TEAS Working Group   J. Dong
Internet-Draft   S. Bryant
Intended status: Informational   Huawei
Expires: July 19, 2019   Z. Li
                                                         China Mobile
                                                         T. Miyasaka
                                                     KDDI Corporation
                                                               Y. Lee
                                                               Huawei
                                                    January 15, 2019

A Framework for Enhanced Virtual Private Networks (VPN+) Service
draft-ietf-teas-enhanced-vpn-00

Abstract

   This document specifies a framework for using existing, modified and
   potential new networking technologies as components to provide an
   Enhanced Virtual Private Networks (VPN+) service. The purpose is to
   enable VPNs to support the needs of new applications, particularly
   applications that are associated with 5G services. Typically, VPN+
   can be used to form the underpinning of network slicing, but will
   also be of use in its own right.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Virtual private networks (VPNs) have served the industry well as a
means of providing different groups of users with logically isolated
access to 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 enhanced VPN services
with additional characteristics such as complete isolation from other
VPNs so that changes in network load have no effect on the throughput
or latency of the service provided to the customer.

Driven largely by needs surfacing from 5G, the concept of network
slicing has gained traction [NGMN-NS-Concept] [TS23501] [TS28530]
[BBF-SD406]. Network slicing requires the underlying network to
support partitioning the network resources to provide the client with
dedicated (private) networking, computing, and storage resources
drawn from a shared pool. The slices may be seen as (and operated
as) virtual networks.

Network abstraction is a technique that can be applied to a network
domain to select network resources by policy to obtain a view of
potential connectivity and a set of service functions.

Network slicing is an approach to network operations that builds on
the concept of network abstraction to provide programmability,
flexibility, and modularity. It may use techniques such as Software
Defined Networking (SDN) [RFC7149] and Network Function
Virtualization (NFV) to create multiple logical (virtual) networks,
each tailored for a set of services or a particular tenant that share
the same set of requirements, on top of a common network. How the
network slices are engineered can be deployment-specific.

Thus, there is a need to create virtual networks with enhanced
characteristics. The tenant of such a virtual network can require a
degree of isolation and performance that previously could only be
satisfied by dedicated networks. Additionally, the tenant may ask
for some level of control to their virtual networks, e.g., to
customize the service forwarding paths in a network slice.

These enhanced properties cannot be met with pure overlay networks, as they require tighter coordination and integration between the underlay and the overlay network. This document introduces a new network service called Enhanced VPN: VPN+. VPN+ refers to a virtual network which has dedicated network resources, including invoked service functions, allocated from the underlay network. Unlike a traditional VPN, an enhanced VPN can achieve greater isolation with strict guaranteed performance. These new properties, which have general applicability, may also be of interest as part of a network slicing solution.

This document specifies a framework for using existing, modified and potential new networking 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 underlay and the overlay of 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) are met, and to take any corrective action to avoid SLA violation, such as switching to an alternate path.

The required network layered structure to achieve this is shown in Section 3.1.

Note that, in this document, the four terms "VPN", "Enhanced VPN" (or "VPN+"), "Virtual Network (VN)", and "Network Slice" may be considered as describing similar concepts dependent on the viewpoint from which they are used.

o An enhanced VPN is clearly a form of VPN, but with additional service-specific commitments.

o A VN is a type of service that connects customer edge points with the additional possibility of requesting further service characteristics in the manner of an enhanced VPN.

o An enhanced VPN or VN is made by creating a slice through the resources of the underlay network.

   o The general concept of network slicing in a TE network is a larger
      problem space than is addressed by VPN+ or VN, but those concepts
      are tools to address some aspects or realizations of network
      slicing.

2. Overview of the Requirements

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

2.1. Isolation between Virtual Networks

   Isolation is a feature requested by some particular customers in the
   network. Such feature is offered by a network operator where the
   traffic from one service instance is isolated from the traffic of
   other services. There are different grades of isolation that range
   from simple separation of 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 introduced to give example of
   different isolation cases. 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 soft isolated VPNs because the
   network delivers the traffic only to the required VPN endpoints.
   However, the traffic from one or more VPNs and regular network
   traffic may congest the network resulting in packet loss and delay
   for other VPNs operating normally. The ability for a VPN to be
   sheltered from this effect is called hard isolation, and this
   property is required by some critical applications.

   The requirement is for an operator to provide both hard and soft
   isolation between the tenants/applications using one enhanced VPN and
   the tenants/applications using another enhanced VPN. Hard isolation
   is needed so that applications with exacting requirements can
   function correctly, despite other demands (perhaps a burst on another
   VPN) competing for the underlying resources. In practice isolation
   may be offered as a spectrum between soft and hard.

   An example of 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 VPNs would require an appropriate degree of isolation to be able
   to continue to operate acceptably.

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

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. Such an
arrangement means that the full cost of the resources must be borne
by the service that is allocated with the resources. On the other
hand, where adequate isolation can be achieved at the packet layer,
this permits the resources to be shared amongst many services and
only dedicated to a service on a temporary basis. This in turn,
allows greater statistical multiplexing of network resources and thus
amortizes the cost over many services, leading to better economy.
However, the degree of isolation required by network slicing cannot
be entirely met with existing mechanisms such as Traffic Engineered
Label Switched Paths (TE-LSPs). This is because most implementations
enforce the bandwidth in the data-plane only at the PEs, but at the P
routers the bandwidth is only reserved in the control plane, thus
bursts of data can accidentally occur at a P router with higher than
committed data rate.

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.

In the remainder of this section we explore how isolation may be
achieved in packet networks.

2.1.1. A Pragmatic Approach to Isolation

A key question is whether it is possible to achieve hard isolation in
packet networks, which 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 application. Pseudowires [RFC3985] emulate services that would

have had hard isolation in their native form. An approximation to
this requirement is sufficient in most cases.

Thus, for example, using FlexE or a channelized sub-interface
together with packet scheduling as interface slicing, optionally
along with the slicing of node resources, a type of hard isolation
can be provided that is adequate for many VPN+ applications. Other
applications may be either satisfied with a classical VPN with or
without reserved bandwidth, or may need dedicated point to point
fiber. The needs of each application must be quantified in order to
provide an economic solution that satisfies those needs without over-
engineering.

This spectrum of isolation is shown in Figure 1:

```
    O=================================================O
    |           \---------------v---------------/
 Statistical              Pragmatic              Absolute
 Multiplexing             Isolation             Isolation
(Traditional VPNs)     (Enhanced VPN)     (Dedicated Network)
```
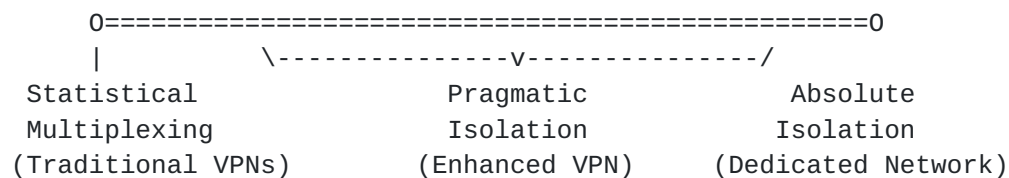
Figure 1: The Spectrum of Isolation

At one end of the above 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 traditional transport networks. The goal of enhanced VPN 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 which is a practical
solution that is good enough for the majority of applications.

2.2. Performance Guarantee

There are several kinds of performance guarantees, including
guaranteed maximum packet loss, guaranteed maximum delay and
guaranteed delay variation. Note that these guarantees apply to the
conformance traffic, the out-of-profile traffic will be handled
following other requirements.

Guaranteed maximum packet loss is a common parameter, and is usually
addressed by setting the packet priorities, queue size and discard
policy. However this becomes more difficult when the requirement is
combined with the latency requirement. 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 is 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 the scope needs
to be extended to methods of enhancing the underlay to better support
the delay guarantee, and to integrate these enhancements with the
overall service provision.

Guaranteed maximum delay variation is a service that may also be
needed. [I-D.ietf-detnet-use-cases] calls up a number of cases where
this is needed, for example electrical utilities have an operational
need for this. 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 virtual networks. Alternatively a dedicated
virtual network 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 TE-LSPs, is
not widely deployed at least partially due to scalability concerns.
Guaranteed latency and enhanced delivery are not yet integrated with
VPNs.

A guaranteed latency service has a latency upper bound provided by the network. Assuring the upper bound is more important than achieving the minimum latency.

In [Section 2.1](#) we considered the work of the IEEE Time Sensitive Networking (TSN) project [TSN] and the work of the IETF DetNet Working group [DETNET] in the context of isolation. The TSN and DetNet work is of greater relevance in assuring end-to-end packet latency. It is also of importance in considering enhanced delivery.

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.

It is these last two characteristics that an enhanced VPN adds to a VPN service.

Flex Ethernet [FLEXE] is a useful underlay to provide these guarantees. This is a method of providing time-slot based channelization over an Ethernet bearer. Such channels are fully isolated from other channels running over the same Ethernet bearer. As noted elsewhere this produces hard isolation but makes the reclamation of unused bandwidth more difficult.

These approaches can be used in tandem. It is possible to use FlexE to provide tenant isolation, and then to use the TSN/Detnet approach to provide a performance guarantee inside the a slice or tenant VPN.

2.3. Integration

A solution to the enhanced VPN problem has to provide close integration of both overlay VPN and the underlay network resource. This needs 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.

2.3.1. 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) described in [RFC8453] as discussed further in Section 4.5.

2.4. Dynamic Configuration

Enhanced VPNs need to be created, modified, and removed from the network according to service demand. An enhanced VPN that requires hard isolation 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 the traffic in flight will be disrupted can be a difficult problem.

The data plane aspects of this problem are discussed further in Section 4.

The control plane aspects of this problem are discussed further in Section 4.4.

The management plane aspects of this problem are discussed further in Section 4.5

Dynamic changes both to the VPN and to the underlay transport network need to be managed to avoid disruption to sensitive services.

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.

2.5. 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 to the resources and functions allocated to this enhanced VPN. 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, or it may be provided with an interface to a control system which is shared with a set of other tenants, or it may be provided with an interface to the control system provided by the network operator.

Further detail on this requirement will be provided in a future
version of the draft. A description of the management plane aspects
of this feature can be found in Section 4.5.

2.6. Applicability

The technologies described in this document should be applicable to a
number types of VPN 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]

o Virtual Networks (VNs) [RFC8453]

Where such VPN or VN types need enhanced isolation and delivery
characteristics, the technology described here can be used to provide
an underlay with the required enhanced performance.

3. 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 specific set of dedicated 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 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 VPN, making use of the data
   plane isolation and 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 identify network slice and the
            associated resources

   o Control plane

         * Collect the underlying network topology and resources available
            and export this to other nodes and/or the centralized
            controller as required.

         * Create the required virtual networks with the resource and
            properties needed by the enhanced VPN services that are
            assigned to it.

         * 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 the life-cycle management (creation, modification,
            decommissioning) of enhanced VPN

         * Provides an interface between the enhanced VPN provider and the
            enhanced VPN clients such that some of the operation requests
            can be met without interfering with the enhanced VPN of other
            clients.

3.1. Layered Architecture

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

```
                   +-------------------+              }
                   | Network Controller|              } Centralized
                   +-------------------+              }   Control

                     .    .   .     .  .
                    .     .   .     .  .
                   .      N----N----N  .              }
                  .      /        /   .               }
             N-----N-----N----N-----N                 }
                   N----N                             }
                  /    / \                            }  Virtual
             N-----N----N----N-----N                  } Networks
                     N----N                           }
                    /    /                            }
             N-----N-----N----N-----N                 }


       +----+ ===== +----+ =====   +----+ ===== +----+  }
       +----+ ===== +----+ =====   +----+ ===== +----+  } Physical
       +----+ ===== +----+ =====   +----+ ===== +----+  } Network
       +----+       +----+         +----+       +----+  }
         N      L     N      L       N      L     N

       N = Partitioned node
       L = Partitioned link

       +----+ = Partition within a node
       +----+

       ====== = Partition within a link
```
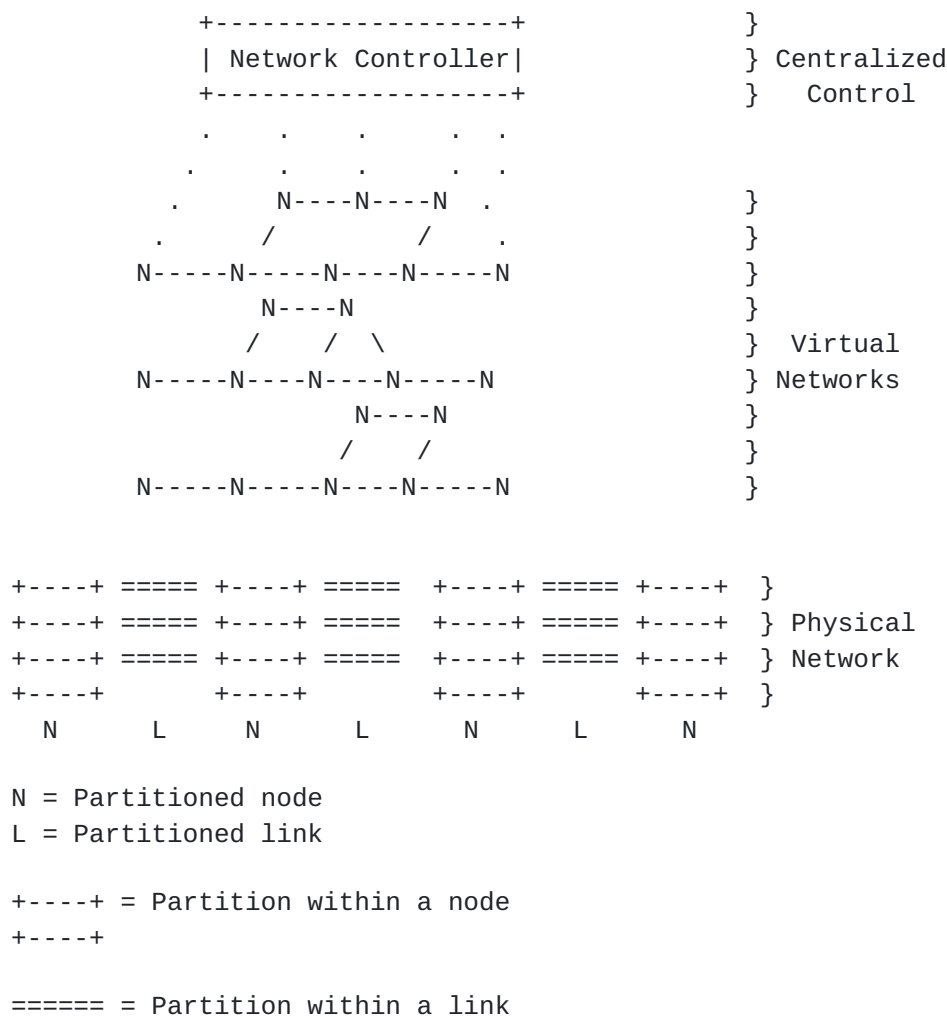
Figure 2: The Layered Architecture

Underpinning everything is the physical infrastructure layer
consisting of partitioned links and nodes which provide the
underlying resources used to provision the separated virtual
networks. Various components and techniques as discussed in
Section 4 can be used to provide the resource partition, such as
FlexE, Time Sensitive Networking, Deterministic Networking, etc.
These partitions may be physical, or virtual so long as the SLA
required by the higher layers is met.

These techniques can be used to provision the virtual networks with
dedicated resources that they need. To get the required
functionality there needs to be integration between these overlays
and the underlay providing the physical resources.

The centralized controller is used to create the virtual networks, to
allocate the resources to each virtual network and to provision the

enhanced VPN services within the virtual networks. A distributed
control plane may also be used for the distribution of the topology
and attribute information of the virtual networks.

The creation and allocation process needs to take a holistic view of
the needs of all of its tenants, and to partition the resources
accordingly. However within a virtual network these resources can if
required be managed via a dynamic control plane. This provides the
required scalability and isolation.

3.2. Multi-Point to Multi-Point

At the VPN service level, the connectivity are usually mesh or
partial-mesh. To support such kind of VPN service, the corresponding
underlay is also an abstract MP2MP medium. However when service
guarantees are provided, the point-to-point path through the underlay
of the enhanced VPN needs to be specifically engineered to meet the
required performance guarantees.

3.3. 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 two type for both PW over MPLS and for
L2TPv3.

4. Candidate Technologies

A VPN is a network created by applying a multiplexing 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 are
dedicated to different enhanced VPNs. Any enhanced VPN solution thus
needs tighter coupling with underlay than is the case with existing
VPNs. We cannot for example share the tunnel between enhanced VPNs
which require hard isolation.

4.1. Underlay Packet and Frame-Based Data Planes

A number of candidate underlay packet or frame-based data plane
solutions which can be used provide the required isolation and
guarantee are described in following sections.

o FlexE

o Time Sensitive Networking

o Dedicated Queues

4.1.1. FlexE

FlexE [FLEXE] is a method of creating a point-to-point Ethernet with
a specific fixed bandwidth. FlexE provides the ability to multiplex
multiple channels over an Ethernet link in a way that provides hard
isolation. FlexE also supports the bonding of multiple links, which
can be used to create larger links out of multiple slower links in a
more efficient way that traditional link aggregation. FlexE also
supports the sub-rating of links, which allows an operator to only
use a portion of a link. However it is a 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 in turn means
that it may be difficult to dynamically redistribute unused bandwidth
to lower priority services. This may increase the cost of providing
services on the network. On the other hand, FlexE can be used to
provide hard isolation between different tenants on a shared
interface. The tenant can then use other methods to manage the
relative priority of their own traffic in each FlexE channel.

Methods of dynamically re-sizing FlexE channels and the implication
for enhanced VPN is for further study.

4.1.2. Dedicated Queues

In order to provide multiple isolated virtual networks for enhanced
VPN, the conventional Diff-Serv based queuing system [RFC2475]
[RFC4594] is insufficient, due to the limited number of queues which
cannot differentiate between traffic of different enhanced VPNs, and
the range of service classes that each need to provide to their
tenants. This problem is particularly acute with an MPLS underlay
due to the small number of traffic class services available. In
order to address this problem and reduce the interference between
enhanced VPNs, it is necessary to steer traffic of VPNs to dedicated
input and output queues. Routers usually have large amount of queues
and sophisticated queuing systems, which could be used or enhanced to
provide the levels of isolation required by the applications of
enhanced VPN. For example, on one physical interface, the queuing
system can provide a set of virtual sub-interfaces, each allocated
with dedicated queueing and buffer resources. Sophisticated queuing
systems of this type may be used to provide end-to-end virtual
isolation between traffic of different enhanced VPNs.

4.1.3. 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
high priority packet stream may be given a scheduled time slot
thereby guaranteeing that it experiences no queuing delay and hence a
reduced latency. However, when no scheduled packet arrives, its
reserved time-slot is handed over to best effort traffic, thereby
improving the economics of the network. 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 a pseudowire.
However the TSN 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 the following section of
this document.

4.2. Packet and Frame-Based Network Layer

We now consider the problem of slice differentiation and resource
representation in the overlay network. The candidate technologies
are:

o Deterministic Networking

   o MPLS-TE

   o Segment Routing

4.2.1. Deterministic Networking

   Deterministic Networking (DetNet) [I-D.ietf-detnet-architecture] 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 Extenstion
   (SCTP-PR) [RFC3758] do 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 packets being lost, and
   trims duplicate packets to prevent excessive flooding of the network
   and to prevent multiple packets being delivered to the destination.
   It also seeks to set an upper bound on latency. The goal is not to
   minimize latency; the optimum upper bound paths may not be the
   minimum latency paths.

   DetNet is based on flows. It currently does not specify the use of
   underlay topology other than the base topology. To be of use for
   enhanced VPN, DetNet needs to be integrated with different virtual
   topologies of enhanced VPNs.

   The detailed design that allows the use DetNet in a multi-tenant
   network, and how to improve the scalability of DetNet in a multi-
   tenant network are topics for further study.

4.2.2. MPLS Traffic Engineering (MPLS-TE)

   MPLS-TE introduces the concept of reserving end-to-end bandwidth for
   a TE-LSP, which can be used as the underlay of VPNs. It also
   introduces the concept of non-shortest path routing through the use
   of the Explicit Route Object [RFC3209]. VPN traffic can be run over
   dedicated TE-LSPs to provide reserved bandwidth for each specific
   connection in a VPN. Some network operators have concerns about the
   scalability and management overhead of RSVP-TE system, and this has
   lead them to consider other solutions for their networks.

4.2.3. Segment Routing

   Segment Routing [RFC8402] is a method that prepends instructions to
   packets at the head-end node and optionally at various points as it
   passes though the network. These instructions allow the packets to
   be routed on paths other than the shortest path for various traffic

engineering reasons. These paths can be strict or loose paths,
depending on the compactness required of the instruction list and the
degree of autonomy granted to the network, for example to support
Equal Cost Multipath load-balancing (ECMP) [RFC2992].

With SR, a path needs to be dynamically created through a set of
segments by simply specifying the Segment Identifiers (SIDs), i.e.
instructions rooted at a particular point in the network. Thus if a
path is to be provisioned from some ingress point A to some egress
point B in the underlay, A is provided with a SID list from A to B
and instructions on how to identify the packets to which the SID list
is to be prepended.

By encoding the state in the packet, as is done in Segment Routing,
per-path state is transitioned out of the network.

However, there are a number of limitations in current SR, which limit
its applicability to enhanced VPNs:

o Segments are shared between different VPNs paths

o There is no reservation of bandwidth

o There is limited differentiation in the data plane.

Thus some extensions to SR are needed to provide isolation between
different enhanced VPNs. This can be achieved by including a finer
granularity of state in the network in anticipation of its future use
by authorized services. We therefore need to evaluate the balance
between this additional state and the performance delivered by the
network.

With current segment routing, the instructions are used to specify
the nodes and links to be traversed. However, in order to achieve
the required isolation between different services, new instructions
can be created which can be prepended to a packet to steer it through
specific network resources and functions.

Traditionally 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) 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. The extent to which these needs
can be satisfied through existing QoS mechanisms is to be determined.
What is clear is that a fine control of which services wait for
which, with a fine granularity of queue management policy is needed.
Note that the concept of a queue is a useful abstraction for many

types of underlay mechanism that may be used to provide enhanced
isolation and latency support.

From the perspective of the control plane, and from the perspective
of the segment routing, the method of steering a packet to a queue
that provides the required properties is an abstraction that hides
the details of the underlying implementation. How the queue
satisfies the requirement is implementation specific and is
transparent to the control plane and data plane mechanisms used.
Thus, for example, a FlexE channel, or a time sensitive networking
packet scheduling slot are abstracted to the same concept and bound
to the data plane in a common manner.

We can also introduce such fine grained packet steering by specifying
the queues through an SR instruction list. Thus new SR instructions
may be created to specify not only which resources are traversed, but
in some cases how they are traversed. For example, it may be
possible to specify not only the queue to be used but the policy to
be applied when enqueuing and dequeuing.

This concept could be further generalized, since as well as queuing
to the output port of a router, it is possible to consider queuing
data to any resource, for example:

o A network processor unit (NPU)

o A central processing unit (CPU) Core

o A Look-up engine

Both SR-MPLS and SRv6 are candidate network layer technologies for
enhanced VPN. In some cases they can be supported by DetNet to meet
the packet loss, delay and jitter requirement of particular service.
However, currently the "pure" IP variant of DetNet
[I-D.ietf-detnet-dp-sol-ip] does not support the Packet Replication,
Elimination, and Re-ordering (PREOF) [I-D.ietf-detnet-architecture]
functions. How to provide the DetNet enhanced delivery in an SRv6
environment needs further study.

4.3. Non-Packet Technologies

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 (and hence slices).

The control and management planes for non-packet data plane
technologies have most in common with MPLS-TE (Section 4.2.2) and
offer the same set of advanced features [RFC3945]. Furthermore,
management techniques such as ACTN ([RFC8453] and Section 4.4) can be
used to aid in the reporting of underlying network topologies, and
the creation of virtual networks with the resource and properties
needed by the enhanced VPN services.

4.4. 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 relies on
distributed control plane to provide scalability, high reliability,
fast reaction, automatic failure recovery etc. Extension and
optimization to the distributed control plane is needed to support
the enhanced properties of VPN+.

RSVP-TE provides the signaling mechanism of establishing a TE-LSP
with end-to-end resource reservation. It can be used to bind the VPN
to specific network resource allocated within the underlay, but there
are the above mentioned scalability concerns.

SR does not have the capability of signaling the resource reservation
along the path, nor do its currently specified distributed link state
routing protocols. On the other hand, the SR approach provides a way
of efficiently binding the network underlay and the enhanced VPN
overlay, as it reduces the amount of state to be maintained in the
network. An SR-based approach with per-slice resource reservation
can easily create dedicated SR network slices, and the VPN services
can be bound to a particular SR network slice. A centralized
controller can perform resource planning and reservation from the
controller's point of view, but this does not ensure resource
reservation is actually done in the network nodes. Thus, if a
distributed control plane is needed, either in place of an SDN
controller or as an assistant to it, the design of the control system
needs to ensure that resources are uniquely allocated in the network
nodes for the correct service, and not allocated to multiple services
causing unintended resource conflict.

4.5. Management Plane

The management plane mechanisms for enhanced VPN can be based on the
VPN service models as defined in [RFC8299] and [RFC8466], possible
augmentations and extensions to these models may be needed, which is
out of the scope of this document.

Abstraction and Control of Traffic Engineered Networks (ACTN)
[RFC8453] specifies the SDN based architecture for the control of TE
networks. The ACTN related data models such as
[I-D.ietf-teas-actn-vn-yang] and
[I-D.lee-teas-te-service-mapping-yang] can be applicable in the
provisioning of enhanced VPN service. The details are described in
Section 4.6.

4.6. Applicability of ACTN to Enhanced VPN

ACTN facilitates end-to-end connections and provides them to the
user. The ACTN framework [RFC8453] highlights how:

o Abstraction of the underlying network resources are provided to
   higher-layer applications and customers.

o Virtualization of underlying resources, whose selection criterion
   is the allocation of those resources for the customer,
   application, or service.

o Creation of a virtualized environment allowing operators to view
   and control multi-domain networks as a single virtualized network.

o The presentation to customers of networks as a virtual network via
   open and programmable interfaces.

The infrastructure managed through ACTN comprises traffic engineered
network resources, which may include:

o Statistical packet bandwidth.

o Physical forwarding plane sources, such as: wavelengths and time
   slots.

o Forwarding and cross-connect capabilities.

The type of network virtualization enabled by ACTN 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.

An ACTN Virtual Network (VN) is a client view of the ACTN managed
infrastructure, and is presented by the ACTN provider as a set of
abstracted resources.

Depending on the agreement between client and provider various VN
operations and VN views are possible.

   o Virtual Network Creation: A VN could be pre-configured and created
      via static or dynamic request and negotiation between customer and
      provider. It must meet the specified SLA attributes which satisfy
      the customer's objectives.

   o Virtual Network Operations: The virtual network may be further
      modified and deleted based on customer request to request changes
      in the network resources reserved for the customer, and used to
      construct the network slice. The customer can further act upon
      the virtual network to manage traffic flow across the virtual
      network.

   o Virtual Network View: The VN topology from a customer point of
      view. These may be a variety of tunnels, or an entire VN
      topology. Such connections may comprise of customer end points,
      access links, intra-domain paths, and inter-domain links.

   Dynamic VN Operations allow a customer to modify or delete the VN.
   The customer can further act upon the virtual network to
   create/modify/delete virtual links and nodes. These changes will
   result in subsequent tunnel management in the operator's networks.

4.6.1. ACTN Used for VPN+ Delivery

   ACTN provides VPN connections between multiple sites as requested via
   a VPN requestor enabled by the Customer Network Controller (CNC).
   The CNC is managed by the customer themselves, and interacts with the
   network provider's Multi-Domain Service Controller (MDSC). The
   Provisioning Network Controllers (PNC) remain entirely under the
   management of the network provider and are not visible to the
   customer.

   The benefits of this model include:

   o Provision of edge-to-edge VPN multi-access connectivity.

   o Management is mostly performed by the network provider, with some
      flexibility delegated to the customer-managed CNC.

```
             ----------------                      ----------------
             | Site-A Users |-----------    -----------| Site-B Users |
             ----------------          |    |          ----------------
                                    -------
                                    | CNC |
                                    -------
    Boundary                           |
    Between    ========================|========================
    Customer &                         |
    Network Operator                   |
                              ---------------
                              |    MDSC      |
                              ---------------
                          _____/      |       _____
                         /               |                \
                        /                |                 \
                   ---------         ---------         ---------
                   |  PNC  |         |  PNC  |         |  PNC  |
                   ---------         ---------         ---------
                       |                 |                /
                       |                 |               /
                     -----             -----           -----
                    (     )           (     )         (     )
         <Site A>---( Phys. )-----------( Phys. )-------( Phys. )---<Site B>
                    ( Net )            ( Net )          ( Net )
                     -----             -----            -----
```
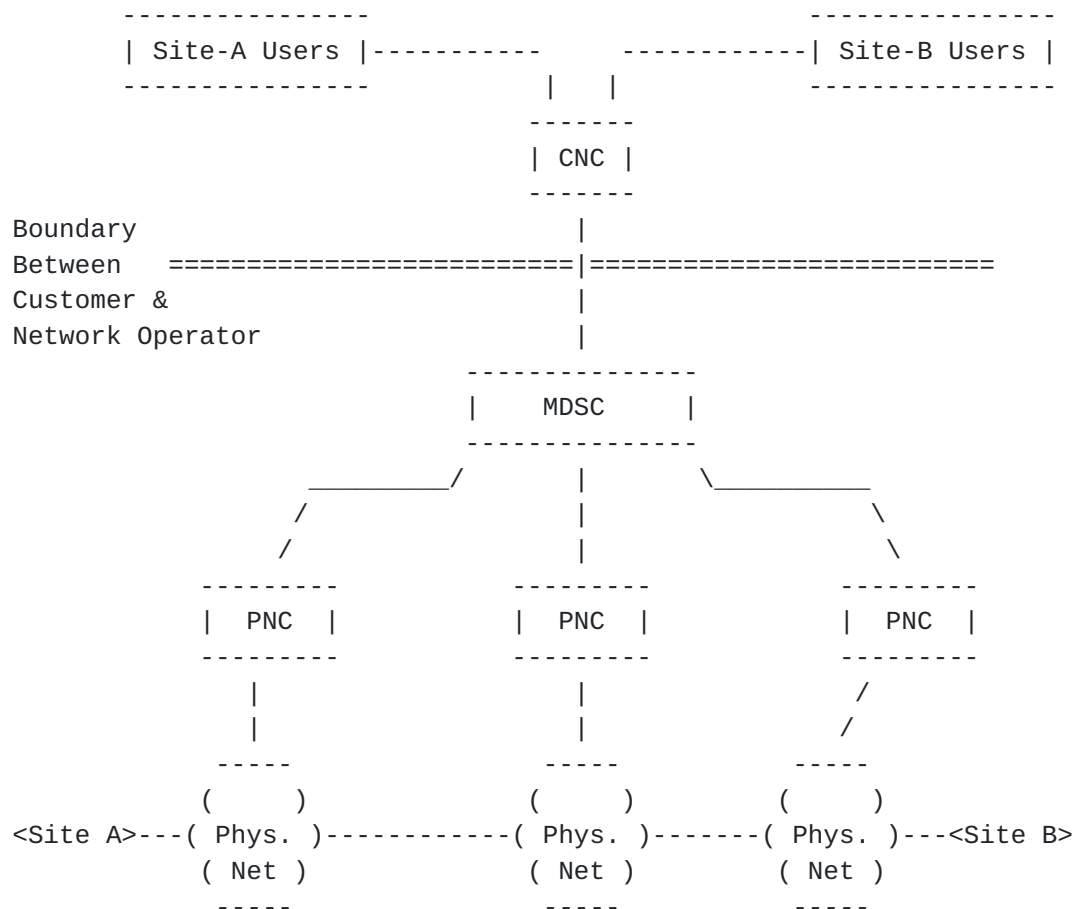
             Figure 3: VPN Delivery in the ACTN Architecture

   Figure 4 presents a more general representation of how multiple
   enhanced VPNs may be created from the resources of multiple physical
   networks using the CNC, MDSC, and PNC components of the ACTN
   architecture. Each enhanced VPN is controlled by its own CNC. The
   CNCs send requests to the provider's MDSC. The provider manages two
   physical networks each under the control of PNC. The MDSC asks the
   PNCs to allocate and provision resources to achieve the enhanced
   VPNs. In this figure, one enhanced VPN is constructed solely from
   the resources of one of the physical networks, while the the VPN uses
   resources from both physical networks.

```
                  --------------          (           )
                 |    CNC      |---------->(    VPN+    )
                  --------^------          (           )
                         |                _(_____ _)
                  --------------          (           ) ^
                 |    CNC      |---------->(    VPN+   ) :
                  ------^--------          (           ) :
                     |    |                (_____) :
                     |    |                 ^      ^    :
  Boundary           |    |                 :      :    :
  Between   =========|====|=================:====:====:========
  Customer &         |    |                 :      :    :
  Network Provider   |    |                 :      :    :
                     v    v                 :      :    :
                  --------------            :    :....:
                 |    MDSC     |            :         :
                  --------------            :         :
                     ^                    ---^------    ...
                     |                   (          )    .
                     v                   (  Physical )    .
                 --------------          ( Network  )      .
                |    PNC      |<-------->(          )    ---^------
                 -------------- |         --------       (          )
                |              |--                       (  Physical )
                |    PNC      |<-------------------------->( Network  )
                 --------------                           (          )
                                                           --------
```
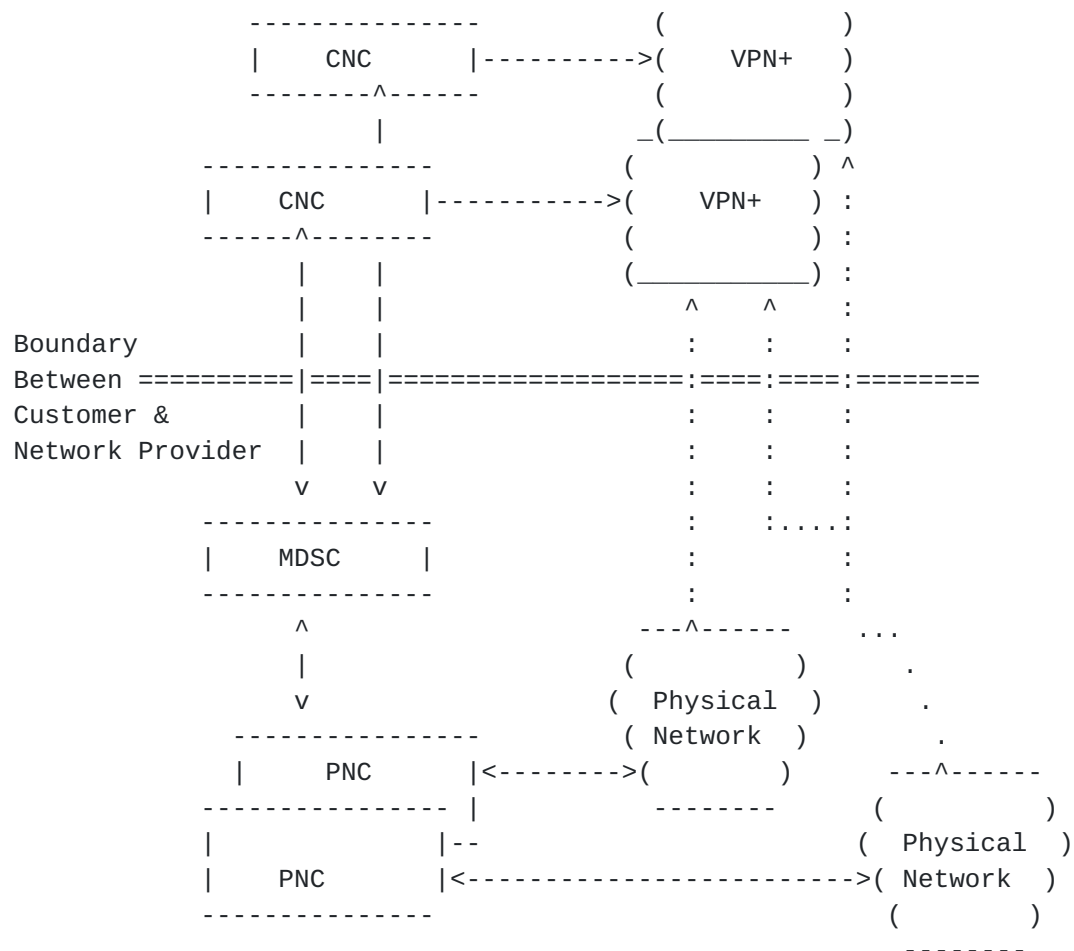
             Figure 4: Generic VPN+ Delivery in the ACTN Architecture

4.6.2. Enhanced VPN Features with ACTN

   This section discusses how the features of ACTN can fulfill the
   enhanced VPN requirements described earlier in this document. As
   previously noted, key requirements of the enhanced VPN include:

   1. Isolation between VPNs

   2. Guaranteed Performance

   3. Integration

   4. Dynamic Configuration

   5. Customized Control Plane

   The subsections that follow outline how each requirement is met using
   ACTN.

4.6.2.1. Isolation Between VPNs

   The ACTN VN YANG model [I-D.ietf-teas-actn-vn-yang] and the TE-
   service mapping model [I-D.lee-teas-te-service-mapping-yang] fulfill
   the VPN isolation requirement by providing the following features for
   the VNs:

   o Each VN is identified with a unique identifier (vn-id and vn-name)
      and so is each VN member that belongs to the VN (vn-member-id).

   o Each instantiated VN is managed and controlled independent of
      other VNs in the network with proper protection level
      (protection).

   o Each VN is instantiated with an isolation requirement described by
      the TE-service mapping model
      [I-D.lee-teas-te-service-mapping-yang]. This mapping supports:

      * Hard isolation with deterministic characteristics (e.g., this
         case may need an optical bypass tunnel or a DetNet/TSN tunnel
         to guarantee latency with no jitter)

      * Hard isolation (i.e., dedicated TE resources in all underlays)

      * Soft isolation (i.e., resource in some layer may be shared
         while in some other layers is dedicated).

      * No isolation (i.e., sharing with other VN).

4.6.2.2. Guaranteed Performance

   Performance objectives of a VN need first to be expressed in order to
   assure the performance guarantee. [I-D.ietf-teas-actn-vn-yang] and
   [I-D.ietf-teas-yang-te-topo] allow configuration of several
   parameters that may affect the VN performance objectives as follows:

   o Bandwidth

   o Objective function (e.g., min cost path, min load path, etc.)

   o Metric Types and their threshold:

      * TE cost, IGP cost, Hop count, or Unidirectional Delay (e.g.,
         can set all path delay <= threshold)

   Once these requests are instantiated, the resources are committed and
   guaranteed through the life cycle of the VN.

4.6.2.3. **Integration**

   ACTN provides mechanisms to correlate customer's VN and the actual TE
   tunnels instantiated in the provider's network. Specifically:

   o Link each VN member to actual TE tunnel.

   o Each VN can be monitored on a various level such as VN level, VN
      member level, TE-tunnel level, and link/node level.

   Service function integration with network topology (L3 and TE
   topology) is in progress in [I-D.ietf-teas-sf-aware-topo-model].
   Specifically, [I-D.ietf-teas-sf-aware-topo-model] addresses a number
   of use-cases that show how TE topology supports various service
   functions.

4.6.2.4. **Dynamic Configuration**

   ACTN provides an architecture that allows the CNC to interact with
   the MDSC which is network provider's SDN controller. This gives the
   customer control of their VNs.

   Specifically, the ACTN VN model [I-D.ietf-teas-actn-vn-yang] allows
   the VN to create, modify, and delete VNs.

4.6.2.5. **Customized Control**

   ACTN provides a YANG model that allows the CNC to control a VN as a
   "Type 2 VN" that allows the customer to provision tunnels that
   connect their endpoints over the customized VN topology.

   For some VN members, the customers are allowed to configure the path
   (i.e., the sequence of virtual nodes and virtual links) over the VN/
   abstract topology.

5. **Scalability Considerations**

   Enhanced VPN provides the performance guaranteed services in packet
   networks, with the cost of introducing necessary additional states
   into the network. There are at least three ways of adding the state
   needed for VPN+:

   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 of course 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 packet
   are shorter.

o Provide a hybrid approach based on using binding SIDs to create
   path fragments, and bind them together with SR.

Dynamic creation of a VPN path using SR requires less state
maintenance in the network core at the expense of larger VPN headers
on the packet. 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 resource 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.

These approaches are for further study.

5.1. Maximum Stack Depth of SR

One of the challenges with SR is the stack depth that nodes are able
to impose on packets [I-D.ietf-isis-segment-routing-msd]. This leads
to a difficult balance between adding state to the network and
minimizing stack depth, or minimizing state and increasing the stack
depth.

5.2. 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. There are ongoing works to improve the scalability
of RSVP-TE LSPs in the control plane [RFC8370].

There is also concern at the scalability of the forwarder footprint
of RSVP as the number of paths through an LSR grows
[I-D.sitaraman-mpls-rsvp-shared-labels] proposes to address this by
employing SR within a tunnel established by RSVP-TE.

6. OAM Considerations

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

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.

These issues will be discussed in a future version of this document.

7. Enhanced Resiliency

Each enhanced VPN has a life-cycle, and needs modification during
deployment as the needs of its tenant change. Additionally, as the
network as a whole evolves, there will 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 meet the required SLA before traffic is
transitioned to it. It is possible for deadlocks arise as a result
of the network becoming fragmented over time, such that it is
impossible to create a new path or modify a 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 the latency transient 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
[I-D.ietf-detnet-architecture].

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.

8. Security Considerations

All types of virtual network require special consideration to be
given to the isolation between the tenants. In this regard enhanced
VPNs neither introduce, no experience a greater security risk than
another VPN of the same base type. However, in an enhanced virtual
network service the isolation requirement needs to be considered. If
a service requires a specific latency then it can be damaged by

simply delaying the packet through the activities of another tenant.
In a network with virtual functions, depriving a function used by
another tenant of compute resources can be just as damaging as
delaying transmission of a packet in the network. The measures to
address these dynamic security risks must be specified as part to the
specific solution.

9. IANA Considerations

There are no requested IANA actions.

10. Contributors

   Daniel King
   Email: daniel@olddog.co.uk

   Adrian Farrel
   Email: adrian@olddog.co.uk

   Jeff Tansura
   Email: jefftant.ietf@gmail.com

   Qin Wu
   Email: bill.wu@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

11. Acknowledgements

12. References

12.1. Normative References

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

12.2. 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-detnet-architecture]
            Finn, N., Thubert, P., Varga, B., and J. Farkas,
            "Deterministic Networking Architecture", draft-ietf-
            detnet-architecture-10 (work in progress), December 2018.

[I-D.ietf-detnet-dp-sol-ip]
            Korhonen, J. and B. Varga, "DetNet IP Data Plane
            Encapsulation", draft-ietf-detnet-dp-sol-ip-01 (work in
            progress), October 2018.

[I-D.ietf-detnet-use-cases]
            Grossman, E., "Deterministic Networking Use Cases", draft-
            ietf-detnet-use-cases-20 (work in progress), December
            2018.

[I-D.ietf-isis-segment-routing-msd]
            Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg,
            "Signaling MSD (Maximum SID Depth) using IS-IS", draft-
            ietf-isis-segment-routing-msd-19 (work in progress),
            October 2018.

   [I-D.ietf-teas-actn-vn-yang]
             Lee, Y., Dhody, D., Ceccarelli, D., Bryskin, I., Yoon, B.,
             Wu, Q., and P. Park, "A Yang Data Model for ACTN VN
             Operation", draft-ietf-teas-actn-vn-yang-03 (work in
             progress), December 2018.

   [I-D.ietf-teas-sf-aware-topo-model]
             Bryskin, I., Liu, X., Lee, Y., Guichard, J., Contreras,
             L., Ceccarelli, D., and J. Tantsura, "SF Aware TE Topology
             YANG Model", draft-ietf-teas-sf-aware-topo-model-02 (work
             in progress), September 2018.

   [I-D.ietf-teas-yang-te-topo]
             Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and
             O. Dios, "YANG Data Model for Traffic Engineering (TE)
             Topologies", draft-ietf-teas-yang-te-topo-18 (work in
             progress), June 2018.

   [I-D.lee-teas-te-service-mapping-yang]
             Lee, Y., Dhody, D., Ceccarelli, D., Tantsura, J.,
             Fioccola, G., and Q. Wu, "Traffic Engineering and Service
             Mapping Yang Model", draft-lee-teas-te-service-mapping-
             yang-13 (work in progress), December 2018.

   [I-D.sitaraman-mpls-rsvp-shared-labels]
             Sitaraman, H., Beeram, V., Parikh, T., and T. Saad,
             "Signaling RSVP-TE tunnels on a shared MPLS forwarding
             plane", draft-sitaraman-mpls-rsvp-shared-labels-03 (work
             in progress), December 2017.

   [NGMN-NS-Concept]
             "NGMN NS Concept", 2016, <https://www.ngmn.org/fileadmin/u
             ser_upload/161010_NGMN_Network_Slicing_framework_v1.0.8.pd
             f>.

   [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>.

   [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>.

   [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path
             Algorithm", RFC 2992, DOI 10.17487/RFC2992, November 2000,
             <https://www.rfc-editor.org/info/rfc2992>.

   [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>.

   [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label
              Switching (GMPLS) Architecture", RFC 3945,
              DOI 10.17487/RFC3945, October 2004,
              <https://www.rfc-editor.org/info/rfc3945>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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>.

   [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
   Huawei

   Email: stewart.bryant@gmail.com

   Zhenqiang Li
   China Mobile

   Email: lizhenqiang@chinamobile.com

   Takuya Miyasaka
   KDDI Corporation

   Email: ta-miyasaka@kddi.com

   Young Lee
   Huawei

   Email: leeyoung@huawei.com