

Workgroup: TEAS Working Group

Internet-Draft:

draft-ietf-teas-enhanced-vpn-11

Published: 20 September 2022

Intended Status: Informational

Expires: 24 March 2023

Authors: J. Dong S. Bryant Z. Li
 Huawei University of Surrey China Mobile
 T. Miyasaka Y. Lee
 KDDI Corporation Samsung

A Framework for Enhanced Virtual Private Network (VPN+)

Abstract

This document describes the framework for Enhanced Virtual Private Network (VPN+). The purpose of VPN+ is to support the needs of new applications (e.g. low latency, bounded jitter, etc.), by utilizing an approach that is based on the VPN and Traffic Engineering (TE) technologies, and adds characteristics that specific services require beyond those provided by existing VPNs. Typically, VPN+ will be used to underpin network slicing, but could also be of use in its own right providing enhanced connectivity services between customer sites. This document also provides an overview of relevant technologies in different network layers, and identifies some areas for potential new work.

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 24 March 2023.

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](#)
2. [Terminology](#)
3. [Overview of the Requirements](#)
 - 3.1. [Performance Guarantees](#)
 - 3.2. [Isolation between VPN+ Services](#)
 - 3.2.1. [A Pragmatic Approach to Isolation](#)
 - 3.3. [Integration with Network Resources and Functions](#)
 - 3.3.1. [Abstraction](#)
 - 3.4. [Dynamic Changes](#)
 - 3.5. [Customized Control](#)
 - 3.6. [Applicability to Overlay Technologies](#)
 - 3.7. [Inter-Domain and Inter-Layer Network](#)
4. [The Architecture of VPN+](#)
 - 4.1. [Layered Architecture](#)
 - 4.2. [Multi-Point to Multi-Point \(MP2MP\) Connectivity](#)
 - 4.3. [Application Specific Data Types](#)
 - 4.4. [Scaling Considerations](#)
5. [Candidate Technologies](#)
 - 5.1. [Forwarding Resource Partitioning](#)
 - 5.1.1. [Flexible Ethernet](#)
 - 5.1.2. [Dedicated Queues](#)
 - 5.1.3. [Time Sensitive Networking](#)
 - 5.2. [Data Plane Encapsulation and Forwarding](#)
 - 5.2.1. [Deterministic Networking](#)
 - 5.2.2. [MPLS Traffic Engineering \(MPLS-TE\)](#)
 - 5.2.3. [Segment Routing](#)
 - 5.3. [Non-Packet Data Plane](#)
 - 5.4. [Control Plane](#)
 - 5.5. [Management Plane](#)
 - 5.6. [Applicability of Service Data Models to VPN+](#)
6. [Applicability in Network Slice Realization](#)
 - 6.1. [VTN Planning](#)
 - 6.2. [VTN Instantiation](#)
 - 6.3. [VPN+ Service Provisioning](#)
 - 6.4. [Network Slice Traffic Steering and Forwarding](#)
7. [Scalability Considerations](#)
 - 7.1. [Maximum Stack Depth of SR](#)

- [7.2. RSVP-TE Scalability](#)
- [7.3. SDN Scaling](#)
- [8. Manageability Considerations](#)
 - [8.1. OAM Considerations](#)
 - [8.2. Telemetry Considerations](#)
- [9. Enhanced Resiliency](#)
- [10. Operational Considerations](#)
- [11. Security Considerations](#)
- [12. IANA Considerations](#)
- [13. Contributors](#)
- [14. Acknowledgements](#)
- [15. References](#)
 - [15.1. Normative References](#)
 - [15.2. Informative References](#)
- [Authors' Addresses](#)

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 common (base) network that is used to provide the VPNs is often referred to as the underlay, and the VPN is often called an overlay.

Customers of a network operator may request connectivity services with advanced characteristics such as low latency guarantees, bounded jitter, or isolation from other services or customers so that changes in some other services (e.g. changes in network load, or events such as congestion or outages) have no or only acceptable effect on the throughput or latency of the services provided to the customer. These services are referred to as "enhanced VPNs" (known as VPN+) in that they are similar to VPN services providing the customer with the required connectivity, but in addition they have enhanced characteristics.

The concept of network slicing has gained traction driven largely by needs surfacing from 5G [[NGMN-NS-Concept](#)] [[TS23501](#)] [[TS28530](#)]. According to [[TS28530](#)], a 5G end-to-end network slice consists of three major types of network segments: Radio Access Network (RAN), Transport Network (TN), and Mobile Core Network (CN). The transport network provides the connectivity between different entities in RAN and CN segments of a 5G end-to-end network slice, with specific performance commitment.

[[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. An IETF Network Slice Service enables connectivity between a set of Service Demarcation Points (SDPs) with

specific Service Level Objectives (SLOs) and Service Level Expectations (SLEs) over a common underlay network. An IETF Network Slice can be realized as a logical network connecting a number of endpoints and is associated with a set of shared or dedicated network resources that are used to satisfy the Service Level Objectives (SLOs) and Service Level Expectations (SLEs) requirements. In this document (which is solely about IETF technologies) we refer to an "IETF Network Slice" simply as a "network slice": a network slice is considered as one typical use case of VPN+.

A network slice may involve multiple technologies (e.g. IP or Optical), and may span multiple administrative domains. Depending on the customer's requirements, a network slice could be isolated from other network slices in terms of data plane, control plane, and management plane resources.

Network slicing can build on the concepts of resource management, network virtualization, and abstraction to provide performance assurance, flexibility, programmability, and modularity. It may use techniques such as Software Defined Networking (SDN) [[RFC7149](#)], network abstraction [[RFC7926](#)], and Network Function Virtualization (NFV) [[RFC8172](#)] [[RFC8568](#)] to create multiple logical (virtual) networks, each tailored for use by a set of services or by a particular tenant or a group of tenants that share the same or similar requirements. These logical networks are created on top of a common underlay network. How the network slices are engineered is deployment-specific.

The requirements of VPN+ services cannot simple be met by overlay networks, as these services require tighter coordination and integration between the overlay and the underlay networks.

In the overlay network, VPN has been defined as the network construct to provide the required connectivity for different services or customers. In the underlay network, this document introduces the concept Virtual Transport Network (VTN) A VTN is a virtual underlay network which is associated with a network topology, and is allocated with a set of dedicated or shared resources from the physical underlay network.

A VPN+ service is realized by integrating a VPN in the overlay and a VTN in the underlay. With this, a VPN+ service can provide enhanced properties such as guaranteed resources and assured or predictable performance. A VPN+ service may also include a set of service functions. These enhanced properties have general applicability, and are also of interest as part of a network slicing solution. Hence VPN+ techniques can be used to instantiate a network slice service,

and they can also be of use in general cases to provide enhanced connectivity services between customer sites or service end points.

[[I-D.ietf-teas-ietf-network-slices](#)] introduces the concept Network Resource Partition (NRP) as a collection of resources in the underlay network that can reliably support specific IETF Network Slice SLAs. An NRP can be associated with a network topology to select or specify the set of links and nodes involved. VTN and NRP are considered as similar concepts, and NRP can be seen as an instantiation of VTN in the context of network slicing.

It is not envisaged that VPN+ services will replace VPN services. VPN services will continue to be delivered using existing mechanisms and can co-exist with VPN+ services.

This document describes a framework for using existing, modified, and potential new technologies as components to provide VPN+ services. Specifically, this document provides:

- *The functional requirements and service characteristics of a VPN+ service.
- *The design of the data plane for VPN+.
- *The necessary control and management protocols in both the underlay and the overlay of VPN+.
- *The mechanisms to achieve integration between overlay and underlay.
- *The necessary Operation, Administration, and Management (OAM) methods to instrument a VPN+ to make sure that the required Service Level Agreement (SLA) between the customer and the network operator is met, and to take any corrective action (such as switching traffic to an alternate path) to avoid SLA violation.

The required layered network structure to achieve these objectives is shown in [Section 4.1](#).

2. Terminology

In this document, the relationship of the four terms "VPN", "VPN+", "VTN", and "Network Slice" are as follows:

- *A Virtual Private Network (VPN) refers to the overlay network service that provides connectivity between different customer sites, and that maintains traffic separation between different customers. The typical VPN technologies are: IPVPN [[RFC2764](#)], L2VPN [[RFC4664](#)], L3VPN [[RFC4364](#)], and EVPN [[RFC7432](#)].

*An enhanced VPN (VPN+) service is an evolution of the VPN service that makes additional service-specific commitments. An enhanced VPN is made by integrating a VPN with a set of network resources allocated in the underlay network.

*A Virtual Transport Network (VTN) is a virtual underlay network which is associated with a logical network topology, and is allocated with a set of dedicated or shared network resources from the physical underlay network. A VTN has the capability of delivering the network resources and performance characteristics required by the VPN+ customers.

*A network slice service could be delivered by provisioning a VPN+ service in the network. Other mechanisms for delivering network slices may exist but are not in scope for this document.

The term "tenant" is used in this document to refer to the customers and all of their associated VPN+ services.

The following terms are also used in this document. Some of them are newly defined, some others reference existing definitions.

SLA: Service Level Agreement. See [[I-D.ietf-teas-ietf-network-slices](#)].

SLO: Service Level Objective. See [[I-D.ietf-teas-ietf-network-slices](#)].

SLE: Service Level Expectation. See [[I-D.ietf-teas-ietf-network-slices](#)].

ACTN: Abstraction and Control of Traffic Engineered Networks [[RFC8453](#)]

DetNet: Deterministic Networking. See [[DETNET](#)] and [[RFC8655](#)]

FlexE: Flexible Ethernet [[FLEXE](#)]

TSN: Time Sensitive Networking [[TSN](#)]

VN: Virtual Network. See [[RFC8453](#)]

VTP: Virtual Transport Path. A VTP is a path through the VTN which provides the required connectivity and performance between two or more customer sites.

3. Overview of the Requirements

This section provides an overview of the requirements of a VPN+ service.

3.1. Performance Guarantees

Performance guarantees are made by network operators to their customers in relation to the services provided to the customers. They are usually expressed in SLAs as a set of SLOs.

There are several kinds of performance guarantee, including guaranteed maximum packet loss, guaranteed maximum delay, and guaranteed delay variation. Note that these guarantees apply to conformance traffic, out-of-profile traffic will be handled according to a separate agreement with the customer. (see for example, section 3.6 of [[RFC7297](#)]).

Guaranteed maximum packet loss is usually addressed by setting packet priorities, queue size, and discard policy. However this becomes more difficult when the requirement is combined with latency requirements. The limiting case is zero congestion loss, and that is the goal of DetNet [[DETNET](#)] and TSN [[TSN](#)]. In modern optical networks, loss due to transmission errors already approaches zero, but there is the possibility of failure of the interface or the fiber itself. This type of fault can be addressed by some form of signal duplication and transmission over diverse paths.

Guaranteed maximum latency is required by a number of applications particularly real-time control applications and some types of virtual reality applications. DetNet [[DETNET](#)] is relevant, however additional methods of enhancing the underlay to better support the delay guarantees may be needed, and these methods will need to be integrated with the overall service provisioning mechanisms.

Guaranteed maximum delay variation is a performance guarantee that may also be needed. [[RFC8578](#)] calls up a number of cases that need this guarantee, for example in electrical utilities. Time transfer is an example service that needs a performance guarantee, although it is in the nature of time that the service might be delivered by the underlay as a shared service and not provided through different VPN+s. Alternatively, a dedicated VPN+ might be used to provide this as a shared service.

This suggests that a spectrum of service guarantees need to be considered when deploying a VPN+. As a guide to understanding the design requirements we can consider four types of service:

- *Best effort
- *Assured bandwidth
- *Guaranteed latency
- *Enhanced delivery

The best effort service is the basic connectivity service that can be provided by current VPNs.

An assured bandwidth service is a connectivity service in which the bandwidth over some period of time is assured. This could be achieved either simply based on a best effort service with over-capacity provisioning, or it can be based on MPLS traffic engineered label switching paths (TE-LSPs) with bandwidth reservations. Depending on the technique used, however, the bandwidth is not necessarily assured at any instant. Providing assured bandwidth to VPNs, for example by using per-VPN TE-LSPs, is not widely deployed at least partially due to scalability concerns. The more common approach of aggregating multiple VPNs onto common TE-LSPs results in shared bandwidth and so may reduce the assurance of bandwidth to any one service. VPN+ aims to provide a more scalable approach for such services.

A guaranteed latency service has an upper bound to edge-to-edge latency. Assuring the upper bound is sometimes more important than minimizing latency. There are several new technologies that provide some assistance with this performance guarantee. Firstly, the IEEE TSN project [[TSN](#)] introduces the concept of scheduling of delay- and loss-sensitive packets. FlexE [[FLEXE](#)] is also useful to help provide a guaranteed upper bound to latency. DetNet is also of relevance in assuring an upper bound of end-to-end packet latency in network layer. The use of these technologies to deliver VPN+ services needs to be considered when a guaranteed latency service is required.

An enhanced delivery service is a connectivity service in which the underlay network (at layer-3) needs to ensure to eliminate packet loss in the event of equipment or media failures. This may be achieved by delivering a copy of the packet through multiple paths. Such a mechanism may need to be used for VPN+ service.

3.2. Isolation between VPN+ Services

One element of the SLA demanded for VPN+ service may be a guarantee that the service offered to the customer will not be affected by any other traffic flows in the network. This is termed "isolation" in section 3.8 of [[RFC7297](#)], and a customer may express the requirement for isolation as an SLE for network slice service [[I-D.ietf-teas-ietf-network-slices](#)].

One way for a network operator to meet the requirement for isolation is setting and conforming to all the SLOs. For example, traffic congestion (interference from other services) might impact on the latency experienced by a VPN+ customer. Thus, in this example, conformance to a latency SLO would be the primary requirement for

delivery of the VPN+ service, and isolation from other services might be only a means to that end.

Another way for a service provider to meet the isolation requirement is to control the degree to which traffic from one service is isolated from other services in the network.

There is a fine distinction between how isolation is requested by a customer and how it is delivered by the service provider. In general, the customer is interested in service performance and not how it is delivered. Thus, for example, the customer wants specific quality guarantees and is not concerned about how the service provider delivers them. However, it should be noted that some aspects of isolation might be directly measurable by a customer if they have information about the traffic patterns on a number of services supported by the same service provider. Furthermore, a customer may be nervous about disruption caused by other services, contamination by other traffic, or delivery of their traffic to the wrong destinations. In this way, the customer may want to specify (and pay for) the level of isolation provided by the service provider.

Isolation is achieved in the realization of a VPN+ through existing technologies that may be supplemented by new mechanisms. The service provider chooses which processes to use to meet this SLE just as they choose how to meet all other SLOs and SLEs. Isolation may be achieved in the network by various forms of resource partitioning ranging from simple separation of service traffic on delivery (ensuring that traffic is not delivered to the wrong customer), through sharing of resources with some form of safeguards, to dedicated allocation of resources for a specific VPN+ service. For example, interference avoidance may be achieved by network capacity planning, allocating dedicated network resources, traffic policing or shaping, prioritizing in using shared network resources, etc.

The terms hard and soft isolation are used to indicate different levels of isolation within the underlay network. A service has soft isolation if the traffic of one service cannot be received by the customers of another service. The existing IP and MPLS VPNs are examples of services with soft isolation: the network delivers the traffic only to the required customer endpoints. However, with soft isolation, as the network resources are shared, traffic from some services may congest the network, resulting in packet loss and delay for other services. The ability for a service or a group of services to be sheltered from this effect is called hard isolation. Hard isolation may be needed so that applications with exacting requirements can function correctly, despite other demands (perhaps a burst of traffic in another service) competing for the underlying resources. A customer may request different degrees of isolation

ranging from soft isolation to hard isolation. In practice isolation may be delivered on a spectrum between soft and hard, and in some cases soft and hard isolation may be used in a hierarchical manner with one VPN+ service being built on another.

To provide the required level of isolation, resources may need to be reserved in the data plane of the underlay network and dedicated to traffic from a specific VPN+ service or a specific group of VPN+ services. This may introduce scalability concerns both in the implementation (as each VPN+ would need to be tracked in the network) and in how many resources need to be reserved and may be under-used (see [Section 4.4](#)). Thus, some trade-off needs to be considered to provide the isolation between VPN+ services while still allowing reasonable resource sharing.

An optical underlay can offer a high degree of isolation, at the cost of allocating resources on a long-term and end-to-end basis. On the other hand, where adequate isolation can be achieved at the packet layer, this permits the resources to be shared amongst a group of services and only dedicated to a service on a temporary basis.

The next section explores a pragmatic approach to isolation in packet networks.

3.2.1. A Pragmatic Approach to Isolation

A key question is whether it is possible to achieve hard isolation in packet networks that were designed to provide statistical multiplexing through sharing of data plane resources, a significant economic advantage when compared to a dedicated or a Time Division Multiplexing (TDM) network. Clearly, there is no need for a customer to request more isolation than their applications require, and no need for a service provider to provide more isolation than requested by their customer, an approximation to full hard isolation is sufficient in most cases when hard isolation is requested. For example, pseudowires [[RFC3985](#)] emulate services that would have had hard isolation in their native form.

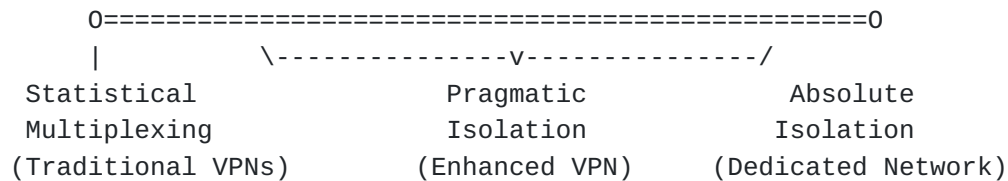


Figure 1: The Spectrum of Isolation

[Figure 1](#) shows a spectrum of isolation that may be delivered by a network. At one end of the spectrum, we see statistical multiplexing technologies that support current VPNs. This is a service type that has served the industry well and will continue to do so. At the opposite end of the spectrum, we have the absolute isolation provided by dedicated transport networks. The goal of VPN+ is "pragmatic isolation". This is isolation that is better than what is obtainable from pure statistical multiplexing, more cost effective and flexible than a dedicated network, but is a practical solution that is good enough for the majority of applications. Mechanisms for both soft isolation and hard isolation are needed to meet different levels of service requirements.

3.3. Integration with Network Resources and Functions

The way to achieve the characteristics demand of a VPN+ service (such as guaranteed or predictable performance) is by integrating the overlay VPN with a particular set of resources in the underlay network which are allocated to meet the service requirements. This needs be done in a flexible and scalable way so that it can be widely deployed in operators' networks to support a good number of VPN+ services.

Taking mobile networks and in particular 5G into consideration, the integration of the network with service functions is likely a requirement. The IETF's work on service function chaining (SFC) [[SFC](#)] provides a foundation for this. Service functions can be considered as part of the VPN+ services. The detailed mechanisms about the integration between service functions and VPN+ are out of the scope of this document.

3.3.1. Abstraction

Integration of the overlay VPN and the underlay network resources and functions does not always need to be a direct mapping. As described in [[RFC7926](#)], abstraction is the process of applying policy to a set of information about a traffic engineered (TE) network to produce selective information that represents the potential ability to connect across the network. The process of abstraction presents the connectivity graph in a way that is independent of the underlying network technologies, capabilities, and topology so that the graph can be used to plan and deliver network services in a uniform way.

With the approach of abstraction, VPN+ may be built on top of an abstracted topology that represents the connectivity capabilities of the underlay TE based network as described in the framework for Abstraction and Control of TE Networks (ACTN) [[RFC8453](#)] as discussed further in [Section 5.5](#).

3.4. Dynamic Changes

VPN+s need to be created, modified, and removed from the network according to service demands. A VPN+ that requires hard isolation ([Section 3.2](#)) must not be disrupted by the instantiation or modification of another VPN+ service. Determining whether modification of an VPN+ can be disruptive to that VPN+, and whether the traffic in flight will be disrupted can be a difficult problem.

Dynamic changes both to the VPN+ and to the underlay network need to be managed to avoid disruption to services that are sensitive to changes in network performance.

In addition to non-disruptively managing the network during changes such as the inclusion of a new VPN+ service endpoint or a change to a link, VPN+ traffic might need to be moved because of changes to traffic patterns and volumes. This means that during the lifetime of a VPN+ service, closed-loop optimization is needed so that the delivered service always matches the ordered service SLA.

The data plane aspects of this problem are discussed further in [Section 5.1](#), [Section 5.2](#), and [Section 5.3](#).

The control plane aspects of this problem are discussed further in [Section 5.4](#).

The management plane aspects of this problem are discussed further in [Section 5.5](#).

3.5. Customized Control

In many cases the customers are delivered with VPN+ services without knowing the information about the underlying VTNs. However, depending on the agreement between the operator and the customer, in some cases the customer may also be provided with some information about the underlying VTNs. Such information can be filtered or aggregated according to the operator's policy. This allows the customer of the VPN+ to have some visibility and even control over how the underlying topology and resources of the VTN are used. For example, the customers may be able to specify the path or path constraints within the VTN for specific traffic flows of their VPN+ service. Depending on the requirements, an VPN+ customer may have his/her own network controller, which may be provided with an interface to the control or management system run by the network operator. Note that such control is within the scope of the customer's VPN+ service, any additional changes beyond this would require some intervention by the network operator.

A description of the control plane aspects of this problem are discussed further in [Section 5.4](#). A description of the management plane aspects of this feature can be found in [Section 5.5](#).

3.6. Applicability to Overlay Technologies

The concept of VPN+ can be applied to any existing and future multi-tenancy overlay technologies including but not limited to :

- *Layer-2 point-to-point services such as pseudowires [[RFC3985](#)]

- *Layer-2 VPNs [[RFC4664](#)]

- *Ethernet VPNs [[RFC7209](#)], [[RFC7432](#)]

- *Layer-3 VPNs [[RFC4364](#)], [[RFC2764](#)]

Where such VPN service types need enhanced isolation and delivery characteristics, the technologies described in [Section 5](#) can be used to tweak the underlay so that to provide the required enhanced performance.

3.7. Inter-Domain and Inter-Layer Network

In some scenarios, a VPN+ service may span multiple network domains. A domain is considered to be any collection of network elements under the responsibility of the same administrative entity, for example, an Autonomous System (AS). In some domains the network operator may manage a multi-layered network, for example, a packet network over an optical network. When VPN+ services are provisioned in such network scenarios, the technologies used in different network planes (data plane, control plane, and management plane) need to provide mechanisms to support multi-domain and multi-layer coordination and integration, so as to provide the required service characteristics for different VPN+ services, and improve network efficiency and operational simplicity.

4. The Architecture of VPN+

A number of VPN+ services will typically be provided by a common network infrastructure. Each VPN+ service is provisioned with an overlay VPN and mapped to a corresponding VTN, which has a specific set of network resources and functions allocated in the underlay to satisfy the needs of the customer. One VTN may support one or more VPN+ services. The integration between the overlay connectivity and the underlay resources ensures the required isolation between different VPN+ services, and achieves the guaranteed performance for different customers.

The VPN+ architecture needs to be designed with consideration given to:

- *An enhanced data plane.
- *A control plane to create VPN+, making use of the data plane isolation and performance guarantee techniques.
- *A management plane for VPN+ service life-cycle management.

These topics are expanded below.

*The enhanced data plane:

- Provides the required packet latency and jitter characteristics.
- Provides the required packet loss characteristics.
- Provides the required resource isolation capability, e.g., bandwidth guarantee.
- Provides the mechanism to associate a packet with the set of resources allocated to a VTN which the VPN+ service packet is mapped to.

*The control plane:

- Collects information about the underlying network topology and network resources, and exports this to network nodes and/or a centralized controller as required.
- Creates VTNs with the network resource and topology properties needed by the VPN+ services.
- Distributes the attributes of VTNs to network nodes which participate in the VTNs and/or a centralized controller.
- Computes and set up network paths in each VTN.
- Maps VPN+ services to an appropriate VTN.
- Determines the risk of SLA violation and takes appropriate avoiding/correction actions.
- Considers the right balance of per-packet and per-node state according to the needs of the VPN+ services to scale to the required size.

*The management plane:

- Provides an interface between the VPN+ service provider (e.g., operator's network management system) and the VPN+ customer (e.g., an organization or a service with VPN+ requirement) such that the operation requests and the related parameters can be exchanged without the awareness of other VPN+ customers.
- Provides an interface between the VPN+ service provider and the VPN+ customers to expose the network capability information toward the customer.
- Provides the service life-cycle management and operation of VPN+ services (e.g., creation, modification, assurance/monitoring, and decommissioning).

*Operations, Administration, and Maintenance (OAM):

- Provides the tools to verify the connectivity and monitor the performance of the VPN+ service.
- Provides the tools to verify whether the underlay network resources are correctly allocated and operating properly.

*Telemetry:

- Provides the mechanisms to collect network information about the operation of the data plane, control plane, and management plane. More specifically, telemetry provides the mechanisms to collect network data:

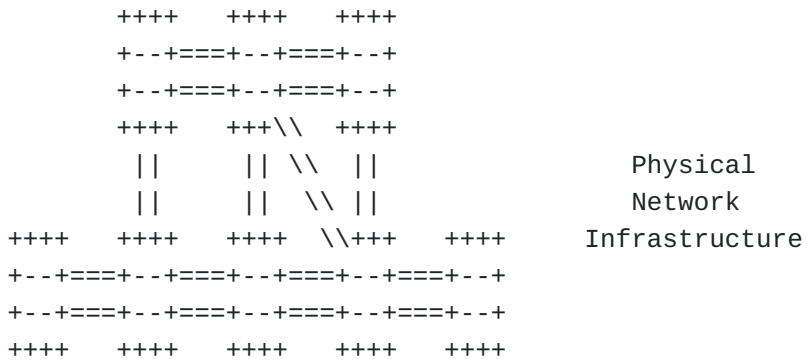
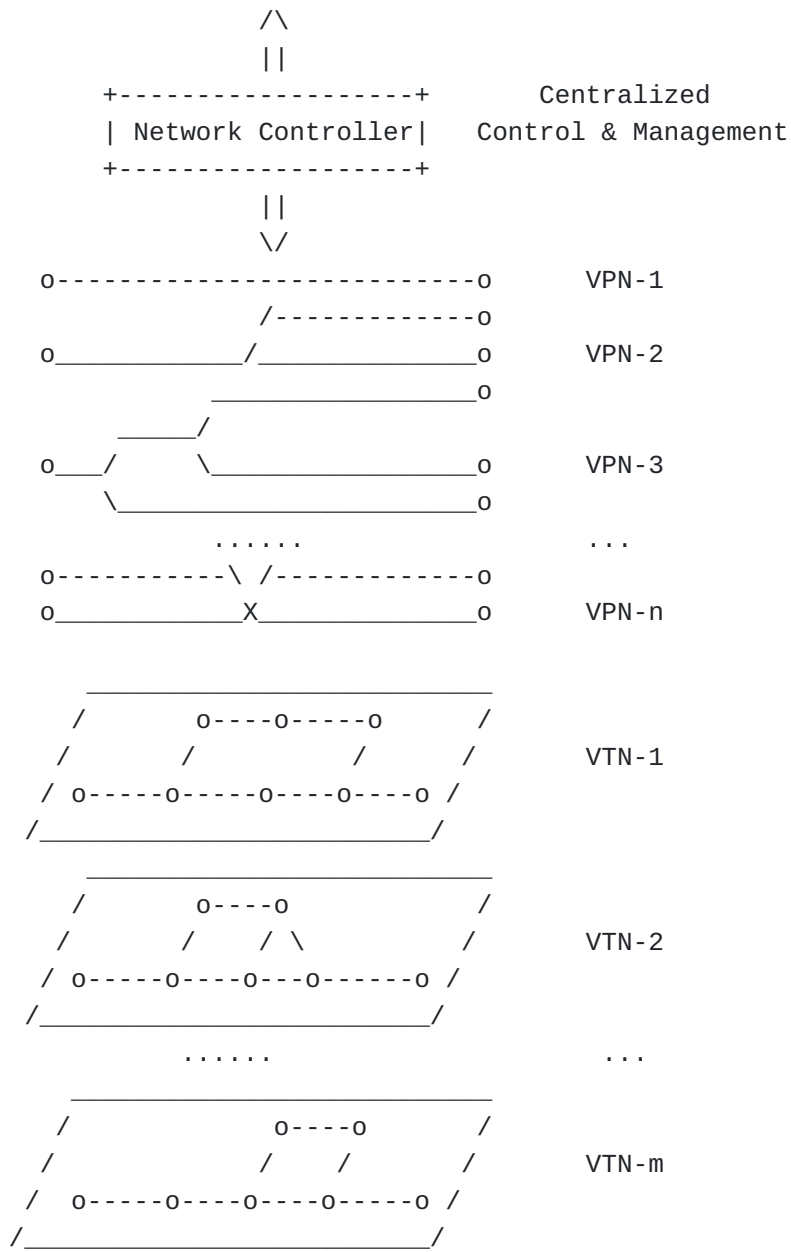
- o from the underlay network for overall performance evaluation and for the planning VPN+ services.

- o from each VPN+ service for monitoring and analytics of the characteristics and SLA fulfillment of the VPN+ services.

4.1. Layered Architecture

The layered architecture of VPN+ is shown in [Figure 2](#).

Underpinning everything is the physical network infrastructure layer which provide the underlying resources used to provision the separated VTNs. This layer is responsible for the partitioning of link and/or node resources for different VTNs. Each subset of link or node resource can be considered as a virtual link or virtual node used to build the VTNs.



o Virtual Node ++++
 +--+ Physical Node with resource partition


```
-- Virtual Link      +--+
                    +---+
                    +---+
                    +---+
== Physical Link with resource partition
```

Figure 2: The Layered Architecture of VPN+

Various components and techniques discussed in [Section 5](#) can be used to enable resource partitioning of the physical network infrastructure, such as FlexE, TSN, dedicated queues, etc. These partitions may be physical or virtual so long as the SLA required by the higher layers is met.

Based on the set of network resource partitions provided by the physical network infrastructure, multiple VTNs can be created, each with a set of dedicated or shared network resources allocated from the physical underlay network, and each can be associated with a customized logical network topology, so as to meet the requirements of different VPN+ services or different groups of VPN+ services. According to the associated logical network topology, each VTN needs to be instantiated on a set of network nodes and links which are involved in the logical topology. And on each node or link, each VTN is associated with a set of local resources which are allocated for the processing of traffic in the VTN. The VTN provides the integration between the logical network topology and the required underlying network resources.

According to the service requirements on connectivity, performance and isolation, etc., VPN services can be mapped to the appropriate VTNs in the network. Different VPN services can be mapped to different VTNs, while it is also possible that multiple VPNs are mapped to the same VTN. Thus the VTN is an essential scaling technique, as it has the potential of eliminating per-path state from the network. In addition, when a group of VPN+ services are mapped to a single VTN, only the network state of the single VTN needs to be maintained in the network (see [Section 4.4](#) for more information).

The network controller is responsible for creating a VTN, instructing the involved network nodes to allocate network resources to the VTN, and provisioning the VPN services on the VTN. A distributed control plane may be used for distributing the VTN resource and topology attributes among nodes in the VTN.

The process used to create VTNs and to allocate network resources for use by the VTNs needs to take a holistic view of the needs of all of the service provider's customers and to partition the resources accordingly. However, within a VTN these resources can, if required, be managed via a dynamic control plane. This provides the required scalability and isolation with some flexibility.

4.2. Multi-Point to Multi-Point (MP2MP) Connectivity

At the VPN service level, the required connectivity for an MP2MP VPN service is usually full or partial mesh. To support such VPN services, the corresponding VTN also needs to provide MP2MP connectivity among the end points.

Other service requirements may be expressed at different granularities, some of which can be applicable to the whole service, while some others may only be applicable to some pairs of end points. For example, when a particular level of performance guarantee is required, the point-to-point path through the underlying VTN of the VPN+ service may need to be specifically engineered to meet the required performance guarantee.

4.3. Application Specific Data Types

Although a lot of the traffic that will be carried over VPN+ will likely be IP based, the design must be capable of carrying other traffic types, in particular Ethernet traffic. This is easily accomplished through the various pseudowire (PW) techniques [[RFC3985](#)]. Where the underlay is MPLS, Ethernet traffic can be carried over VPN+ encapsulated according to the method specified in [[RFC4448](#)]. Where the underlay is IP, Layer Two Tunneling Protocol - Version 3 (L2TPv3) [[RFC3931](#)] can be used with Ethernet traffic carried according to [[RFC4719](#)]. Encapsulations have been defined for most of the common layer-2 types for both PW over MPLS and for L2TPv3.

4.4. Scaling Considerations

VPNs are instantiated as overlays on top of an operator's network and offered as services to the operator's customers. An important feature of overlays is that they can deliver services without placing per-service state in the core of the underlay network.

VPN+ may need to install some additional state within the network to achieve the features that they require. Solutions must consider minimizing and controlling the scale of such state, and deployment architectures should constrain the number of VPN+ services so that the additional state introduced to the network is acceptable and under control. It is expected that the number of VPN+ services will be small at the beginning, and even in the future the number of VPN+ services will be fewer than traditional VPNs because existing VPN techniques are good enough to meet the needs of most existing VPN-type services.

In general, it is not required that the state in the network be maintained in a 1:1 relationship with the VPN+ services. It will usually be possible to aggregate a set or group of VPN+ services so

that they share the same VTN and the same set of network resources (much in the same way that current VPNs are aggregated over transport tunnels) so that collections of VPN+ services that require the same behavior from the network in terms of resource reservation, latency bounds, resiliency, etc. can be grouped together. This is an important feature to assist with the scaling characteristics of VPN+ deployments.

[[I-D.ietf-teas-nrp-scalability](#)] provides more details of scalability considerations for the network resource partitions used to instantiate VTNs, and [Section 7](#) includes a greater discussion of scalability considerations.

5. Candidate Technologies

A VPN is a virtual network created by applying a demultiplexing technique to the underlying network (the underlay) to distinguish the traffic of one VPN from that of another. The connections of VPN are supported by a set of underlay paths. A path that travels by other than the shortest path through the underlay normally requires state to specify that path. The state of the paths could be applied to the underlay through the use of the RSVP-TE signaling protocol, or directly through the use of an SDN controller. Based on Segment Routing, state could be maintained at the ingress node of the path, and carried in the data packet. Other techniques may emerge as this problem is studied. This state gets harder to manage as the number of paths increases. Furthermore, as we increase the coupling between the underlay and the overlay to support the VPN+ service, this state is likely to increase further. We cannot, for example, share the paths and network resource between VPN+ services which require hard isolation.

VTN can be used to provide a group of virtual underlay paths (VTP) with a common set of network resources. Through the use of VTNs, a subset of underlay network resource can be either dedicated for a particular VPN+ service or shared among a group of VPN+ services. This section describes the candidate technologies in different network planes which can be used to build VTNs.

5.1. Forwarding Resource Partitioning

Several candidate layer-2 packet- or frame-based forwarding plane mechanisms which can provide the required resource isolation and performance guarantees are described in the following sections.

5.1.1. Flexible Ethernet

FlexE [[FLEXE](#)] provides the ability to multiplex channels over an Ethernet link to create point-to-point fixed-bandwidth connections

in a way that provides hard isolation. FlexE also supports bonding links to create larger links out of multiple low capacity links.

However, FlexE is only a link level technology. When packets are received by the downstream node, they need to be processed in a way that preserves that isolation in the downstream node. This in turn requires a queuing and forwarding implementation that preserves the end-to-end isolation.

If different FlexE channels are used for different services, then no sharing is possible between the FlexE channels. This means that it may be difficult to dynamically redistribute unused bandwidth to lower priority services in another FlexE channel. If one FlexE channel is used by one customer, the customer can use some methods to manage the relative priority of their own traffic in the FlexE channel.

5.1.2. Dedicated Queues

DiffServ based queuing systems are described in [[RFC2475](#)] and [[RFC4594](#)]. This approach is not sufficient to provide isolation for VPN+ services because DiffServ does not provide enough markers to differentiate between traffic of a large number of VPN+ services. Nor does DiffServ offer the range of service classes that each VPN+ service needs to provide to its tenants. This problem is particularly acute with an MPLS underlay, because MPLS only provides eight traffic classes.

In addition, DiffServ, as currently implemented, mainly provides per-hop priority-based scheduling, and it is difficult to use it to achieve quantitative resource reservation for different VPN+ services.

To address these problems and to reduce the potential interference between VPN+ services, it would be necessary to steer traffic to dedicated input and output queues per VPN+ service or per group of VPN+ services: some routers have a large number of queues and sophisticated queuing systems which could support this, while some routers may struggle to provide the granularity and level of isolation required by the applications of VPN+.

5.1.3. Time Sensitive Networking

Time Sensitive Networking (TSN) [[TSN](#)] is an IEEE project to provide a method of carrying time sensitive information over Ethernet. It introduces the concept of packet scheduling where a packet stream may be given a time slot guaranteeing that it experiences no queuing delay or increase in latency beyond the very small scheduling delay. The mechanisms defined in TSN can be used to meet the requirements of time sensitive traffic flows of VPN+ service.

Ethernet can be emulated over a layer-3 network using an IP or MPLS pseudowire. However, a TSN Ethernet payload would be opaque to the underlay and thus not treated specifically as time sensitive data. The preferred method of carrying TSN over a layer-3 network is through the use of deterministic networking as explained in [Section 5.2.1](#).

5.2. Data Plane Encapsulation and Forwarding

This section considers the problem of VPN+ service differentiation and the representation of underlying network resources in the network layer. More specifically, it describes the possible data plane mechanisms to determine the network resources and the logical network topology or paths associated with a VTN.

5.2.1. Deterministic Networking

Deterministic Networking (DetNet) [[RFC8655](#)] is a technique being developed in the IETF to enhance the ability of layer-3 networks to deliver packets more reliably and with greater control over the delay. The design cannot use re-transmission techniques such as TCP since that can exceed the delay tolerated by the applications. DetNet pre-emptively sends copies of the packet over various paths to minimize the chance of all copies of a packet being lost. It also seeks to set an upper bound on latency, but the goal is not to minimize latency. Detnet can be realized over IP data plane [[RFC8939](#)] or MPLS data plane [[RFC8964](#)], and may be used to provide Virtual Transport Paths (VTPs) for VPN+ services.

5.2.2. MPLS Traffic Engineering (MPLS-TE)

MPLS-TE [[RFC2702](#)][[RFC3209](#)] introduces the concept of reserving end-to-end bandwidth for a TE-LSP, which can be used to provide a point-to-point Virtual Transport Path (VTP) across the underlay network to support VPN services. VPN traffic can be carried over dedicated TE-LSPs to provide reserved bandwidth for each specific connection in a VPN, and VPNs with similar behavior requirements may be multiplexed onto the same TE-LSPs. Some network operators have concerns about the scalability and management overhead of MPLS-TE system, especially with regard to those systems that use an active control plane, and this has lead them to consider other solutions for traffic engineering in their networks.

5.2.3. Segment Routing

Segment Routing (SR) [[RFC8402](#)] is a method that prepends instructions to packets at the head-end of a path. These instructions are used to specify the nodes and links to be traversed, and allow the packets to be routed on paths other than the shortest path. By encoding the state in the packet, per-path

state is transitioned out of the network. SR can be instantiated using MPLS data plane (SR-MPLS) or IPv6 data plane (SRv6).

An SR traffic engineered path operates with a granularity of a link. Hints about priority are provided using the Traffic Class (TC) field in the packet header. However, to achieve the performance and isolation characteristics that are sought by VPN+ customers, it will be necessary to steer packets through specific virtual links and/or queues on the same link and direct them to use specific resources. With SR, it is possible to introduce such fine-grained packet steering by specifying the queues and the associated resources through an SR instruction list.

Note that the concept of a queue is a useful abstraction for different types of underlay mechanism that may be used to provide enhanced isolation and performance support. How the queue satisfies the requirement is implementation specific and is transparent to the layer-3 data plane and control plane mechanisms used.

With Segment Routing, the SR instruction list could be used to build a P2P path, and a group of SR SIDs could also be used to represent an MP2MP network. Thus, the SR based mechanism could be used to provide both a Virtual Transport Path (VTP) and a Virtual Transport Network (VTN) for VPN+ services.

5.3. Non-Packet Data Plane

Non-packet underlay data plane technologies often have TE properties and behaviors, and meet many of the key requirements in particular for bandwidth guarantees, traffic isolation (with physical isolation often being an integral part of the technology), highly predictable latency and jitter characteristics, measurable loss characteristics, and ease of identification of flows. The cost is that the resources are allocated on a long-term and end-to-end basis. Such an arrangement means that the full cost of the resources has to be borne by the client that is allocated with the resources. When a VTN built with this data plane can be used to support multiple VPN+ services, the cost could be distributed among such group of services.

5.4. Control Plane

The control plane of VPN+ would likely be based on a hybrid control mechanism that takes advantage of a logically centralized controller for on-demand provisioning and global optimization, whilst still relying on a distributed control plane to provide scalability, high reliability, fast reaction, automatic failure recovery, etc. Extension to and optimization of the centralized and distributed control plane is needed to support the enhanced properties of VPN+.

As described in section 4, the VPN+ control plane needs to provide the following functions:

- *Collect information about the underlying network topology and network resources, and exports this to network nodes and/or a centralized controller as required.
- *Create VTNs with the network resource and topology properties needed by the VPN+ services.
- *Distribute the attributes of VTNs to network nodes which participate in the VTNs and/or the centralized controller.
- *Map VPN+ services to an appropriate VTN.
- *Compute and set up VTPs in each VTN to meet VPN+ service requirements.

The collection of underlying network topology and resource information can be done using existing the IGP and BGP-LS based mechanisms. The creation of VTN and the distribution of VTN attributes may need further control protocol extensions. The computation of VTPs based on the attributes and constraints of the VTN can be performed either by the headend node of the path or a centralized Path Computation Element (PCE).

There are two candidate mechanisms for the setup of VTPs in the VTN: RSVP-TE and Segment Routing (SR).

- *RSVP-TE [[RFC3209](#)] provides the signaling mechanism for establishing a TE-LSP in an MPLS network with end-to-end resource reservation. This can be seen as an approach of providing a Virtual Transport Path (VTP) which could be used to bind the VPN to specific network resources allocated within the underlay, but there remain scalability concerns as mentioned in [Section 5.2.2](#).
- *The SR control plane [[RFC8665](#)] [[RFC8667](#)] [[RFC9085](#)] does not have the capability of signaling resource reservations along the path. On the other hand, the SR approach provides a potential way of binding the underlay network resource and the VTNs without requiring per-path state to be maintained in the network. A centralized controller can perform resource planning and reservation for VTNs, and it needs to instruct the network nodes to ensure that resources are correctly allocated for the VTN. The controller could provision the SR paths based on the mechanism in [[RFC9256](#)] to the headend nodes of the paths.

According to the service requirements for connectivity, performance and isolation, one VPN+ service may be mapped a dedicated VTN, or a group of VPN+ services may be mapped to the same VTN. The mapping of

VPN+ services to VTN can be achieved using existing control mechanisms with possible extensions, and it can be based on either the characteristics of the data packet or the attributes of the VPN service routes.

5.5. Management Plane

The management plane provides the interface between the VPN+ service provider and the customers for life-cycle management of the VPN+ service (i.e., creation, modification, assurance/monitoring, and decommissioning). It relies on a set of service data models for the description of the information and operations needed on the interface.

As an example, in the context of 5G end-to-end network slicing [[TS28530](#)], the management of the transport network segment of the 5G end-to-end network slice can be realized with the management plane of VPN+. The 3GPP management system may provide the connectivity and performance related parameters as requirements to the management plane of the transport network. It may also require the transport network to expose the capabilities and status of the network slice. Thus, an interface between the VPN+ management plane and the 5G network slice management system, and relevant service data models are needed for the coordination of 5G end-to-end network slice management.

The management plane interface and data models for VPN+ services can be based on the service models described in [Section 5.6](#).

It is important that the management life-cycle supports in-place modification of VPN+ services. That is, it should be possible to add and remove end points, as well as to change the requested characteristics of the service that is delivered. The management system needs to be able to assess the revised VPN+ requests and determine whether they can be provided by the existing VTNs or whether changes must be made, and it will additionally need to determine whether those changes to the VTN are possible. If not, then the customer's modification request may be rejected.

When the modification of a VPN+ service is possible, the management system must make every effort to make the changes in a non-disruptive way. That is, the modification of the VPN+ service or the underlying VTN must not perturbate traffic on the VPN+ service in a way that causes the service level to drop below the agreed levels. Furthermore, in the spirit of isolation, changes to one VPN+ service should not cause disruption to other VPN+ services.

The network operator for the underlay network (i.e., the provider of the VPN+ service) may delegate some operational aspects of the

overlay VPN and the underlying VTN to the customer. In this way, the VPN+ is presented to the customer as a virtual network, and the customer can choose how to use that network. Some mechanisms in the operator's network is needed, so that a customer cannot exceed the capabilities of the virtual links and nodes, but can decide how to load traffic onto the network, for example, by assigning different metrics to the virtual links so that the customer can control how traffic is routed through the virtual network. This approach requires a management system for the virtual network, but does not necessarily require any coordination between the management systems of the virtual network and the physical network, except that the virtual network management system might notice when the VTN is close to capacity or considerably under-used and automatically request changes in the service provided by the underlay network.

5.6. Applicability of Service Data Models to VPN+

This section describes the applicability of the existing and in-progress service data models to VPN+. [\[RFC8309\]](#) describes the scope and purpose of service models and shows where a service model might fit into an SDN based network management architecture. New service models may also be introduced for some of the required management functions.

Service data models are used to represent, monitor, and manage the virtual networks and services enabled by VPN+. The VPN customer service models (e.g., the Layer-3 VPN Service Model (L3SM) [\[RFC8299\]](#), the Layer-2 VPN Service Model (L2SM) [\[RFC8466\]](#)), or the ACTN Virtual Network (VN) model [\[I-D.ietf-teas-actn-vn-yang\]](#)) are service models which can provide the customer's view of the VPN+ service. The Layer-3 VPN Network Model (L3NM) [\[RFC9182\]](#), the Layer-2 VPN network model (L2NM) [\[I-D.ietf-opsawg-l2nm\]](#) provide the operator's view of the managed infrastructure as a set of virtual networks and the associated resources. The NRP model [\[I-D.wd-teas-nrp-yang\]](#) further provides the management of the virtual underlay network topology and resources both in the controller and in the network devices to instantiate the VTNs needed for the VPN+ services.

6. Applicability in Network Slice Realization

One of the typical use cases of VPN+ is to deliver IETF Network Slice Service. This section describes the applicability of VPN+ in network slice realization.

In order to provide IETF network slices to customers, a technology-agnostic network slice service model [\[I-D.ietf-teas-ietf-network-slice-nbi-yang\]](#) is needed for the customers to communicate the requirements of IETF network slices (end points, connectivity, SLOs,

and SLEs). These requirements may be realized using technology specified in this document to instruct the network to instantiate a VPN+ service to meet the requirements of the IETF network slice customers.

6.1. VTN Planning

According to the network operators' network resource planning policy, or based on the requirement of one or a group of customers or services, a VTN may need to be created to support the requested VPN+ services. One of the basic requirements for a VTN is to provide a set of dedicated network resources to avoid unexpected interference from other services in the same network. Other possible requirements may include the required topology and connectivity, bandwidth, latency, reliability, etc.

A centralized network controller can be responsible for calculating a subset of the underlay network topology (which is called a logical topology) to support the VTN requirement. And on the network nodes and links within the logical topology, the set of network resources to be allocated to the VTN can also be determined by the controller. Normally such calculation needs to take the underlay network connectivity information and the available network resource information of the underlay network into consideration. The network controller may also take the status of the existing VTNs into consideration in the planning and calculation of a new VTN.

6.2. VTN Instantiation

According to the result of the VTN planning, the network nodes and links involved in the logical topology of the VTN are instructed to allocate the required set of network resources for the VTN. One or multiple mechanisms as specified in section 5.1 can be used to partition the forwarding plane network resources and allocate different subsets of resources to different VTNs. In addition, the data plane identifiers which are used to identify the set of network resources allocated to the VTN are also provisioned on the network nodes. Depends on the data plane technologies used, the set of network resources of a VTN can be identified using either resource aware SR segments as specified in [[I-D.ietf-spring-resource-aware-segments](#)], or a dedicated VTN resource ID as specified in [[I-D.ietf-6man-enhanced-vpn-vtn-id](#)] can be introduced. The network nodes involved in a VTN may distribute the logical topology information, the VTN specific network resource information and the VTN resource identifiers using the control plane. Such information could be used by the controller and the network nodes to compute the TE or shortest paths within the VTN, and install the VTN specific forwarding entries to network nodes.

6.3. VPN+ Service Provisioning

According to the connectivity requirements of an IETF network slice service, an overlay VPN can be created using the existing or future multi-tenancy overlay technologies as described in [Section 3.6](#).

Then according to the SLO and SLE requirements of the network slice, the overlay VPN is mapped to an appropriate VTN as the virtual underlay. The integration of the overlay VPN and the underlay VTN together provide a VPN+ service which can meet the network slice service requirements.

6.4. Network Slice Traffic Steering and Forwarding

At the edge of the operator's network, traffic of IETF network slices can be classified based on the rules defined by the operator's policy, so that the traffic is treated as a specific VPN+ service, which is further mapped to a underlay VTN. Packets belonging to the VPN+ service will be processed and forwarded by network nodes based the TE or shortest path forwarding entries and the set of network resources of the corresponding VTN.

7. Scalability Considerations

VPN+ provides performance guaranteed services in packet networks, but with the potential cost of introducing additional state into the network. There are at least three ways that this additional state might be brought into the network:

- *Introduce the complete state into the packet, as is done in SR. This allows the controller to specify the detailed series of forwarding and processing instructions for the packet as it transits the network. The cost of this is an increase in the packet header size. The cost is also that systems will have capabilities enabled in case they are called upon by a service. This is a type of latent state, and increases as the path and resources that need to be exclusively available to a VPN are specified more precisely.

- *Introduce the state to the network. This is normally done by creating a path using signaling such as RSVP-TE. This could be extended to include any element that needs to be specified along the path, for example explicitly specifying queuing policy. It is also possible to use other methods to introduce path state, such as via an SDN controller, or possibly by modifying a routing protocol. With this approach there is state per path: per-path characteristic that needs to be maintained over the life of the path. This is more network state than is needed using SR, but the packets are usually shorter.

*Provide a hybrid approach. One example is based on using binding SIDs [[RFC8402](#)] to represent path fragments, and bind them together with SR. Dynamic creation of a VPN service path using SR requires less state maintenance in the network core at the expense of larger packet headers. The packet size can be lower if a form of loose source routing is used (using a few nodal SIDs), and it will be lower if no specific functions or resources on the routers are specified.

Reducing the state in the network is important to VPN+, as it requires the overlay to be more closely integrated with the underlay than with traditional VPNs. This tighter coupling would normally mean that more state needs to be created and maintained in the network, as the state about fine granularity processing would need to be loaded and maintained in the routers. Aggregation is a well-established approach to reduce the amount of state and improve scaling, and VTN is considered as the network construct to aggregate the states of VPN+ services. In addition, an SR approach allows much of the state to be spread amongst the network ingress nodes, and transiently carried in the packets as SIDs.

The following sections describe some of the scalability concerns that need to be considered. Further discussion of the scalability considerations of the underlying network construct of VPN+ can be found in [[I-D.ietf-teas-nrp-scalability](#)].

7.1. Maximum Stack Depth of SR

One of the challenges with SR is the stack depth that nodes are able to impose on packets [[RFC8491](#)]. This leads to a difficult balance between adding state to the network and minimizing stack depth, or minimizing state and increasing the stack depth.

7.2. RSVP-TE Scalability

The established method of creating a resource allocated path through an MPLS network is to use the RSVP-TE protocol. However, there have been concerns that this requires significant continuous state maintenance in the network. Work to improve the scalability of RSVP-TE LSPs in the control plane can be found in [[RFC8370](#)].

There is also concern at the scalability of the forwarder footprint of RSVP-TE as the number of paths through a label switching router (LSR) grows. [[RFC8577](#)] addresses this by employing SR within a tunnel established by RSVP-TE.

7.3. SDN Scaling

The centralized approach of SDN requires state to be stored in the network, but does not have the overhead of also requiring control

plane state to be maintained. Each individual network node may need to maintain a communication channel with the SDN controller, but that compares favorably with the need for a control plane to maintain communication with all neighbors.

However, SDN may transfer some of the scalability concerns from the network to the centralized controller. In particular, there may be a heavy processing burden at the controller, and a heavy load in the network surrounding the controller. A centralized controller also presents a single point of failure within the network.

8. Manageability Considerations

Operations, Administration, and Maintenance (OAM) are necessary functionalities for network and service management. The OAM mechanisms can be further classified as OAM and telemetry. This section describes the considerations of both classes of OAM mechanisms.

8.1. OAM Considerations

The design of OAM for VPN+ services needs to consider the following requirements:

- *Instrumentation of the underlay so that the network operator can be sure that the resources committed to a customer are operating correctly and delivering the required performance.
- *Instrumentation of the overlay by the customer. This is likely to be transparent to the network operator and to use existing methods. Particular consideration needs to be given to the need to verify the isolation and the various committed performance characteristics.
- *Instrumentation of the overlay by the network provider to proactively demonstrate that the committed performance is being delivered. This needs to be done in a non-intrusive manner, particularly when the tenant is deploying a performance sensitive application.
- *Verification of the conformity of the path to the service requirement. This may need to be done as part of a commissioning test.

A study of OAM in SR networks has been documented in [[RFC8403](#)].

8.2. Telemetry Considerations

Network visibility is essential for network operation. Network telemetry has been considered as an ideal means to gain sufficient

network visibility with better flexibility, scalability, accuracy, coverage, and performance than conventional OAM technologies.

As defined in [[RFC9232](#)], the objective of Network Telemetry is to acquire network data remotely for network monitoring and operation. It is a general term for a large set of network visibility techniques and protocols. Network telemetry addresses the current network operation issues and enables smooth evolution toward intent-driven autonomous networks. Telemetry can be applied on the forwarding plane, the control plane, and the management plane in a network.

How the telemetry mechanisms could be used or extended for the VPN+ service is out of the scope of this document.

9. Enhanced Resiliency

Each VPN+ service has a life cycle, and may need modification during deployment as the needs of its tenant change. This is discussed in [Section 5.5](#). Additionally, as the network evolves, there may need to be garbage collection performed to consolidate resources into usable quanta.

Systems in which the path is imposed, such as SR or some form of explicit routing, tend to do well in these applications, because it is possible to perform an atomic transition from one path to another. That is, a single action by the head-end that changes the path without the need for coordinated action by the routers along the path. However, implementations and the monitoring protocols need to make sure that the new path is operational and meets the required SLA before traffic is transitioned to it. It is possible for deadlocks to arise as a result of the network becoming fragmented over time, such that it is impossible to create a new path or to modify an existing path without impacting the SLA of other paths. Resolution of this situation is as much a commercial issue as it is a technical issue and is outside the scope of this document.

There are, however, two manifestations of the latency problem that are for further study in any of these approaches:

- *The problem of packets overtaking one another if a path latency reduces during a transition.

- *The problem of transient variation in latency in either direction as a path migrates.

There is also the matter of what happens during failure in the underlay infrastructure. Fast reroute is one approach, but that still produces a transient loss with a normal goal of rectifying this within 50ms [[RFC5654](#)]. An alternative is some form of N+1

delivery such as has been used for many years to support protection from service disruption. This may be taken to a different level using the techniques of DetNet with multiple in-network replication and the culling of later packets [[RFC8655](#)].

In addition to the approach used to protect high priority packets, consideration should be given to the impact of best effort traffic on the high priority packets during a transition. Specifically, if a conventional re-convergence process is used there will inevitably be micro-loops and whilst some form of explicit routing will protect the high priority traffic, lower priority traffic on best effort shortest paths will micro-loop without the use of a loop prevention technology. To provide the highest quality of service to high priority traffic, either this traffic must be shielded from the micro-loops, or micro-loops must be prevented completely.

10. Operational Considerations

It is likely that VPN+ services will be introduced in networks which already have VPN services deployed. Depending on service requirements, the tenants or the operator may choose to use a VPN or a VPN+ to fulfill a service requirement. The information and parameters to assist such a decision needs to be reflected on the management interface between the tenant and the operator.

11. Security Considerations

All types of virtual network require special consideration to be given to the isolation of traffic belonging to different tenants. That is, traffic belonging to one VPN must not be delivered to end points outside that VPN. In this regard VPN+ neither introduce, nor experience greater security risks than other VPNs.

However, in a VPN+ service the additional service requirements need to be considered. For example, if a service requires a specific upper bound to latency then it can be damaged by simply delaying the packets through the activities of another tenant, i.e., by introducing bursts of traffic for other services. In some respects this makes the VPN+ more susceptible to attacks since the SLA may be broken. But another view is that the operator must, in any case, perform monitoring of the VPN+ to ensure that the SLA is met, and this means that the operator may be more likely to spot the early onset of a security attack and be able to take pre-emptive protective action.

The measures to address these dynamic security risks must be specified as part to the specific solution are form part of the isolation requirements of a service.

While a VPN+ service may be sold as offering encryption and other security features as part of the service, customers would be well advised to take responsibility for their own security requirements themselves possibly by encrypting traffic before handing it off to the service provider.

The privacy of VPN+ service customers must be preserved. It should not be possible for one customer to discover the existence of another customer, nor should the sites that are members of an VPN+ be externally visible.

A VPN+ service (even one with hard isolation requirements) does not provide any additional guarantees of privacy for customer traffic compared to regular VPNs: the traffic within the network may be intercepted and errors may lead to mis-delivery. Users who wish to ensure the privacy of their traffic must take their own precautions including end-to-end encryption.

12. IANA Considerations

There are no requested IANA actions.

13. Contributors

Daniel King
Email: daniel@olddog.co.uk

Adrian Farrel
Email: adrian@olddog.co.uk

Jeff Tansura
Email: jefftant.ietf@gmail.com

Zhenbin Li
Email: lizhenbin@huawei.com

Qin Wu
Email: bill.wu@huawei.com

Bo Wu
Email: lana.wubo@huawei.com

Daniele Ceccarelli
Email: daniele.ceccarelli@ericsson.com

Mohamed Boucadair
Email: mohamed.boucadair@orange.com

Sergio Belotti
Email: sergio.belotti@nokia.com

Haomian Zheng
Email: zhenghaomian@huawei.com

14. Acknowledgements

The authors would like to thank Charlie Perkins, James N Guichard, John E Drake, Shunsuke Homma, and Luis M. Contreras for their review and valuable comments.

This work was supported in part by the European Commission funded H2020-ICT-2016-2 METRO-HAUL project (G.A. 761727).

15. References

15.1. Normative References

[RFC2764] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., and A. Malis, "A Framework for IP Based Virtual Private Networks", RFC 2764, DOI 10.17487/RFC2764, February 2000, <<https://www.rfc-editor.org/info/rfc2764>>.

[RFC3985]

Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <<https://www.rfc-editor.org/info/rfc3985>>.

[RFC4664]

Andersson, L., Ed. and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, DOI 10.17487/RFC4664, September 2006, <<https://www.rfc-editor.org/info/rfc4664>>.

15.2. Informative References

[DETNET]

"Deterministic Networking", March , <<https://datatracker.ietf.org/wg/detnet/about/>>.

[FLEXE]

"Flex Ethernet Implementation Agreement", March 2016, <<http://www.oiforum.com/wp-content/uploads/OIF-FLEXE-01.0.pdf>>.

[I-D.ietf-6man-enhanced-vpn-vtn-id]

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

[I-D.ietf-opsawg-l2nm]

Boucadair, M., Dios, O. G. D., Barguil, S., and L. A. Munoz, "A YANG Network Data Model for Layer 2 VPNs", Work in Progress, Internet-Draft, draft-ietf-opsawg-l2nm-19, 2 June 2022, <<https://www.ietf.org/archive/id/draft-ietf-opsawg-l2nm-19.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-05, 8 September 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-05.txt>>.

[I-D.ietf-teas-actn-vn-yang]

Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Y. Yoon, "A YANG Data Model for Virtual Network (VN) Operations", Work in Progress, Internet-Draft, draft-ietf-teas-actn-vn-yang-15, 11 July 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-actn-vn-yang-15.txt>>.

[I-D.ietf-teas-applicability-actn-slicing]

King, D., Drake, J., Zheng, H., and A. Farrel, "Applicability of Abstraction

and Control of Traffic Engineered Networks (ACTN) to Network Slicing", Work in Progress, Internet-Draft, draft-ietf-teas-applicability-actn-slicing-02, 6 September 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-applicability-actn-slicing-02.txt>>.

[I-D.ietf-teas-ietf-network-slice-nbi-yang] Wu, B., Dhody, D., Rokui, R., Saad, T., and L. Han, "IETF Network Slice Service YANG Model", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-nbi-yang-02, 11 July 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slice-nbi-yang-02.txt>>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-14, 3 August 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-14.txt>>.

[I-D.ietf-teas-nrp-scalability]

Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J. N., Mishra, G., Qin, F., Saad, T., and V. P. Beeram, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-00, 11 July 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-nrp-scalability-00.txt>>.

[I-D.wd-teas-nrp-yang] Wu, B., Dhody, D., and Y. Cheng, "A YANG Data Model for Network Resource Partition (NRP)", Work in Progress, Internet-Draft, draft-wd-teas-nrp-yang-01, 11 July 2022, <<https://www.ietf.org/archive/id/draft-wd-teas-nrp-yang-01.txt>>.

[NGMN-NS-Concept] hao ,, "NGMN NS Concept", 2016, <https://www.ngmn.org/fileadmin/user_upload/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf>.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <<https://www.rfc-editor.org/info/rfc2475>>.

[RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, DOI 10.17487/RFC2702, September 1999, <<https://www.rfc-editor.org/info/rfc2702>>.

- [RFC3209]** Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC3758]** Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, DOI 10.17487/RFC3758, May 2004, <<https://www.rfc-editor.org/info/rfc3758>>.
- [RFC3931]** Lau, J., Ed., Townsley, M., Ed., and I. Goyret, Ed., "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, DOI 10.17487/RFC3931, March 2005, <<https://www.rfc-editor.org/info/rfc3931>>.
- [RFC4364]** Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4448]** Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, DOI 10.17487/RFC4448, April 2006, <<https://www.rfc-editor.org/info/rfc4448>>.
- [RFC4594]** Babiarz, J., Chan, K., and F. Baker, "Configuration Guidelines for DiffServ Service Classes", RFC 4594, DOI 10.17487/RFC4594, August 2006, <<https://www.rfc-editor.org/info/rfc4594>>.
- [RFC4719]** Aggarwal, R., Ed., Townsley, M., Ed., and M. Dos Santos, Ed., "Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4719, DOI 10.17487/RFC4719, November 2006, <<https://www.rfc-editor.org/info/rfc4719>>.
- [RFC5151]** Farrel, A., Ed., Ayyangar, A., and JP. Vasseur, "Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 5151, DOI 10.17487/RFC5151, February 2008, <<https://www.rfc-editor.org/info/rfc5151>>.
- [RFC5654]** Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, DOI 10.17487/RFC5654,

September 2009, <<https://www.rfc-editor.org/info/rfc5654>>.

- [RFC7149] Boucadair, M. and C. Jacquenet, "Software-Defined Networking: A Perspective from within a Service Provider Environment", RFC 7149, DOI 10.17487/RFC7149, March 2014, <<https://www.rfc-editor.org/info/rfc7149>>.
- [RFC7209] Sajassi, A., Aggarwal, R., Uttaro, J., Bitar, N., Henderickx, W., and A. Isaac, "Requirements for Ethernet VPN (EVPN)", RFC 7209, DOI 10.17487/RFC7209, May 2014, <<https://www.rfc-editor.org/info/rfc7209>>.
- [RFC7297] Boucadair, M., Jacquenet, C., and N. Wang, "IP Connectivity Provisioning Profile (CPP)", RFC 7297, DOI 10.17487/RFC7297, July 2014, <<https://www.rfc-editor.org/info/rfc7297>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7926] Farrel, A., Ed., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange between Interconnected Traffic-Engineered Networks", BCP 206, RFC 7926, DOI 10.17487/RFC7926, July 2016, <<https://www.rfc-editor.org/info/rfc7926>>.
- [RFC8172] Morton, A., "Considerations for Benchmarking Virtual Network Functions and Their Infrastructure", RFC 8172, DOI 10.17487/RFC8172, July 2017, <<https://www.rfc-editor.org/info/rfc8172>>.
- [RFC8299] Wu, Q., Ed., Litkowski, S., Tomotaki, L., and K. Ogaki, "YANG Data Model for L3VPN Service Delivery", RFC 8299, DOI 10.17487/RFC8299, January 2018, <<https://www.rfc-editor.org/info/rfc8299>>.
- [RFC8309] Wu, Q., Liu, W., and A. Farrel, "Service Models Explained", RFC 8309, DOI 10.17487/RFC8309, January 2018, <<https://www.rfc-editor.org/info/rfc8309>>.
- [RFC8370] Beeram, V., Ed., Minei, I., Shakir, R., Pacella, D., and T. Saad, "Techniques to Improve the Scalability of RSVP-

TE Deployments", RFC 8370, DOI 10.17487/RFC8370, May 2018, <<https://www.rfc-editor.org/info/rfc8370>>.

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8403] Geib, R., Ed., Filsfils, C., Pignataro, C., Ed., and N. Kumar, "A Scalable and Topology-Aware MPLS Data-Plane Monitoring System", RFC 8403, DOI 10.17487/RFC8403, July 2018, <<https://www.rfc-editor.org/info/rfc8403>>.
- [RFC8453] Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018, <<https://www.rfc-editor.org/info/rfc8453>>.
- [RFC8466] Wen, B., Fioccola, G., Ed., Xie, C., and L. Jalil, "A YANG Data Model for Layer 2 Virtual Private Network (L2VPN) Service Delivery", RFC 8466, DOI 10.17487/RFC8466, October 2018, <<https://www.rfc-editor.org/info/rfc8466>>.
- [RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <<https://www.rfc-editor.org/info/rfc8491>>.
- [RFC8568] Bernardos, C.J., Rahman, A., Zuniga, J.C., Contreras, L.M., Aranda, P., and P. Lynch, "Network Virtualization Research Challenges", RFC 8568, DOI 10.17487/RFC8568, April 2019, <<https://www.rfc-editor.org/info/rfc8568>>.
- [RFC8577] Sitaraman, H., Beeram, V., Parikh, T., and T. Saad, "Signaling RSVP-TE Tunnels on a Shared MPLS Forwarding Plane", RFC 8577, DOI 10.17487/RFC8577, April 2019, <<https://www.rfc-editor.org/info/rfc8577>>.
- [RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <<https://www.rfc-editor.org/info/rfc8578>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF

Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

[RFC8939] Varga, B., Ed., Farkas, J., Berger, L., Fedyk, D., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: IP", RFC 8939, DOI 10.17487/RFC8939, November 2020, <<https://www.rfc-editor.org/info/rfc8939>>.

[RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <<https://www.rfc-editor.org/info/rfc8964>>.

[RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.

[RFC9182] Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <<https://www.rfc-editor.org/info/rfc9182>>.

[RFC9232] Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Network Telemetry Framework", RFC 9232, DOI 10.17487/RFC9232, May 2022, <<https://www.rfc-editor.org/info/rfc9232>>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture",

RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

[SFC] "Service Function Chaining", March , <<https://datatracker.ietf.org/wg/sfc/about>>.

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

[TS28530] "3GPP TS28.530", 2016, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3273>>.

[TSN] "Time-Sensitive Networking", March , <<https://1.ieee802.org/tsn/>>.

Authors' Addresses

Jie Dong
Huawei

Email: jie.dong@huawei.com

Stewart Bryant
University of Surrey

Email: stewart.bryant@gmail.com

Zhenqiang Li
China Mobile

Email: lizhenqiang@chinamobile.com

Takuya Miyasaka
KDDI Corporation

Email: ta-miyasaka@kddi.com

Young Lee
Samsung

Email: younglee.tx@gmail.com