Routing Area Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

S. Bryant
J. Dong
Huawei

July 3, 2017

Enhanced Virtual Private Networks (VPN+)
draft-bryant-rtgwg-enhanced-vpn-00

Abstract

This draft describes a number of enhancements that need to be made to virtual private networks (VPNs) to support the needs of new applications, particularly applications that are associated with 5G services. A network enhanced with these properties may form the underpin 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 networks, often referred to as 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.

Driven largely by needs surfacing from 5G, the concept of network slicing has gained traction. The network slicing problem is described in [I-D.galis-netslices-revised-problem-statement] and the network slicing architecture is described in [I-D.geng-netslices-architecture]. A study of the new work needed in the IETF to address the gap between the requirements and existing IETF protocols is discussed in [I-D.qiang-netslices-gap-analysis].

Setting aside the details of the life-cycle management of a network slice instance (NSI), the underpinning technology in the transport network is a type of virtual network which provides the client with dedicated (private) networking, computing and storage resources drawn from a shared pool. The tenant of the NSI 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 of the network slice, e.g. to customize the service paths in the network slice.

These new network layer properties are of interest as part of a network slicing solution, as a precursor to the full roll-out of network slicing, and in their own right. These 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 allocated from the underlay network, and can achieve a greater isolation and lower latency than traditional VPN.

In this draft we identify the new and modified components that need to be provided in the network layer and their associated control and monitoring of an enhanced VPN. Specifically we are concerned with the technology needed to be provided by the enhanced VPN underlay, the enhanced VPN data-plane and the necessary protocols in both the underlay and the overlay of enhanced VPN. It is likely that these enhanced VPNs will be used to create network slices with different isolation requirements. Different service types, e.g. emergency services, enterprise service and broadband services etc. may be partitioned into different "hard" slices according to the isolation requirement. These "hard" slices might then be used to carry one or multiple VPNs. VPNs on such a hard slice may be only partially isolated (so called "soft" slices).

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Overview of the Requirements

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

3.1. Isolation between VPNs

The requirement is 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 a flash demand being created on another VPN competing for the

underlying resources. An example might be 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.

We introduce the terms hard and soft isolation to cover cases such as the above. A VPN has soft isolation if the traffic of one VPN cannot be inspected by the traffic of another. An 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 delays 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. A network slice that experiences only soft isolation is said to be soft sliced, and a network slice that has hard isolation is said to be hard sliced.

Although these isolation requirements are triggered by the needs of network slicing in support of 5G networks, they have general utility.

It is of course possible to achieve high degrees of isolation in the optical layer. However this is done at the cost of allocating resources on a long term basis 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 the resources. On the other hand, isolation at the packet layer allows the resources to be shared amongst many services and only dedicated to a service on a temporary basis. This allows greater statistical multiplexing of network resources and amortizes the cost over many services, leading to better economy. However, the degree of isolation required by network slicing cannot easily be met with MPLS-TE packet LSPs. Thus some trade-off between the two approaches needs to be considered to provide the required isolation between virtual networks while still allows reasonable sharing inside each VPN.

3.2. Guaranteed Performance

There are several aspects to guaranteed performance, guaranteed maximum packet loss, guaranteed maximum delay and guaranteed delay variation.

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 combine with the latency requirement.

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 is relevant, however the scope needs to be extended to methods of enhancing the IP/MPLS 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. Time transfer is one example of a service that needs this, although the fungible nature of time means that it might be delivered by the underlay and not provided through different virtual networks. The need for guaranteed maximum delay variation as a general requirement is for further study.

A possible mechanism to address these guarantees is to provide enhancement to the underlay network through technologies such as Flexible Ethernet [FLEXE].

3.3. Customized Control Plane

In some cases it is desirable that an enhanced VPN has a custom control plane, so that the tenant of the enhanced VPN can have some control to the resources and functions partitioned for this VPN.

Further detail on this requirement will be provided in a future version of the draft.

4. Components of VPN+

4.1. Use of Segment Routing Constructs

Clearly we can use traditional constructs to create a VPN, but there are advantages to the use of Segment Routing (SR) in the creation of virtual networks with enhanced properties.

Segment Routing [I-D.ietf-spring-segment-routing] is a method that prepends instructions to packets at entry and possibly at various points as it passes though the network. These instructions allow 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 ECMP). 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 dedicated network resources and functions, e.g. queues, processors, links, services etc. New instructions can also 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.

With SR, a path is dynamically created through a set of resources by simply specifying the Segment IDs (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 the A..B SID list and instructions on how to identify the packets to which the SID list is to be prepended.

The SIDs may be used to specify both network paths, or service functions as described in [I-D.xu-mpls-unified-source-routing-instruction].

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 scaling properties will reduce roughly from a function of (N/2)^2 to a function of N, where N is the VPN path length in intervention points (hops plus network functions). Reducing the state in the network is important to VPN+, as VPN+ requires the overlay to be more closely integrated with the underlay than with traditional VPNs. This tighter coupling would normally mean that significant state needed to be created and maintained in the core. 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.

4.2. Latency Support

The IETF has ongoing work on support for a latency ceiling [I-D.ietf-detnet-architecture]. The provision of a latency ceiling is a requirement of the application seeking the use of enhanced virtual networks. The current design of DetNet assumes the design of the underlay network is unchanged. In this section we look at some changes that could be used to assist in achieving low latency ceiling across the wide area.

Traditionally a traffic engineered path operates with a granularity of a link with hints about priority provided through the use of the traffic class field in the header. However to achieve the latency and isolation characteristics that are sought, steering packets through specific queues may be required. This allows a much finer

control of which services wait for which, and a much finer granularity of queue management policy.

This may be introduced into traditional path construction techniques such as RSVP-TE and MPLS-TP, or it may be introduced by specifying the queue in an SR instruction list.

4.3. Support of an IP underlay

Where an underlay needs to be provided by IP, a number of options present themselves. We could allocate an IP address to that path and construct a path through the network for that IP address. The path could be laid in with RSVP signaling or through SDN controller. This path could have all of the required properties including specifying resources to use and functions to visit. Although this construct has been considered many time over the years, such a mechanism, at least to the author's knowledge, has not found favor in deployment.

There are two ways that segment routing might be used to adapt the system described above to an IP context. One is to use the method described in [I-D.ietf-6man-segment-routing-header] in which each segment (instruction) is encoded as a normal IPv6 address. An alternative is to use the more compact representation considered in [I-D.xu-mpls-unified-source-routing-instruction] and [I-D.bryant-mpls-unified-ip-sr].

4.4. Application Specific Network Types

Although the transport service that underpins the extended VPN is likely MPLS/IP based, it needs to be able to carry application specific non-MPLS/IP traffic. This can be accommodated through the use of pseudowires (PWs).

4.5. A Hybrid Control Plane

It is expected that VPN+ would 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+. [I-D.king-teas-applicability-actn-slicing] describes the use of ACTN to network slicing. This approach may be considered as part of the centralized control plane of VPN+ in some applications.

5. Applicability to Network Slicing

In [I-D.geng-netslices-architecture] a network slice is defined to be:

"A managed group of subsets of resources, network functions / network virtual functions at the data, control, management/orchestration planes and services at a given time. Network slice is programmable and has the ability to expose its capabilities. The behaviour of the network slice realized via network slice instance(s)."

A network slice instance (NSI) is then defined as:

"An activated network slice. It is created based on network template.
A set of managed run-time network functions, and resources to run these network functions, forming a complete instantiated logical network to meet certain network characteristics required by the service instance(s).
It provides the network characteristics that are required by a service instance. A network slice instance may also be shared across multiple service instances provided by the network operator. The network slice instance may be composed by none, one or more sub-network instances, which may be shared by another network slice instance."

A network slice can thus be thought of as a customized set of logical network and compute resources required by the service the slice is supporting. These resources support both virtual services in the data path and operation of the slice, for example by providing routing services. The customization includes the connectivity, performance, and isolation characteristics. These characteristics can be provided by the enhanced VPN described in this draft.

6. Security Considerations

All types of virtual network require special consideration to be given to the isolation between the tenants. However in an enhanced virtual network service hard isolation 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.

7. IANA Considerations

There are no requested IANA actions.

8. References

8.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

8.2. Informative References

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

[I-D.bryant-mpls-unified-ip-sr]
Bryant, S., Xu, X., Chen, M., Farrel, A., and J. Drake, "A Unified Approach to IP Segment Routing", draft-bryant-mpls-unified-ip-sr-00 (work in progress), June 2017.

[I-D.galis-netslices-revised-problem-statement]
Galis, A., "Network Slicing - Revised Problem Statement", draft-galis-netslices-revised-problem-statement-00 (work in progress), June 2017.

[I-D.geng-netslices-architecture]
67, 4., Dong, J., Bryant, S., kiran.makhijani@huawei.com, k., Galis, A., Foy, X., and S. Kuklinski, "Network Slicing Architecture", draft-geng-netslices-architecture-01 (work in progress), June 2017.

[I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Raza, K., Leddy, J., Field, B., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Matsushima, S., Leung, I., Linkova, J., Aries, E., Kosugi, T., Vyncke, E., Lebrun, D., Steinberg, D., and R. Raszuk, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-06 (work in progress), March 2017.

[I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", draft-ietf-detnet-architecture-02 (work in progress), June 2017.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-12 (work in progress), June 2017.

[I-D.king-teas-applicability-actn-slicing]
King, D., "Applicability of Abstraction and Control of TE Networks (ACTN) to Network Slicing", draft-king-teas-applicability-actn-slicing-00 (work in progress), June 2017.

[I-D.qiang-netslices-gap-analysis]
Qiang, L., Martinez-Julia, P., 67, 4., Dong, J., kiran.makhijani@huawei.com, k., Galis, A., Hares, S., and S. Slawomir, "Gap Analysis for Network Slicing", draft-qiang-netslices-gap-analysis-00 (work in progress), June 2017.

[I-D.xu-mpls-unified-source-routing-instruction]
Xu, X., Bryant, S., Raszuk, R., Chunduri, U., Contreras, L., Jalil, L., Assarpour, H., Velde, G., Tantsura, J., and S. Ma, "Unified Source Routing Instruction using MPLS Label Stack", draft-xu-mpls-unified-source-routing-instruction-02 (work in progress), June 2017.

Authors' Addresses

Stewart Bryant
Huawei

Email: stewart.bryant@gmail.com

Jie Dong
Huawei

Email: jie.dong@huawei.com