Network Working Group
Internet Draft
Intended status: Informational
Expires: January 2018

L. Dunbar
A. Malis
Huawei
C. Jacquenet
Orange

October 30, 2017

VPN Extension to Dynamic Cloud DC Problem Statement
draft-dm-vpn-ext-to-cloud-dc-problem-statement-01

Abstract

This document describes the problems associated with extending existing VPN that interconnects Enterprise customers' multiple sites to dynamic workloads instantiated in cloud data centers. This document further describes a set of requirements that a solution would need to fulfill to address the problems discussed herein.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on April 30, 2009.

Table of Contents

1. Introduction

Cloud-based applications and services continue to change how businesses of all sizes work and share information. "Cloud based applications & workloads" are those that are instantiated in third party Data Centers which also host services for other customers. The benefits of these cloud-based applications and services are numerous, including fueling mobility and access to applications anytime, anywhere, and on any device, making collaboration more efficient and easier to manage.

With the advent of widely available third party cloud data centers in diverse geographic locations and the advancement of tools for monitoring and predicting application behaviors, it is technically feasible for enterprises to instantiate applications and workloads geographically closest to their end users. This property aids in improving end-to-end latency and overall user experience. Conversely, an enterprise may wish to shutdown applications and workloads that are too far from their end users (therefore removing the networking connection to those deleted applications and workloads). In addition, an Enterprise may wish to take advantage of more and more business applications offered by third party private cloud data centers, such as SAP HANA, Oracle Cloud, Salesforce Cloud, etc.

However, given the nature of how most Enterprise VPN networks are built, whether SD-WAN, MPLS-based, or a combination of both, it is difficult (or impossible) for many Enterprises to utilize these cloud-based resources in a flexible and scalable manner with reasons to be elaborated in subsequent sections of this documents.

This document describes a number of issues with existing VPN technologies, either SD-WAN or MPLS-based, related to connectivity of Enterprise sites to dynamic workloads instantiated in a cloud data center. The Enterprise Sites include HQ, spokes, on premise data centers, and branch offices as their corresponding VPN features can be different.

2. Definition of terms

Cloud DC: Off-Premise Data Centers that usually host applications and workload owned by different organizations or tenants.

Controller: Used interchangeably with SD-WAN controller to manage SD-WAN overlay path creation/deletion and monitoring the path conditions between two sites.

SD-WAN: Software Defined Wide Area Network, which can mean many different things. In this document, "SD-WAN" refers to the solutions specified by ONUG (Open Network User Group), which build point-to-point IPsec overlay paths between two end-points (or branch offices) that need to intercommunicate.

3. Problems associated with current SD-WAN solutions

A software-defined wide area network (SD-WAN) VPN is an overlay network that decouples the network management function from the physical hardware, using a centralized controller to set policies, prioritize network traffic, establish IPsec [RFC6071] tunnels between enterprise locations to carry the VPN traffic, and to map between internal addresses on the VPN and external addresses on the public Internet. Many enterprises use SD-WAN VPNs as an alternative, or in addition to, more traditional VPNs (such as MPLS-based VPNs [RFC4364] or [RFC4664]).

SD-WAN is typically used to control traffic distribution among multiple paths between two end-points, e.g. some paths being MPLS path, others being via public internet. SD-WAN depends on logically centralized network control to utilize real-time traffic management over multiple paths between the two end-points. The virtual overlay,

possibly secured by IPsec tunnels, is transported over the public Internet using fiber, cable, or DSL-based Internet access, but can use other types of WAN connections as well, including private or public WAN connections like MPLS, or wireless technologies such as Wi-Fi or 4G/Long Term Evolution (LTE).

SD-WAN lets enterprises augment their current network with cost-effective, readily available Broadband Internet connectivity. When used in conjunction with MPLS VPNs, some traffic can be offloaded to SD-WAN overlay paths based on traffic forwarding policy (application-based or otherwise), or when the MPLS VPN connection between the two locations is congested, or otherwise undesirable.

3.1. Complexity of multi-point any-to-any interconnection

The dynamic workload instantiated in cloud DC needs to communicate with multiple branch offices and on premise data centers. Most enterprises need multi-point interconnection among multiple locations, as done by MPLS L2/L3 VPNs.

Using SD-WAN overlay paths to achieve any-to-any mesh interconnection among all branches not only requires all branches CPEs to support SD-WAN, but also require CPEs to manage routing among other CPEs located at other locations, which can increase the complexity of the CPEs when compared to MPLS-based VPN solutions.

Today's industry so called "SD-WAN" solutions build point-to-point IPsec overlay paths between two end-points (or branch offices) that need to intercommunicate. This overlay path can serve as a backup to an MPLS path in a hybrid solution, or as the primary path in a stand-alone solution.

Whereas, MPLS-based VPNs have their PEs directly connected to the CEs. Therefore, CEs only need to forward all traffic with destinations attached to remote CEs to the directly attached PEs, leaving all the routing to remote destinations to the PEs, thereby reducing the complexity of CPEs. Even for multi-home CPEs, the CPEs only need to route traffic among the PEs that the CPEs are directly attached to. But the SD-WAN's CPE needs to route the traffic among all other CPEs.

For an enterprise with multiple sites, using SD-WAN overlay paths among sites requires each CPE to manage all the addresses that local hosts have the potential to reach, i.e. map internal VPN addresses

to appropriate SD-WAN paths. This is similar to the complexity of Frame Relay-based VPNs, where each CPE needed to maintain mesh routing for all destinations if they were to avoid an extra hop through a hub router. That is one of the primary reasons most enterprise networks transitioned to MPLS-based VPNs from Frame Relay. Even though SD-WAN CPEs can get assistance from a central controller (instead of performing their own routing protocol) to resolve the mapping between destinations and SD-WAN paths, SD-WAN CPEs are still responsible for routing table maintenance as remote destinations change their attachments, e.g. the dynamic workload in other DCs are de-commissioned or added.

3.2. Poor performance over long distance

When CPEs are far apart from each other or across particular boundaries, whether political (e.g. country boundary) or related to Internet topology, performance over the public Internet can be problematic and unpredictable. Even though there are many monitoring tools available to measure delay and various performance characteristics of the network, the measurement for paths over the Internet is passive and past measurements may not represent future performance.

To compensate for delay over the Internet, most content today is hosted by data centers closest to end users. E2E services usually do not traverse long distances, but rather between end users and local data centers. Content distribution to the edges has transformed user experience of accessing content over the Internet.

However, SD-WAN is about connecting two geographically different locations, which is very different from today's experience of accessing various websites over the Internet.

3.3. Scaling Issues with IPsec Tunnels

IPsec is used by SD-WAN to achieve secure overlay connections between two locations over any underlay network.
For a simple SD-WAN overlay between a small number of fixed branch offices, each CPE only needs to terminate a very small number of IPsec tunnels, which will be for the most part static in nature. However, for multiple branches to reach workloads hosted in cloud DCs, the SD-WAN solution requires a Cloud DC gateway to maintain individual IPsec tunnels between the Cloud DC gateway and each individual branch office. For a company with hundreds or thousands of locations, there could be hundreds (or even thousands) of IPsec tunnels terminating at the Cloud DC gateway, which can be very processing intensive for the gateway. Many routers today have limited capacity to support a large number of IPsec tunnels.

3.4. End-to-End Security Concern for data flows

When IPsec tunnels from enterprise on-premise CPEs are terminated at the Cloud DC gateway where the workloads or applications are hosted, some enterprises have concerns regarding traffic to/from their workload being exposed to others behind the data center gateway (e.g., exposed to other organizations that have workloads in the same data center).
To ensure that traffic to/from workloads is not exposed to unwanted entities, it's necessary to have the IPsec tunnels go all the way to the workload (say servers, or VMs) within the DC.

4. Problems associated with MPLS-based VPNs for dynamic applications in the cloud

Traditional VPNs (e.g., MPLS based L2/L3 VPNs) that most businesses use to connect their branch offices are isolated, secure and reliable, but they cannot keep up in the cloud-based world. One of the key roadblocks for achieving this dynamic workload instantiation is the lack of flexible & secure network connectivity to workloads in third party cloud data centers. Another roadblock is the lack of a standard way to express and enforce consistent security policies to their workload [RFC8192]. The traditional VPN path and bandwidth

are not flexible enough in supporting the need for enterprises to connect to dynamically instantiated (or removed) workloads and applications at any place (i.e., third party cloud data centers).

The current business environment is characterized by dramatic and constant changes, especially around the IT space. Those changes include:

- The movement on the part of most businesses to adopt digital business initiatives, such as variety of data & behavior analytic tools for end customers, employees, and products. Those tools need to be running close to their targets to achieve optimal results.
- Broad adoption of public cloud computing services.
- Deployment of WAN solutions based on new architectural approaches.
- The dramatic increase in the number and sophistication of security attacks.

Traditional MPLS-based VPNs have been widely deployed as an effective way to support business and organizations that require network performance and reliability. MPLS shifted the burden of managing a VPN service from enterprises to service providers. The CPEs for MPLS VPN are also simpler and less expensive, since they do not need to manage how to send packets to remote sites, they simply pass all outbound traffic to the MPLS VPN PEs to which the CPEs are attached. MPLS has addressed the problems of scale, availability, and fast recovery from network faults, and incorporates traffic engineering to ensure guaranteed bandwidth for high priority traffic.

However, traditional MPLS-based VPNs are not optimized for connecting to dynamic applications in cloud data centers because:

- It's not easy to add/remove VPN's PEs at dynamic locations. It takes a relatively long time to deploy provider edge routers at new locations. When enterprise's workloads are changed from one cloud DC to another (i.e., removed from one DC and re-instantiated to another location when demand changes), the enterprise branch offices need to be connected to the new cloud DC.

The big drive for moving workloads into the cloud comes from widely available cloud DCs at geographically diverse locations that apps can be instantiated so that they can be close to their users. When the user base changes, the applications can be quickly moved to a new cloud DC location closest to the new user base.

- The third party cloud data center where an enterprise chooses to host workloads for easy access to its clients may not be connected to the Provider Edge (PE) nodes of the enterprise's VPN.
- Many cloud data centers do not expose its internal network for provider MPLS based VPNs to reach the workload natively & securely.
- Many data centers use some forms of encapsulation, i.e. VxLAN, STT, NSH, etc. There has not been any standard to address the interworking to those encapsulations.
-

5. Requirements for Dynamic Cloud Data Center VPNs

In order to address the aforementioned issues, any solution for enterprise VPNs that includes connectivity to dynamic workloads or applications in cloud data centers should satisfy a set of requirements:

- The solution should allow enterprises to take advantage of the current state of the art in VPN technology, both in traditional MPLS-based VPNs and IPsec-based VPNs (or any combination thereof) that run over the top of the public Internet.
- The solution shouldn't require enterprise to upgrade all their existing CPEs. .
- The solution shouldn't require either CPEs or routers to support more than xx IPsec tunnels simultaneously.
- The solution needs to support easy and fast VPN connections to dynamic workloads and applications in third party data centers, and easily allow these workloads to migrate both within a data center and between data centers.

- Allow VPNs to provide bandwidth and other performance guarantees.
- Be a cost-effective solution for enterprises to incorporate dynamic cloud-based applications and workloads into their existing VPN environment.

6. Security Considerations

For the most part, we introduce no new security concerns beyond those of existing MPLS=based VPNs, which are extremely widely deployed. The one addition to MPLS VPNs is selective use of SD-WAN, which uses IPsec tunnels.

7. IANA Considerations

This document requires no IANA actions. RFC Editor: Please remove this section before publication.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[RFC8192] S. Hares, et al "Interface to Network Security Functions (I2NSF) Problem Statement and Use Cases", July 2017

[ITU-T-X1036] ITU-T Recommendation X.1036, "Framework for creation, storage, distribution and enforcement of policies for network security", Nov 2007.

[RFC6071] S. Frankel and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", Feb 2011.

[RFC4364] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", Feb 2006

[RFC4664] L. Andersson and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", Sept 2006.

9. Acknowledgments

Acknowledgements to Jim Guichard for his review and contributions.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Linda Dunbar
Huawei
Email: Linda.Dunbar@huawei.com

Andrew G. Malis
Huawei
Email: agmalis@gmail.com

Christian Jacquenet
France Telecom
Rennes, 35000
France
Email: Christian.jacquenet@orange.com