

TEAS Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 14, 2021

J. Dong  
Huawei  
S. Bryant  
Futurewei  
Z. Li  
China Mobile  
T. Miyasaka  
KDDI Corporation  
Y. Lee  
Samsung  
July 13, 2020

**A Framework for Enhanced Virtual Private Networks (VPN+) Service  
draft-ietf-teas-enhanced-vpn-06**

**Abstract**

This document describes the framework for Enhanced Virtual Private Network (VPN+) service. The purpose is to support the needs of new applications, particularly applications that are associated with 5G services, by utilizing an approach that is based on existing VPN and Traffic Engineering (TE) technologies and adds features that specific services require over and above traditional VPNs.

Typically, VPN+ will be used to form the underpinning of network slicing, but could also be of use in its own right providing enhanced connectivity services between customer sites.

It is envisaged that enhanced VPNs will be delivered using a combination of existing, modified, and new networking technologies. This document provides an overview of relevant technologies and identifies some areas for potential new work.

Comparing to traditional VPNs, It is not envisaged that quite large numbers of VPN+ services will be deployed in a network. In other word, it is not intended that all existing VPNs supported by a network will use VPN+ related techniques.

**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 January 14, 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="#">3</a>
<a href="#">2.</a>	Terminologies . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Overview of the Requirements . . . . .	<a href="#">6</a>
<a href="#">3.1.</a>	Isolation between Enhanced VPN Services . . . . .	<a href="#">6</a>
<a href="#">3.1.1.</a>	A Pragmatic Approach to Isolation . . . . .	<a href="#">8</a>
<a href="#">3.2.</a>	Performance Guarantee . . . . .	<a href="#">9</a>
<a href="#">3.3.</a>	Integration . . . . .	<a href="#">10</a>
<a href="#">3.3.1.</a>	Abstraction . . . . .	<a href="#">11</a>
<a href="#">3.4.</a>	Dynamic Management . . . . .	<a href="#">11</a>
<a href="#">3.5.</a>	Customized Control . . . . .	<a href="#">12</a>
<a href="#">3.6.</a>	Applicability . . . . .	<a href="#">12</a>
<a href="#">3.7.</a>	Inter-Domain and Inter-Layer Network . . . . .	<a href="#">12</a>
<a href="#">4.</a>	Architecture of Enhanced VPN . . . . .	<a href="#">13</a>
<a href="#">4.1.</a>	Layered Architecture . . . . .	<a href="#">14</a>
<a href="#">4.2.</a>	Multi-Point to Multi-Point (MP2MP) Connectivity . . . . .	<a href="#">17</a>
<a href="#">4.3.</a>	Application Specific Network Types . . . . .	<a href="#">17</a>
<a href="#">4.4.</a>	Scaling Considerations . . . . .	<a href="#">17</a>
<a href="#">5.</a>	Candidate Technologies . . . . .	<a href="#">18</a>
<a href="#">5.1.</a>	Layer-Two Data Plane . . . . .	<a href="#">18</a>
<a href="#">5.1.1.</a>	Flexible Ethernet . . . . .	<a href="#">18</a>
<a href="#">5.1.2.</a>	Dedicated Queues . . . . .	<a href="#">19</a>
<a href="#">5.1.3.</a>	Time Sensitive Networking . . . . .	<a href="#">19</a>
<a href="#">5.2.</a>	Layer-Three Data Plane . . . . .	<a href="#">20</a>
<a href="#">5.2.1.</a>	Deterministic Networking . . . . .	<a href="#">20</a>



5.2.2.	MPLS Traffic Engineering (MPLS-TE)	20
5.2.3.	Segment Routing	20
5.3.	Non-Packet Data Plane	21
5.4.	Control Plane	21
5.5.	Management Plane	22
5.6.	Applicability of Service Data Models to Enhanced VPN	22
5.6.1.	Network Slice Delivery via Coordinated Service Data Models	23
6.	Scalability Considerations	24
6.1.	Maximum Stack Depth of SR	25
6.2.	RSVP Scalability	25
6.3.	SDN Scaling	25
7.	OAM Considerations	25
8.	Telemetry Considerations	26
9.	Enhanced Resiliency	26
10.	Operational Considerations	27
11.	Security Considerations	27
12.	IANA Considerations	28
13.	Contributors	28
14.	Acknowledgements	29
15.	References	29
15.1.	Normative References	29
15.2.	Informative References	30
Authors'	Addresses	35

## 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 or 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 a connectivity services with advanced characteristics such as enhanced isolation from other services so that changes in some other service (such as changes in network load, or events such as congestion or outages) have no or acceptable effect on the throughput or latency of the services provided to the customer. These services are "enhanced VPNs" (known as VPN+) in that they are similar to VPN services as they provide the customer with required connectivity, but have enhanced characteristics.

Driven largely by needs surfacing from 5G, the concept of network slicing has gained traction [NGMN-NS-Concept] [TS23501] [TS28530] [BBF-SD406]. According to [TS28530], a 5G end-to-end network slice consists of three major types network segments: Radio Access Network (RAN), Transport Network (TN) and Mobile Core Network (CN). The



transport network provides the required connectivity between different entities in RAN and CN segments of an end-to-end network slice, with specific performance commitment.

A transport network slice is a virtual (logical) network with a particular network topology and a set of shared or dedicated network resources, which are used to provide the network slice consumer with the required connectivity, appropriate isolation and specific Service Level Objective (SLO).

A transport network slice could span multiple technologies (such as IP or Optical) and multiple administrative domains. Depending on the consumer's requirement, a transport network slice could be isolated from other, often concurrent transport network slices in terms of data plane, control plane, and management plane resources.

In this document the term "network slice" refers to a transport network slice, and is considered as one typical use case of enhanced VPN.

Network slicing builds on the concept 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 a set of services or a particular tenant or a group of tenants that share the same or similar set of requirements, on top of a common network. How the network slices are engineered can be deployment-specific.

VPN+ could be used to form the underpinning of transport network slice, but could also be of use in general cases providing enhanced connectivity services between customer sites.

The requirement of enhanced VPN services cannot be met by simple overlay networks, as they require tighter coordination and integration between the underlay and the overlay network. VPN+ is built from a VPN overlay and a underlying Virtual Transport Network (VTN) which has a customized network topology and a set of dedicated or shared network resources. It may optionally include a set of invoked service functions allocated from the underlay network. Thus an enhanced VPN can achieve greater isolation with strict performance guarantees. These new properties, which have general applicability, may also be of interest as part of a network slicing solution. It is not envisaged that VPN+ services will replace traditional VPN services that can continue to be deployed using pre-existing mechanisms.



This document specifies a framework for using existing, modified, and potential new technologies as components to provide a VPN+ service. Specifically we are concerned with:

- o The design of the enhanced data plane.
- o The necessary protocols in both the underlay and the overlay of the enhanced VPN.
- o The mechanisms to achieve integration between overlay and underlay.
- o The necessary Operation, Administration, and Management (OAM) methods to instrument an enhanced VPN to make sure that the required Service Level Agreement (SLA) is met, and to take any corrective action to avoid SLA violation, such as switching to an alternate path.

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

Note that, in this document, the relationship of the four terms "VPN", "VPN+", "VTN", and "Transport Network Slice" are described as below:

- o A VPN refers to the overlay virtual private network which provides the required service connectivity and traffic separation between different VPN customers.
- o A Virtual Transport Network (VTN) is a virtual underlay network that connects customer edge points with the additional capability of providing the isolation and performance characteristics required by an enhanced VPN customer.
- o An enhanced VPN (VPN+) can be considered as an evolution of VPN service, but with additional service-specific commitments. An enhanced VPN (VPN+) is made by integrating an overlay VPN and a VTN with a set of network resources allocated in the underlay network.
- o A transport network slice could be provided with an enhanced VPN (VPN+).

## [2.](#) Terminologies

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





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

Detnet: Deterministic Networking [[DETNET](#)]

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

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

VN: Virtual Network [[I-D.ietf-teas-actn-vn-yang](#)]

VPN: Virtual Private Network. IPVPN is defined in [[RFC2764](#)], L2VPN is defined in [[RFC4664](#)].

VPN+: Enhanced VPN service. An enhanced VPN service (VPN+) can be considered as an evolution of VPN service, but with additional service-specific commitments such as enhanced isolation and performance guarantee.

VTP: Virtual Transport Path. A VTP is a virtual underlay path which connects two customer edge points with the capability of providing the isolation and performance characteristics required by an enhanced VPN customer. A VTP usually has a customized path with a set of reserved network resources along the path.

VTN: Virtual Transport Network. A VTN is a virtual underlay network that connects customer edge points with the capability of providing the isolation and performance characteristics required by an enhanced VPN customer. A VTN usually has a customized topology and a set of dedicated or shared network resources.

### **3. Overview of the Requirements**

In this section we provide an overview of the requirements of an enhanced VPN service.

#### **3.1. Isolation between Enhanced VPN Services**

One element of the SLA demanded for an enhanced VPN is a guarantee that the service offered to the customer will not be perturbed by any other traffic flows in the network. One way for a service provider to guarantee the customer's SLA is by controlling the degree of isolation from other services in the network. Isolation is a feature that can be requested by customers. There are different grades of how isolation may be enabled by a network operator and that may result in different levels of service perceived by the customer. These range from simple separation of service traffic on delivery (ensuring that traffic is not delivered to the wrong customer), all



the way to complete separation within the underlay so that the traffic from different services use distinct network resources.

The terms hard and soft isolation are used to illustrate different levels of isolation. A VPN has soft isolation if the traffic of one VPN cannot be received by the customers of another VPN. Both IP and MPLS VPNs are examples of VPNs with soft isolation: the network delivers the traffic only to the required VPN endpoints. However, with soft isolation, traffic from VPNs and regular non-VPN traffic may congest the network resulting in packet loss and delay for other VPNs operating normally. The ability for a VPN service or a group of VPN services to be sheltered from this effect is called hard isolation, and this property is required by some applications. Hard isolation is needed so that applications with exacting requirements can function correctly, despite other demands (perhaps a burst of traffic in another VPN) competing for the underlying resources. In practice isolation may be offered as a spectrum between soft and hard, and in some cases soft and hard isolation may be used in a hierarchical manner. An operator may offer its customers a choice of different degrees of isolation ranging from soft isolation up to hard isolation.

An example of the requirement for hard isolation is a network supporting both emergency services and public broadband multi-media services. During a major incident the VPNs supporting these services would both be expected to experience high data volumes, and it is important that both make progress in the transmission of their data. In these circumstances the VPN services would require an appropriate degree of isolation to be able to continue to operate acceptably. On the other hand, VPNs servicing ordinary bulk data may expect to contest for network resources and queue packets so that traffic is delivered within SLAs, but with some potential delays and interference.

In order to provide the required level of isolation, resources may have to be reserved in the data plane of the underlay network and dedicated to traffic from a specific VPN or a specific group of VPNs to form different enhanced VPNs in the network. This may introduce scalability concerns, thus some trade-off needs to be considered to provide the required isolation between some enhanced VPNs while still allowing reasonable sharing.

An optical layer 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.



There are several new technologies that provide some assistance with these data plane issues. Firstly there is the IEEE project on Time Sensitive Networking [TSN] which introduces the concept of packet scheduling of delay and loss sensitive packets. Then there is [FLEXE] which provides the ability to multiplex multiple channels over one or more Ethernet links in a way that provides hard isolation. Finally there are advanced queueing approaches which allow the construction of virtual sub-interfaces, each of which is provided with dedicated resource in a shared physical interface. These approaches are described in more detail later in this document.

[Section 3.1.1](#) explores pragmatic approaches to isolation in packet networks.

### **[3.1.1. A Pragmatic Approach to Isolation](#)**

A key question is whether it is possible to achieve hard isolation in packet networks that were never designed to support hard isolation. On the contrary, they were designed to provide statistical multiplexing, a significant economic advantage when compared to a dedicated, or a Time Division Multiplexing (TDM) network. However, there is no need to provide any harder isolation than is required by the applications. An approximation to this requirement is sufficient in most cases. Pseudowires[RFC3985] emulate services that would have had hard isolation in their native form.

This spectrum of isolation is shown in Figure 1:




Figure 1: The Spectrum of Isolation

Figure 1 shows the spectrum of isolation that may be delivered by a network. At one end of the figure, we have traditional statistical multiplexing technologies that support 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 enhanced VPNs is "pragmatic isolation". This is isolation that is better than 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 would be needed to meet different levels of service requirement.

### **3.2. Performance Guarantee**

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 other requirements.

Guaranteed maximum packet loss is a common parameter, and 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 the Deterministic Networking work that the IETF [[DETNET](#)] and IEEE [[TSN](#)] are pursuing. In modern optical networks, loss due to transmission errors already approaches zero, but there are the possibilities of failure of the interface or the fiber itself. This can only be addressed by some form of signal duplication and transmission over diverse paths.

Guaranteed maximum latency is required in a number of applications particularly real-time control applications and some types of virtual reality applications. The work of the IETF Deterministic Networking (DetNet) Working Group [[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 service that may also be needed. [[RFC8578](#)] calls up a number of cases where this is needed, for example in electrical utilities. Time transfer is one example of a service that needs this, 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 enhanced VPNs. Alternatively a dedicated enhanced VPN may be used to provide this as a shared service.

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

- o Best effort
- o Assured bandwidth
- o Guaranteed latency





- o Enhanced delivery

Best effort service is the basic service that current VPNs can provide.

An assured bandwidth service is one in which the bandwidth over some period of time is assured. This can be achieved either simply based on best effort with over-capacity provisioning, or it can be based on TE-LSPs with bandwidth reservation. The instantaneous bandwidth is however, not necessarily assured, depending on the technique used. Providing assured bandwidth to VPNs, for example by using per-VPN TE-LSPs, is not widely deployed at least partially due to scalability concerns. VPN+ aims to provide a more scalable approach for such kind of service.

A guaranteed latency service has a latency upper bound provided by the network. Assuring the upper bound is sometimes more important than minimizing latency. There are several new technologies that provide some assistance with performance guarantee. Firstly there is the IEEE project on Time Sensitive Networking [[TSN](#)] which introduces the concept of packet scheduling of delay and loss sensitive packets. Then the DetNet work is also of greater relevance in assuring upper bound of end-to-end packet latency. Flex Ethernet [[FLEXE](#)] is also useful to provide these guarantees. The usage of such underlying technologies for VPN+ service needs to be considered.

An enhanced delivery service is one in which the underlay network (at Layer 3) attempts to deliver the packet through multiple paths in the hope of eliminating packet loss due to equipment or media failures. Such mechanism may need to be used for VPN+ service.

### **[3.3. Integration](#)**

The only way to achieve the enhanced characteristics provided by an enhanced VPN (such as guaranteed or predicted performance) is by integrating the overlay VPN with a particular set of network resources in the underlay network which are allocated to meet the service requirement. This needs to be done in a flexible and scalable way so that it can be widely deployed in operator networks to support a reasonable number of enhanced VPN customers.

Taking mobile networks and in particular 5G into consideration, the integration of network and the service functions is a likely requirement. The work in IETF SFC working group [[SFC](#)] provides a foundation for this integration.



### **3.3.1. Abstraction**

Integration of the overlay VPN and the underlay network resources does not need to be a tight mapping. As described in [\[RFC7926\]](#), abstraction is the process of applying policy to a set of information about a 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.

Virtual networks can be built on top of an abstracted topology that represents the connectivity capabilities of the underlay 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 Management**

Enhanced VPNs need to be created, modified, and removed from the network according to service demand. An enhanced VPN that requires hard isolation ([Section 3.1](#)) must not be disrupted by the instantiation or modification of another enhanced VPN. Determining whether modification of an enhanced VPN can be disruptive to that VPN, and in particular whether the traffic in flight will be disrupted can be a difficult problem.

The data plane aspects of this problem are discussed further in Sections [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](#).

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

In addition to non-disruptively managing the network as a result of gross change such as the inclusion of a new VPN endpoint or a change to a link, VPN traffic might need to be moved as a result of traffic volume changes.



### **[3.5.](#) Customized Control**

In some cases it is desirable that an enhanced VPN has a customized control plane, so that the tenant of the enhanced VPN can have some control of how the resources and functions allocated to this enhanced VPN are used. For example, the tenant may be able to specify the service paths in his own enhanced VPN. Depending on the requirement, an enhanced VPN may have its own dedicated controller, which may be provided with an interface to the control system provided by the network operator. Note that such control is within the scope of the tenant's enhanced VPN, any change beyond that would require some intervention of the 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**

The technologies described in this document should be applicable to a number types of VPN overlay services such as:

- o Layer 2 point-to-point services such as pseudowires [[RFC3985](#)]
- o Layer 2 VPNs [[RFC4664](#)]
- o Ethernet VPNs [[RFC7209](#)]
- o Layer 3 VPNs [[RFC4364](#)], [[RFC2764](#)]

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

### **[3.7.](#) Inter-Domain and Inter-Layer Network**

In some scenarios, an enhanced VPN services may span multiple network domains. A domain is considered to be any collection of network elements within a common realm of address space or path computation responsibility [[RFC5151](#)]. In some domains the operator may manage a multi-layered network, for example, a packet network over an optical network. When enhanced VPNs 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 enhanced VPNs, and improve network efficiency and operational simplicity.



#### **4. Architecture of Enhanced VPN**

A number of enhanced VPN services will typically be provided by a common network infrastructure. Each enhanced VPN consists of both the overlay and a corresponding VTN with a specific set of network resources and functions allocated in the underlay to satisfy the needs of the VPN tenant. The integration between overlay and various underlay resources ensures the required isolation between different enhanced VPNs, and achieves the guaranteed performance for different services.

An enhanced VPN needs to be designed with consideration given to:

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

These required characteristics are expanded below:

- o Enhanced data plane
  - \* Provides the required resource isolation capability, e.g. bandwidth guarantee.
  - \* Provides the required packet latency and jitter characteristics.
  - \* Provides the required packet loss characteristics.
  - \* Provides the mechanism to associate a packet with the set of resources allocated to the enhanced VPN which the packet belongs.
- o Control plane
  - \* Collect information about the underlying network topology and resources available and export this to nodes in the network and/or the centralized controller as required.
  - \* Create the required virtual transport networks (VTNs) with the resource and properties needed by the enhanced VPN services that are assigned to them.
  - \* Determine the risk of SLA violation and take appropriate avoiding action.





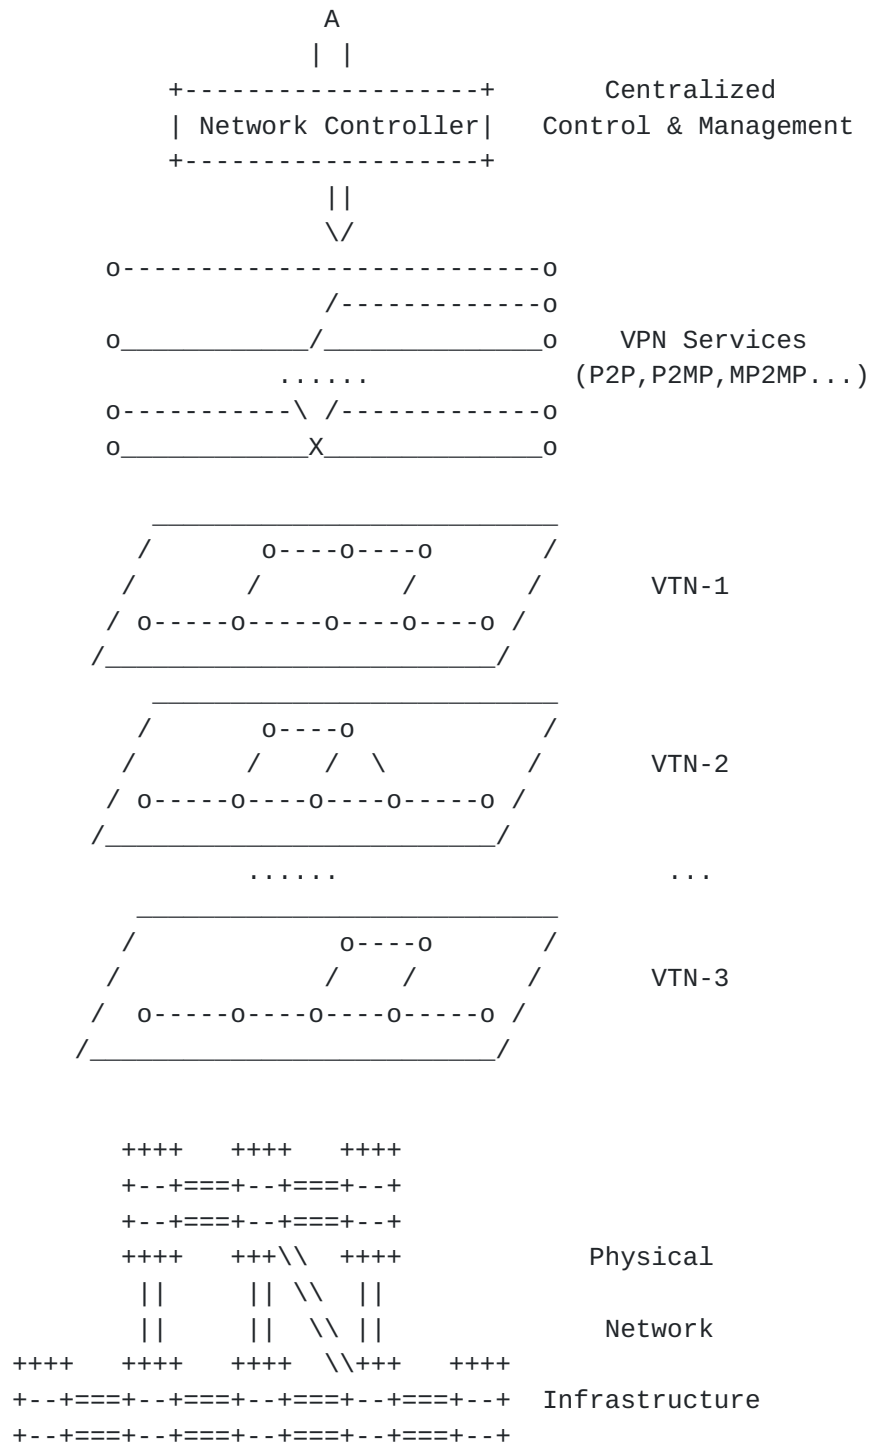
- \* Determine the right balance of per-packet and per-node state according to the needs of enhanced VPN service to scale to the required size.
- o Management plane
  - \* Provides an interface between the enhanced VPN provider (e.g. the Transport Network (TN) Manager) and the enhanced VPN clients (e.g. the 3GPP Management System) such that some of the operation requests can be met without interfering with the enhanced VPN of other clients.
  - \* Provides an interface between the enhanced VPN provider and the enhanced VPN clients to expose transport network capability information toward the enhanced VPN client.
  - \* Provides the service life-cycle management and operation of enhanced VPN (e.g. creation, modification, assurance/monitoring and decommissioning).
- o Operations, Administration, and Maintenance (OAM)
  - \* Provides the OAM tools to verify the connectivity and performance of the enhanced VPN.
  - \* Provide the OAM tools to verify whether the underlay network resources are correctly allocated and operated properly.
- o Telemetry
  - \* Provides the mechanism to collect the data plane, control plane and management plane data of the network, more specifically:
    - \*
    - + Provides the mechanism to collect network data from the underlay network for overall performance evaluation and the enhanced VPN service planning.
    - + Provides the mechanism to collect network data of each enhanced VPN for the monitoring and analytics of the characteristics and SLA fulfilment of enhanced VPN services.

#### **4.1. Layered Architecture**

The layered architecture of an enhanced VPN is shown in Figure 2.



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





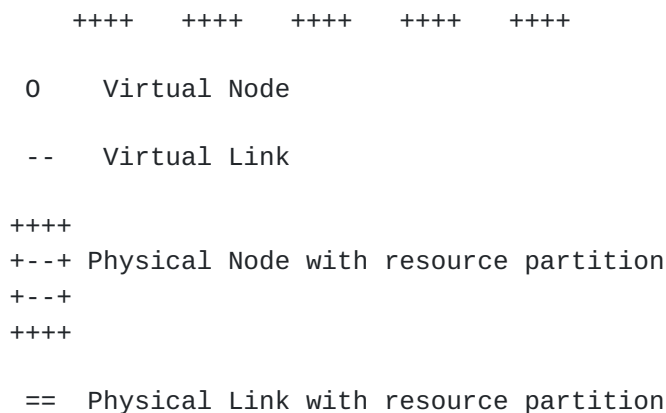


Figure 2: The Layered Architecture of VPN+

Various components and techniques discussed in [Section 5](#) can be used to enable resource partition, such as FlexE, Time Sensitive Networking, Deterministic Networking, 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 network resources provided by the physical network infrastructure, multiple VTNs can be provisioned, each with customized topology and other attributes to meet the requirement of different enhanced VPNs or different groups of enhanced VPNs. To get the required characteristic, each VTN needs to be mapped to a set of network nodes and links in the network infrastructure. And on each node or link, the VTN is associated with a set of resources which are allocated for the processing of traffic in the VTN. VTN provides the integration between the virtual network topology and the required underlying network resources.

The centralized controller is used to create the VTN, and to instruct the network nodes to allocate the required resources to each VTN and to provision the enhanced VPN services on the VTNs. A distributed control plane may also be used for the distribution of the VTN topology and attribute information between nodes within the VTNs.

The process used to create VTNs and to allocate network resources for use by VTNs needs to take a holistic view of the needs of all of its tenants (i.e., of all customers and their associated VTNs), 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.



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

At the VPN service level, the required connectivity is usually mesh or partial-mesh. To support such kinds of VPN service, the corresponding VTN in underlay is also an abstract MP2MP medium. Other service requirements may be expressed at different granularity, some of which can be applicable to the whole service, while some others may be only applicable to some pairs of end points. For example, when particular level of performance guarantee is required, the point-to-point path through the underlay of the enhanced VPN may need to be specifically engineered to meet the required performance guarantee.

#### **4.3. Application Specific Network Types**

Although a lot of the traffic that will be carried over the enhanced VPN will likely be IPv4 or IPv6, the design has to 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 can be carried over the enhanced 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 are able to deliver services without placing per-service state in the core of the underlay network.

Enhanced VPNs may need to install some additional state within the network to achieve the additional features that they require. Solutions must consider minimizing and controlling the scale of such state, and deployment architectures should constrain the number of enhanced VPNs that would exist where such services would place additional state in the network. It is expected that the number of enhanced VPN would be small in the beginning, and even in future the number of enhanced VPN will be much fewer than traditional VPNs, because pre-existing VPN techniques are be 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 of VPN+ services so that they





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

See [Section 6](#) for a greater discussion of scalability considerations.

## 5. Candidate Technologies

A VPN is a network created by applying a demultiplexing technique to the underlying network (the underlay) in order to distinguish the traffic of one VPN from that of another. A VPN path that travels by other than the shortest path through the underlay normally requires state in the underlay to specify that path. State is normally applied to the underlay through the use of the RSVP signaling protocol, or directly through the use of an SDN controller, although other techniques may emerge as this problem is studied. This state gets harder to manage as the number of VPN paths increases. Furthermore, as we increase the coupling between the underlay and the overlay to support the enhanced VPN service, this state will increase further.

In an enhanced VPN different subsets of the underlay resources can be dedicated to different enhanced VPNs or different groups of enhanced VPNs. An enhanced VPN solution thus needs tighter coupling with underlay than is the case with existing VPNs. We cannot, for example, share the network resource between enhanced VPNs which require hard isolation.

### 5.1. Layer-Two Data Plane

A number of candidate Layer 2 packet or frame-based data plane solutions which can be used provide the required isolation and guarantees are described in 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 tenant, the tenant can use some methods to manage the relative priority of his own traffic in the FlexE channel.

#### **5.1.2. Dedicated Queues**

DiffServ based queuing systems are described in [[RFC2475](#)] and [[RFC4594](#)]. This is considered insufficient to provide isolation for enhanced VPNs because DiffServ does not always provide enough markers to differentiate between traffic of many enhanced VPNs, or offer the range of service classes that each VPN 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.

In order to address these problems and to reduce the potential interference between enhanced VPNs, it would be necessary to steer traffic to dedicated input and output queues per enhanced VPN: 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 enhanced VPN.

#### **5.1.3. Time Sensitive Networking**

Time Sensitive Networking (TSN) [[TSN](#)] is an IEEE project that is designing 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. The mechanisms defined in TSN can be used to meet the requirements of time sensitive services of an enhanced VPN.

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. Layer-Three Data Plane**

We now consider the problem of slice differentiation and resource representation in the network layer.

### **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. Even the delay improvements that are achieved with Stream Control Transmission Protocol Partial Reliability Extension (SCTP-PR) [[RFC3758](#)] may not meet the bounds set by application demands. 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.

### **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 point-to-point Virtual Transport Path (VTP) across the underlay network to support VPNs. VPN traffic can be carried over dedicated TE-LSPs to provide reserved bandwidth for each specific connection in a VPN, and VPNs with similar behaviour requirements may be multiplexed onto the same TE-LSPs. Some network operators have concerns about the scalability and management overhead of MPLS-TE system, and this has lead them to consider other solutions for 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.

An SR traffic engineered path operates with a granularity of a link with hints about priority provided through the use of the traffic class (TC) or Differentiated Services Code Point (DSCP) field in the header. However to achieve the latency and isolation characteristics that are sought by the enhanced VPN users, steering packets through specific queues and resources will likely be required. With SR, it



is possible to introduce such fine-grained packet steering by specifying the queues and resources through an SR instruction list.

Note that the concept of queue is a useful abstraction for different types of underlay mechanism that may be used to provide enhanced isolation and latency 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, a group of SR SIDs could also be used to represent a MP2MP network. Thus the SR based mechanism could be used to provide both Virtual Transport Path (VTP) and Virtual Transport Network (VTN) for enhanced VPN services.

### **5.3. Non-Packet Data Plane**

Non-packet underlay data plane technologies often have TE properties and behaviours, 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 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 service that is allocated with the resources.

### **5.4. Control Plane**

Enhanced VPN would likely be based on a hybrid control mechanism, which takes advantage of the 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 distributed control plane is needed to support the enhanced properties of VPN+.

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 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 mentioned in [Section 5.2.2](#).

The control plane of SR [[RFC8665](#)] [[RFC8667](#)] [[I-D.ietf-idr-bgp-ls-segment-routing-ext](#)] 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 enhanced VPN service without requiring per-path state to be maintained in the network. A centralized controller can perform resource planning and reservation for enhanced VPNs, while it needs to ensure that resources are correctly allocated in network nodes for the enhanced VPN service. The controller could also compute the SR paths based on the planned or collected network resource and other attributes, and provision the SR paths based on the mechanism in [\[I-D.ietf-spring-segment-routing-policy\]](#) to the ingress nodes of the enhanced VPN services. The distributed control plane may be used to advertise the network attributes associated with enhanced VPNs, and compute the SR paths with specific constraints of enhanced VPN services.

### **[5.5.](#) Management Plane**

The management plane provides the interface between the enhanced VPN provider and the clients for the service life-cycle management (e.g. 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 enhanced VPNs is considered as the management of the transport network part of the end-to-end network slice. 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 capability and status of the transport network slice. Thus, an interface between the enhanced VPN management plane and the 3GPP network slice management system, and relevant service data models are needed for the coordination of end-to-end network slice management.

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

### **[5.6.](#) Applicability of Service Data Models to Enhanced VPN**

ACTN supports operators in viewing and controlling different domains and presenting virtualized networks to their customers. The ACTN framework [\[RFC8453\]](#) highlights how:

- o Abstraction of the underlying network resources is provided to higher-layer applications and customers.
- o Underlying resources are virtualized and allocated for the customer, application, or service.



- o A virtualized environment is created allowing operators to view and control multi-domain networks as a single virtualized network.
- o Networks can be presented to customers as a virtual network via open and programmable interfaces.

The type of network virtualization enabled by ACTN managed infrastructure provides customers and applications (tenants) with the capability to utilize and independently control allocated virtual network resources as if they were physically their own resources. Service Data models are used to represent, monitor, and manage the virtual networks and services enabled by ACTN. The Customer VPN model (e.g. L3SM [[RFC8299](#)], L2SM [[RFC8466](#)]) or an ACTN Virtual Network (VN) [[I-D.ietf-teas-actn-vn-yang](#)] model is a customer view of the ACTN managed infrastructure, and is presented by the ACTN provider as a set of abstracted services or resources. The L3VPN network model [[I-D.ietf-opsawg-l3sm-l3nm](#)] and [[I-D.ietf-opsawg-l2nm](#)] provide a network view of the ACTN managed infrastructure presented by the ACTN provider as a set of transport resources.

#### **5.6.1. Network Slice Delivery via Coordinated Service Data Models**

In order to support network slice service in transport network, a Transport Slice (TS) Northbound Interface (NBI) data model may be needed for a consumer to express the requirements for transport slices, which can be technology-agnostic. Then these requirements may be realized using technology-specific Southbound Interface (SBI).

As per [[RFC8453](#)] and [[I-D.ietf-teas-actn-yang](#)], the CNC-MDSC Interface (CMI) of ACTN is used to convey the virtual network service requirements, which is a generic interface to deliver various TE based VN services. In the context of network slice northbound interface, there may be some gaps in L3SM/L2SM or VN model, or the combination of them. The TS NBI is required to communicate the connectivity of the transport slice, along with the service level objective (SLO) parameters and traffic selection rules, and provides a way to monitor the state of the transport slice. This can be described in a more abstracted manner, so as to reduce the association with specific realization technologies of transport network slice, such as the VPN and TE technologies. The transport slice model as defined in [[I-D.wd-teas-transport-slice-yang](#)] provides an abstracted and generic approach to meet the transport slice NBI requirement.

The MDSC-PNC Interface (MPI) models in the ACTN architecture can be used for the realization of transport slices, for example, in a TE enabled transport network, and may also be used for cross-layer or cross-domain implementation of transport slice.



## 6. Scalability Considerations

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

- o 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 we more precisely specify the path and resources that need to be exclusively available to a VPN.
- o Introduce the state to the network. This is normally done by creating a path using RSVP-TE, which can be extended to introduce any element that needs to be specified along the path, for example explicitly specifying queuing policy. It is possible to use other methods to introduce path state, such as via a Software Defined Network (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 its life-cycle. This is more state than is needed using SR, but the packets are shorter.
- o Provide a hybrid approach. One example is based on using binding SIDs [[RFC8402](#)] to create 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 enhanced 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 needed 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. However, a segment routed approach allows much of this state to be spread amongst the network ingress nodes, and transiently carried in the packets as SIDs.



### **6.1. 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.

### **6.2. RSVP Scalability**

The traditional method of creating a resource allocated path through an MPLS network is to use the RSVP 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 as the number of paths through an LSR grows. [[RFC8577](#)] proposes to address this by employing SR within a tunnel established by RSVP-TE.

### **6.3. 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 favourably 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.

## **7. OAM Considerations**

The enhanced VPN OAM design needs to consider the following requirements:

- o Instrumentation of the underlay so that the network operator can be sure that the resources committed to a tenant are operating correctly and delivering the required performance.
- o Instrumentation of the overlay by the tenant. 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.





- o 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.
- o 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. 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 [[I-D.ietf-opsawg-ntf](#)], 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 enhanced VPN service is out of the scope of this document.

## **9. Enhanced Resiliency**

Each enhanced VPN has a life-cycle, and may need modification during deployment as the needs of its tenant change. Additionally, as the network as a whole 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. This is a single action by the head-end 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 up 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:

- o The problem of packets overtaking one and other if a path latency reduces during a transition.
- o 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 proposed by the IETF deterministic network work with multiple in-network replication and the culling of later packets [[RFC8655](#)].

In addition to the approach used to protect high priority packets, consideration has to be given to the impact of best effort traffic on the high priority packets during a transient. 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.

## **[10.](#) Operational Considerations**

It is likely that enhanced VPN service will be introduced in networks which already have traditional VPN services deployed. Depends on service requirement, the tenants or the operator may choose to use traditional VPN or enhanced VPN to fulfil the service requirement. The information and parameters to assist such decision needs to be reflected on the management interface between the tenants 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 enhanced VPNs neither introduce, nor experience a greater security risks than other VPNs.

However, in an enhanced Virtual Private Network 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.

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 an enhanced 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 enhanced 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 enhanced VPN be externally visible.

## **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 and Shunsuke Homma 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**

- [BBF-SD406] "BBF SD-406: End-to-End Network Slicing", 2016, <<https://wiki.broadband-forum.org/display/BBF/SD-406+End-to-End+Network+Slicing>>.
- [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-idr-bgp-ls-segment-routing-ext] Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", [draft-ietf-idr-bgp-ls-segment-routing-ext-16](#) (work in progress), June 2019.
- [I-D.ietf-opsawg-l2nm] Barguil, S., Dios, O., Boucadair, M., Munoz, L., Jalil, L., and J. Ma, "A Layer 2 VPN Network YANG Model", [draft-ietf-opsawg-l2nm-00](#) (work in progress), July 2020.
- [I-D.ietf-opsawg-l3sm-l3nm] Barguil, S., Dios, O., Boucadair, M., Munoz, L., and A. Aguado, "A Layer 3 VPN Network YANG Model", [draft-ietf-opsawg-l3sm-l3nm-03](#) (work in progress), April 2020.
- [I-D.ietf-opsawg-ntf] Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Network Telemetry Framework", [draft-ietf-opsawg-ntf-03](#) (work in progress), April 2020.



- [I-D.ietf-spring-segment-routing-policy]  
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-08](#) (work in progress), July 2020.
- [I-D.ietf-teas-actn-vn-yang]  
Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., and B. Yoon, "A Yang Data Model for VN Operation", [draft-ietf-teas-actn-vn-yang-08](#) (work in progress), March 2020.
- [I-D.ietf-teas-actn-yang]  
Lee, Y., Zheng, H., Ceccarelli, D., Yoon, B., Dios, O., Shin, J., and S. Belotti, "Applicability of YANG models for Abstraction and Control of Traffic Engineered Networks", [draft-ietf-teas-actn-yang-05](#) (work in progress), February 2020.
- [I-D.wd-teas-transport-slice-yang]  
Bo, W., Dhody, D., Han, L., and R. Rokui, "A Yang Data Model for Transport Slice NBI", [draft-wd-teas-transport-slice-yang-02](#) (work in progress), July 2020.
- [NGMN-NS-Concept]  
"NGMN NS Concept", 2016, <[https://www.ngmn.org/fileadmin/user\\_upload/161010\\_NGMN\\_Network\\_Slicing\\_framework\\_v1.0.8.pdf](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>>.



- [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>>.
- [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>>.
- [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  
Futurewei

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

