about the network control plane and data plane in enabling VPN+ services and the VTNs, and proposes several scalability optimization mechanisms. The security considerations in [I-D.ietf-teas-enhanced-vpn] applies 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-08, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-08.txt>.

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 Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-dong-6man-enhanced-vpn-vtn-id-05, 8 September 2021, <https://www.ietf.org/archive/id/draft-dong-6man-enhanced-vpn-vtn-id-05.txt>.

[I-D.dong-lsr-sr-enhanced-vpn]
Dong, J., Hu, Z., Li, Z., Tang, X., Pang, R., JooHeon, L., and S. Bryant, "IGP Extensions for Scalable Segment Routing based Enhanced VPN", Work in Progress, Internet-Draft, draft-dong-lsr-sr-enhanced-vpn-06, 11 July 2021, <https://www.ietf.org/archive/id/draft-dong-lsr-sr-enhanced-vpn-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-17, 6 July 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-17.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-01, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-isis-sr-vtn-mt-01.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, 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, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-01.txt>.

[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-04, 23 August 2021, <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-04.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, 11 July 2021, <https://www.ietf.org/archive/id/draft-zhu-lsr-isis-sr-vtn-flexalgo-03.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>.

[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>.

[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>.

[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