

TEAS Working Group  
Internet-Draft  
Intended status: Informational  
Expires: May 6, 2021

J. Dong  
Z. Li  
Huawei Technologies  
F. Qin  
China Mobile  
G. Yang  
China Telecom  
November 2, 2020

**Scalability Considerations for Enhanced VPN (VPN+)**  
**draft-dong-teas-enhanced-vpn-vtn-scalability-01**

**Abstract**

Enhanced VPN (VPN+) aims to provide enhancements to existing VPN services to support the needs of new applications, particularly including the applications that are associated with 5G services. VPN+ could be used to provide network slicing in 5G, and may also be of use in more generic scenarios, such as enterprise services which have demanding requirement. With the requirement for VPN+ services increase, scalability would become an important factor for deployment of VPN+. 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 May 6, 2021.

**Copyright Notice**

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	VPN+ Scalability Requirements . . . . .	<a href="#">3</a>
<a href="#">3.</a>	VPN+ Scalability Considerations . . . . .	<a href="#">4</a>
<a href="#">3.1.</a>	Control Plane Scalability . . . . .	<a href="#">5</a>
<a href="#">3.1.1.</a>	Distributed Control Plane . . . . .	<a href="#">5</a>
<a href="#">3.1.2.</a>	Centralized Control Plane . . . . .	<a href="#">5</a>
<a href="#">3.2.</a>	Data Plane Scalability . . . . .	<a href="#">6</a>
<a href="#">3.3.</a>	Gap Analysis of Existing Mechanisms . . . . .	<a href="#">6</a>
<a href="#">4.</a>	Possible Optimizations . . . . .	<a href="#">7</a>
<a href="#">4.1.</a>	Control Plane Optimization . . . . .	<a href="#">7</a>
<a href="#">4.2.</a>	Data Plane Optimization . . . . .	<a href="#">9</a>
<a href="#">5.</a>	Solution Evolution for Improved Scalability . . . . .	<a href="#">10</a>
<a href="#">6.</a>	Security Considerations . . . . .	<a href="#">11</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">11</a>
<a href="#">8.</a>	Contributors . . . . .	<a href="#">11</a>
<a href="#">9.</a>	Acknowledgments . . . . .	<a href="#">11</a>
<a href="#">10.</a>	References . . . . .	<a href="#">11</a>
<a href="#">10.1.</a>	Normative References . . . . .	<a href="#">11</a>
<a href="#">10.2.</a>	Informative References . . . . .	<a href="#">11</a>
	Authors' Addresses . . . . .	<a href="#">12</a>

## [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 infrastructure. The VPN service is provided with two network layers: the overlay and the underlay. The underlay is responsible for establishing 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 service delivery between different tenants.

Enhanced VPN service (VPN+) [[I-D.ietf-teas-enhanced-vpn](#)] is targeted at new applications which require better isolation between tenants and/or services, and have more stringent performance requirements



than can be provided with existing VPNs. To meet the requirement of VPN+ services, 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 requirements of one or a group of VPN+ services. The VPN in overlay together with the corresponding VTN in the underlay provide the VPN+ service.

[I-D.ietf-teas-enhanced-vpn] provides some general analysis of the scalability of VPN+. This document gives detailed analysis of the scalability considerations when enabling VPN+ services. The focus of this document is mainly on the scalability of the underlay of VPN+, i.e. the VTN.

## 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 introduces some scalability considerations to the network. This section gives some analysis of the number of VPN+ services that might be needed in a network.

There are several use cases where VPN+ may be necessary, and these determine how many will be required in a network. One typical use case of VPN+ is to provide network slicing for applications or services in 5G, thus the number of network slices needed could reflect the number of VPN+ services. In the future, with the development and evolution of 5G, it is expected that more and more network slices will be deployed. The number of network slices required is relevant to how network slicing will be used, and the progress of 5G for vertical industrial services. The potential number of network slices is analyzed by classifying the network slicing deployment into three typical scenarios:

1. Network slicing can be used by a network operator internally to isolate different types of services. For example, in a converged multi-service network, different 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. It is also possible that an infrastructure network operator provides network slices to other network operators as a wholesale service. In this scenario, the number of network slices in a network would be relatively small, such as on the order of 10 or so. This could be the typical case in the beginning of the network slicing deployment.



2. Network slicing can be used to provide isolated and customized virtual networks for tenants in 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, manufacturing, public safety, on-line gaming, 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 the order of 100.
3. With the evolution of 5G, network slicing could be widely used by both vertical industrial tenants and enterprise tenants which require guaranteed or predictable service performance. 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 still be less than the number of traditional VPN services in the network.

In 3GPP [[TS23501](#)], a 5G 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 multiple network slices in RAN and CN can be mapped to the same transport network slice, the amount of transport slice still needs to be comparable with the number of 5G network slices. Thus the scalability of transport network slices needs to be taken into consideration from the beginning.

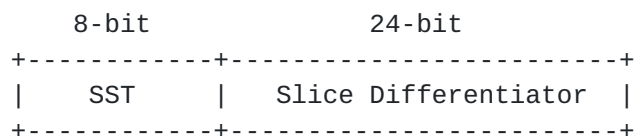


Figure 1. Format of S-NSSAI in 3GPP

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

### 3. VPN+ Scalability Considerations

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



### **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 centralized controller and distributed control plane.

#### **3.1.1. Distributed Control Plane**

At part of the construction of VPN+ services, it is necessary to create different VTNs that provide customized topology and resource attributes. The attributes and state information of each VTN needs to be exchanged in the control plane. The scalability of the distributed control plane for the establishment and maintenance of VTNs needs to be considered in the following aspects:

- o The number of control protocol instances maintained on each node
- o The number of protocol sessions maintained on each link
- o The number of routes advertised by each node
- o The amount of attributes associated with each route
- o 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 the control plane may increase dramatically. For example, the overhead of maintaining separated control protocol instances for each VTN 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 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**

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 network nodes to the centralized controller, thus the scalability of the controller also needs to be considered.

To provide global optimization for 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 and 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.

### **3.2. Data Plane Scalability**

To provide different VPN+ services with the required isolation and performance characteristics, it is necessary to allocate different sets of network resources to different VTNs. As the number of VPN+ increases, the number of VTNs will increase accordingly. This requires the underlying network to provide finer-granular network resource partitioning, which means the amount of state about the reserved network resources to be maintained on network nodes will also increase accordingly.

In data plane, traffic of different VPN+ services need to be processed separately according to the topology and resource constraints of the associated VTN, thus the identifier of VTN needs to be carried either directly or implicitly in the data packet. Different representations of the VTN identifier in data packet have different scalability implication. One approach is to reuse some existing fields in packet headers to additionally identify the VTN the packet belongs to. As this introduces additional semantics to an existing identifier, it may increase the amount of the identifiers to be allocated and managed, which may not be expected in its original design and could cause scalability problem. An alternative is to introduce a dedicated identifier in the packet for VTN identification.

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 needs to maintain separate forwarding entries for a target node in each VTN it participates.

### **3.3. Gap Analysis of Existing Mechanisms**

One candidate approach to build VTN is using Segment Routing (either SR-MPLS or SRv6) as the data plane, and distributing the customized topology and resource attribute based on Multi-topology [[RFC4915](#)] [[RFC5120](#)], Flex-Algo [[I-D.ietf-lsr-flex-algo](#)] or the combination of these mechanisms in the control plane. If the number of VTNs increases to a certain extent, such approach may have several scalability issues:



1. The number of SR SIDs needed will increase dependent upon the number of VTNs in the network, which will bring challenges both to the SID information distribution in control plane and to the installation of forwarding entries for the SIDs in data plane.
2. The number of SPF computation will also increase in proportion to the number of VTNs in the network, which can introduce significant overhead of the computing resources on network nodes.
3. The maximum number of network topology supported by OSPF Multi-topology is 128, the maximum number of Flex-Algo is 128, which may not meet the required number of VTNs in some networks.

#### **4. Possible Optimizations**

##### **4.1. Control Plane Optimization**

For the distributed control plane, several optimizations can be considered to reduce the 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 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 a large number of control protocol sessions, and could also reduce the amount of control plane message flooding in the network.

The second optimization 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 control plane. For a VTN, there are two basic types of attributes: the topology attribute and the associated 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 resource on particular network segments. It is more efficient if only one copy of the topology attribute is advertised, then multiple VTNs sharing the same topology could refer to the topology information, and share the result of topology-based route computation. Similarly, information of a subset of network resource reserved on network segments could be advertised once and then be used by multiple VTNs. This methodology could also apply to other attributes of VTN which may be introduced later and can be processed independently.



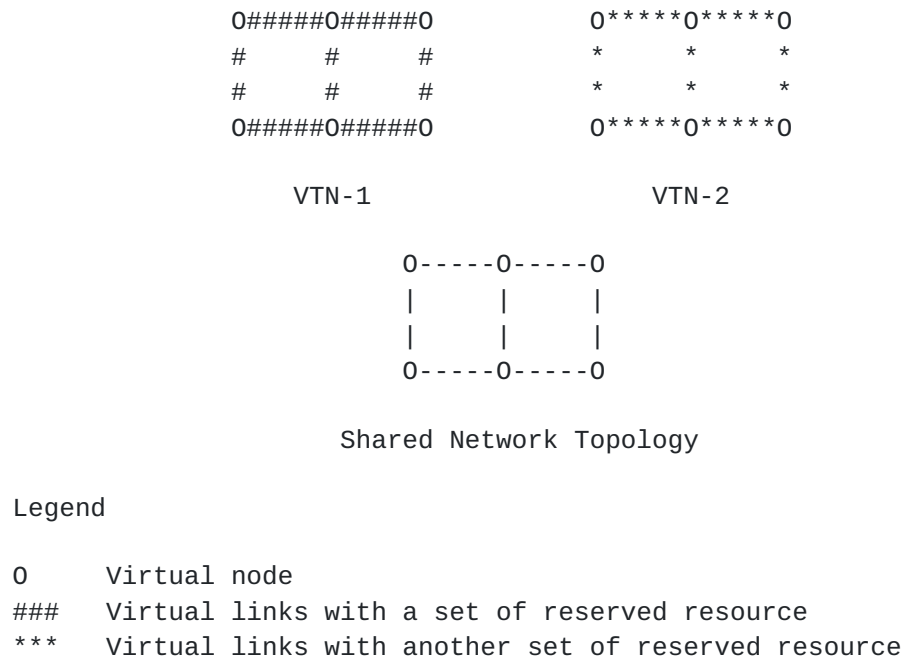


Figure 2. Topology Sharing between VTNs

FIG-2

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

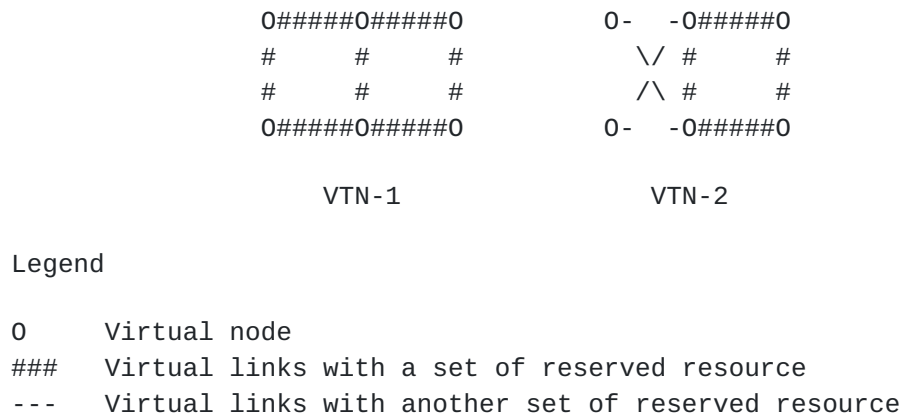


Figure 3. Resource Sharing between VTNs



Figure 3 gives another example of multiple VTNs which shares the same set of network resources on some links. In this case, information about the reserved resource on each link only needs to be advertised once, then both VTN-1 and VTN-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 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 systems could be improved.

#### **4.2. Data Plane Optimization**

To support more VPN+ services while keeping the amount of data plane state in a reasonable scale, one possible approach is to classify a set of VPN+ services which has similar service characteristics and performance requirements into a group, and such group of VPN+ is mapped to one VTN, which is allocated with an aggregated set of network topology and resources to meet the service requirement of the whole group of VPN+. Different groups of VPN+ need to 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 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 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 packet. Then the existing identifier in the packet header used for topology based forwarding is kept unchanged. The benefit is the number of existing topology-specific identifiers will only increase in proportion to the number of topologies rather than the number of VTNs, so that its scalability 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 4 shows the concept of using different data plane identifiers for topology-based and VTN resource-based packet processing respectively.





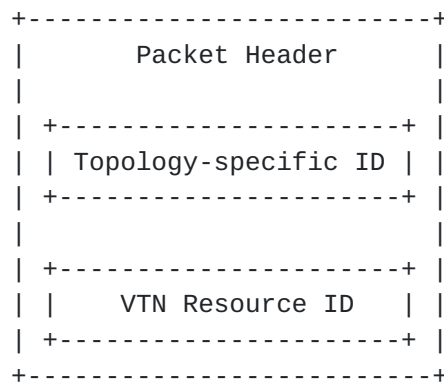


Figure 4. Decoupled Data Plane Identifiers

In an IPv6 [[RFC8200](#)] based network, this could be achieved by introducing a dedicated field in either 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 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.

## 5. Solution Evolution for Improved Scalability

Based on the analysis in this document, the control plane and data plane for VPN+ needs to evolve to support the increasing number of VPN+ services in the network.

For example, by introducing resource-awareness to segment routing SIDs [[I-D.ietf-spring-resource-aware-segments](#)], 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 VPN+ in the network. Such mechanism is considered as basic SR-VTN.

As the number of required VPN+ services increases, more VTNs needs to be created, then the control plane scalability could be improved by decoupling the topology attribute from other attributes (e.g. resource attribute) of VTN, so that multiple VTNs could share the same topology or resource attribute.



To further improve the data plane scalability, dedicated data plane identifiers of VTN can be introduced to decouple the topology-specific forwarding and the VTN resource-based processing in data plane.

## **6. Security Considerations**

TBD

## **7. IANA Considerations**

This document makes no request of IANA.

## **8. Contributors**

Zhibo Hu  
Email: huzhibo@huawei.com

## **9. Acknowledgments**

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

## **10. References**

### **10.1. Normative References**

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

### **10.2. Informative References**

- [I-D.dong-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 for Enhanced VPN", [draft-dong-spring-sr-for-enhanced-vpn-10](#) (work in progress), August 2020.
- [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-13](#) (work in progress), October 2020.



[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", [draft-ietf-spring-resource-aware-segments-00](#) (work in progress), July 2020.

[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-06](#) (work in progress), July 2020.

[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](mailto:jie.dong@huawei.com)



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

Email: lizhenbin@huawei.com

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

Email: qinfengwei@chinamobile.com

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

Email: yangguangm@chinatelecom.cn



