TEAS working group
Internet-Draft
Intended status: Standard Track
Expires: August 2020

J. Dong
Z. Li
Huawei
F. Qin
China Mobile

February 10, 2020

Virtual Transport Network (VTN) Scalability Considerations for Enhanced VPN

draft-dong-teas-enhanced-vpn-vtn-scalability-00

Abstract

Enhanced VPN (VPN+) is an enhancement to VPN services to support the needs of new applications, particularly including the applications that are associated with 5G services. An enhanced VPN could be used for transport network slicing in 5G, and will also be of use in more generic scenarios. I-D.ietf-teas-enhanced-vpn describes the framework and candidate component technologies for providing enhanced VPN services. This document describes the scalability considerations in the control plane and data plane to enable VPN+ services, some optimization mechanisms are also described.

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 August 10, 2020.

Copyright Notice

Table of Contents

1. Introduction

Virtual private networks (VPNs) have served the industry well as a means of providing different groups of users with logically isolated connectivity over a common network. The VPN service is provided with two network layers: the overlay and the underlay. The underlay is responsible for establishing the network connectivity and managing network resources to meet the service requirement. The overlay is used to distribute the membership and reachability information of the tenants, and provide logical separation of services between different tenants.

Enhanced VPN service (VPN+) [I-D.ietf-teas-enhanced-vpn] is targeted at new applications which require better isolation and have more stringent performance requirements than can be provided with existing overlay VPNs. To meet the requirement of enhanced VPN services, a number of virtual transport networks (VTN) need to be created, each with a subset of the underlay network topology and a set of network resources allocated to meet the requirement of a specific VPN+ service or a group of VPN+ services. The overlay VPN together with the corresponding VTN in the underlay provide the enhanced VPN service.

[I-D.ietf-teas-enhanced-vpn] provides some general analysis to the scalability of VPN+. This document gives detailed analysis to the scalability considerations to enable enhanced VPN service. The focus of this document is on the underlay of the enhanced VPN, i.e. the virtual transport network.

In the context of 5G, enhanced VPN can be used to provide network slicing in transport network.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Scalability Requirement

As mentioned in [I-D.ietf-teas-enhanced-vpn], VPN+ services may require to install some additional state within the network to achieve the additional features. This introduces some scalability concerns to the network. This section gives some analysis about the number of VPN+ services needed in a network.

The number of enhanced VPNs required in a network is determined by the use cases. One typical use case of enhanced VPN is to provide transport network slicing for applications or services in 5G. With the development and evolution of 5G, it is expected that more network slices will be needed. The number of network slices required in a network is relevant to how network slicing is used in the network and the evolution of 5G for vertical industrial services. The potential number of network slices is analyzed by classifying the network slicing deployment into three typical types of scenarios:

1. Network slicing can be used to isolate different types of business of the network operator. For example, in a converged multi-service network, different network slices can be created to carry mobile service, fixed broadband service and enterprise service respectively, each type of service could be managed by a separate department or management team. Some particular service types, such as multicast service may also be deployed in a dedicated network slice. It is also possible that a infrastructure network operator can provide network slices to other network operators as wholesale service. In this scenario, the number of network slices in a network would be relatively small, such as in the order of 10 or so. This could be the typical case in the beginning of network slicing deployment.

2. Network slicing can be used to provide isolated and customized virtual networks for tenants of different vertical industries. At the early stage of the vertical industrial service deployment, a few top tenants in some typical industries will begin to use network slicing to support their business, such as smart grid, manufacture, public safety, on-line games etc. Considering the number of the vertical industries, and the number of top tenants in each industry, the number of network slices may increase to around 100.

3. With the evolution of 5G, network slicing could be widely used by both vertical industrial tenants and premium business tenants. The total amount of network slices may increase to the order of 1000 or more. While it is expected that the number of network slices would be still less than the number of traditional VPN services in the network.

In 3GPP [TS23501], a network slice is identified using 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 network (RAN and CN) to provide a large number of network slices. Although it is possible that several network slices in RAN and CN can be mapped to the same transport network slice, the scalability of transport network slices needs to be taken into consideration from the beginning.

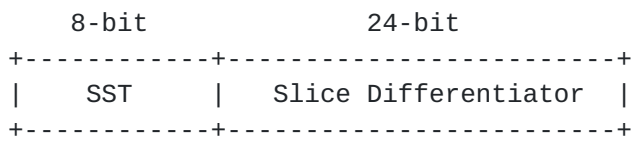

Figure 1 Format of Network Slice Identifier in 3GPP

Enhanced VPN needs to meet the scalability requirement of network slicing in different scenarios. The increased number of enhanced VPNs will introduce additional complexity and overhead to both the control plane and data plane, especially for the underlying virtual transport network.

4. Scalability Considerations

In this section, the scalability of control plane and data plane is analyzed to understand whether the existing mechanisms could meet the scalability requirement of enhanced VPNs, and to identify possible optimizations.

4.1. Control Plane Scalability Considerations

As described in section 3.1 of [I-D.ietf-teas-enhanced-vpn], the control plane of enhanced VPN could be based on a hybrid of centralized controller and distributed control plane.

4.1.1. Distributed Control Plane

As the underlay of VPN+ service, it is required that the different VTNs need to be created to provide customized topology and resource attributes for different applications or tenants, and the state of each VTN needs to be exchanged in control plane. The scalability of the distributed control plane for the establishment and maintenance of VTNs needs to be considered in the following aspects:

- The number of control protocol instances maintained on each node
- The number of the protocol sessions maintained on each node
- The number of routes advertised by each node
- The amount of attributes associated with each route
- The number of route computation (i.e. SPF) executed on each node

As the number of VTNs increases, it is expected that for some of the above aspects, the overhead in control plane may become unaffordable. For example, the overhead of maintaining separated routing instances for different VTNs is considered higher than maintaining separated virtual network topologies for different VTNs in the same routing instance, and the overhead of maintaining separate protocol sessions for each VTN is higher than using a shared protocol session for the information exchange of multiple VTNs. In order 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.

4.1.2. Centralized Control Plane

Although the SDN approach can reduce the amount of control plane overhead in the distributed control plane, it may transfer some of the scalability concerns from the network to the centralized controller, thus the scalability of the controller also needs to be considered.

In order to provide global optimization for 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 and multiple VTNs requires global optimization concurrently, there may be a heavy processing burden at the controller, and also a heavy load in the network surrounding the controller for the distribution of the updated network state.

4.2. Data Plane Scalability Considerations

To provide different enhanced VPNs with the required isolation and performance, it is necessary to allocate different set of network resources to different VTNs to provide the underlay for different enhanced VPNs. As the number of enhanced VPNs increases, the number of VTNs would increase accordingly. This requires the underlying network to provide finer-granular network resource partitioning, which means the amount of states about the reserved network resources to be maintained on network nodes will also increase.

In a network, traffic of different VPN+ services need to be processed separately according to the topology and resource constraints of the corresponding VTN, thus the identifier of the corresponding VTN needs to be carried either directly or implicitly in the data packet. Different representations of the VTN ID in data

packet has different scalability characteristics. It is possible to reuse some existing fields in packet header to additionally identify the VTN the packet belongs to, while this may result in more of the existing identifiers being allocated than expected in the original design. An alternative is to introduce a new identifier in the packet for VTN identification.

In addition, the introduction of per VTN forwarding has impact on the scalability of the forwarding entries on network nodes, as a network node needs to maintain separate forwarding entries for each VTN it participates.

4.3. Gap Analysis of Existing Mechanism

One candidate approach to build VTN is using Segment Routing (either SR-MPLS or SRv6) as data plane and distributing the customized topology and resource attribute based on Multi-topology [RFC4915] [RFC5120] and/or Flex-Algo [I-D.ietf-lsr-flex-algo] mechanism in control plane. While if the number of VTNs increases to more than 100, such approach may have several scalability issues:

1. The number of SR SIDs needed would increase proportional to the number of VTNs in the network, which would bring challenge both to the control plane distribution of the SIDs and to the installation of data plane forwarding entries for the SIDs.

2. The number of SPF computation would also increase proportional to the number of VTNs in the network, which can introduce significant overhead to the control plane of network nodes.

3. The maximum number of virtual network instances supported by OSPF Multi-topology and Flex-Algo is 128, which may not meet the required number of VTNs in a network.

5. Possible Optimization

5.1. Control Plane Optimization

For the distributed control plane, several optimizations are proposed to reduce the overhead and improve the control plane scalability.

The first proposed 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 session is used for the establishment of multiple VTNs. Information of different VTNs can be exchanged over the same control session, with necessary identification information to distinguish them in the control messages. This could reduce the overhead of maintaining large amount of control sessions, and could also reduce the amount of control message flooding in the network.

The second proposed mechanism is to decompose the attributes of a VTN into different groups, so that different types of attribute can be advertised and processed separately in the control plane. For a VTN, there are two basic types of attributes, the topology attribute and the associated network resource attribute. In a network, multiple VTNs could share the same topology, and multiple VTNs may share the same set of network resource on particular network segments. It would be more efficient if only one copy of the topology attribute is advertised, then multiple VTNs referring to the same topology could share the topology information and the result of topology based route computation. Similarly, information of a subset of reserved network resource could be advertised once and then be used by multiple VTNs. This methodology also applies to other attributes of VTN which may be introduced later and can be processed independently.

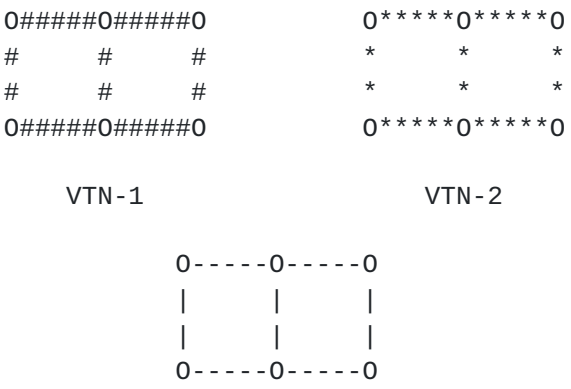

Shared Network Topology

Legend

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

Figure 2 Topology Sharing between VTNs

Figure 1 gives an example of multiple VTNs which shares the same topology attribute. As shown in the figure, VTN-1 and VTN-2 have the same topology, while the resource attributes on links 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 used by both VTNs to generate their routing tables.

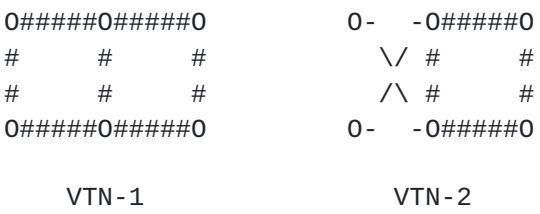

Legend

O Virtual node
Virtual links with a set of reserved resource
--- Virtual links with another set of reserved resource

Figure 3 Resource Sharing between VTNs

Figure 2 gives another example of multiple VTNs which shares the same set of network resources on some links. Similarly, information about the reserved resource on each link only needs to be advertised once, then both VTN-1 and VPN-2 could refer to the link resource for constraint based computation.

For the centralized control plane, it is suggested that the centralized controller is deployed as a complementary mechanism to the distributed control plane rather than a total replacement, so that the computation burden in control plane could be shared by both the centralized controller and the network nodes, thus the scalability of both system could be improved.

5.2. Data Plane Optimization

In order to support more enhanced VPNs services while keeping the amount of data plane state in a reasonable scale, one possible approach is to classify a set of enhanced VPN services which has similar service characteristics and performance requirements into a group, and such group of enhanced VPNs is mapped to one VTN which is allocated with an aggregated set of network resources to meet the service requirement of the whole group of enhanced VPNs. Different

groups of enhanced VPNs need to be mapped to different VTNs with different set of network resources allocated. With appropriate grouping of enhanced VPN services, a reasonable number of VTNs with network resources aggregation could still meet the service requirements.

Another optimization in data plane is to decouple the identifier 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 field in packet header to uniquely identify the set of local network resources allocated to the VTN on each network node for the processing of the received packet. Then the existing identifier in packet header used for topology based forwarding is kept unchanged. The benefit is the number of existing topology-specific identifiers will only increase as the number of the virtual network topologies increases, so that the scalability of the existing identifier will not be impacted by the increase of VTN. Note this probably requires network nodes to support a hierarchical forwarding table in the data plane. Figure 3 shows

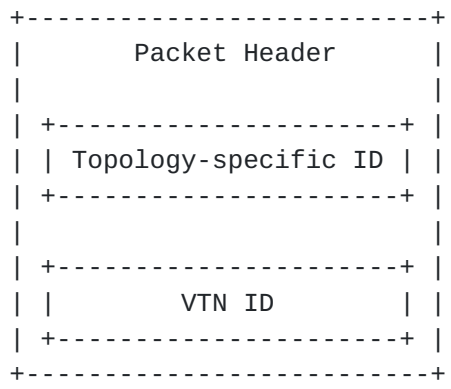

Figure 4 Decoupled Data Plane Identifiers

In an IPv6 [RFC8200] based network, this could be achieved by introducing a dedicated field in the IPv6 fixed header or one of the extension headers to carry the VTN 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 ID needs to be parsed by every node along the path which is capable of VTN specific forwarding. In an MPLS [RFC3032] based network, this may be achieved by introducing a new dedicated MPLS label to identify the VTN instance, while the existing MPLS labels could be used for topology based packet

forwarding towards the associated destination prefix. This requires that both labels be parsed by each node along the forwarding path of the packet. The detailed extensions in IPv6 and MPLS encapsulation are out of the scope of this document.

6. Solution Evolution for Improved Scalability

Based on the analysis in this document, the solution for enhanced VPN needs to evolve to support the increasing number of enhanced VPNs in the network.

For example, by introducing resource awareness to segment routing SIDs [I-D.dong-spring-sr-for-enhanced-vpn], and using Multi-Topology or Flex-Algo as control plane could provide a solution for building a limited set of VTNs in the network to meet the requirement of a small number of enhanced VPNs in the network. Such mechanism can be called SR-VTN.

As the number of required enhanced VPNs increases, more VTNs needs to be created, then the control plane scalability could be improved by introducing topology sharing between multiple VTNs. Such mechanism can be called Topology Independent (TR) SR-VTN.

In order to further improve the data plane scalability, dedicated data plane identifiers of VTN can be introduced to decouple the topology based forwarding and the resource based processing in the data plane. Such mechanism can be called Resource Independent (RI) SR-VTN.

7. Security Considerations

TBD

IANA Considerations

This document makes no request of IANA.

Acknowledgments

The authors would like to thank XXX for the review and valuable comments.

References

Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Informative References

[I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Service", draft-ietf-teas-enhanced-vpn-04 (work in progress), January 2020.

[I-D.dong-spring-sr-for-enhanced-vpn] Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Segment Routing for Enhanced VPN Service", draft-dong-spring-sr-for-enhanced-vpn-06 (work in progress), December 2019.

[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>.

[I-D.ietf-lsr-flex-algo] Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", draft-ietf-lsr-flex-algo-05 (work in progress), November 2019.

[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>.

[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>.

[TS23501] "3GPP TS23.501", 2019, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

Authors' Addresses

Jie Dong
Huawei

Email: jie.dong@huawei.com

Zhenbin Li
Huawei

Email: lizhenbin@huawei.com

Fengwei Qin
China Mobile

Email: qinfengwei@chinamobile.com