

Network Working Group
Internet Draft
Intended status: Informational
Expires: November 1, 2020

L. Dunbar
Futurewei
A. Malis
Independent

C. Jacquenet
Orange
May 1, 2020

Networks Connecting to Hybrid Cloud DCs: Gap Analysis
draft-ietf-rtgwg-net2cloud-gap-analysis-06

Abstract

This document analyzes the technical gaps that may affect the dynamic connection to workloads and applications hosted in hybrid Cloud Data Centers from enterprise premises.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), 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 November 1, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction.....	3
2.	Conventions used in this document.....	3
3.	Gap Analysis for Accessing Cloud Resources.....	4
3.1.	Multiple PEs connecting to virtual CPEs in Cloud DCs.....	7
3.2.	Access Control for workloads in the Cloud DCs.....	7
3.3.	NAT Traversal.....	8
3.4.	BGP between PEs and remote CPEs via Internet.....	8
3.5.	Multicast traffic from/to the remote edges.....	9
4.	Gap Analysis of Overlay Edge Node's WAN Port Management.....	9
5.	Aggregating VPN paths and Internet paths.....	11
5.1.	Control Plane for Overlay over Heterogeneous Networks....	12
5.2.	Using BGP UPDATE Messages.....	13
5.2.1.	Lacking SD-WAN Segments Identifier.....	13
5.2.2.	Missing attributes in Tunnel-Encap.....	13
5.3.	SECURE-L3VPN/EVPN.....	15
5.4.	Preventing attacks from Internet-facing ports.....	16
6.	Gap Summary.....	16
7.	Manageability Considerations.....	17
8.	Security Considerations.....	17
9.	IANA Considerations.....	18
10.	References.....	18
10.1.	Normative References.....	18
10.2.	Informative References.....	18
11.	Acknowledgments.....	19

1. Introduction

[Net2Cloud-Problem] describes the problems enterprises face today when interconnecting their branch offices with dynamic workloads hosted in third party data centers (a.k.a. Cloud DCs). In particular, this document analyzes the routing protocols to identify whether there are any gaps that may impede such interconnection which may for example justify additional specification effort to define proper protocol extensions.

For the sake of readability, an edge, an endpoint, C-PE, or CPE are used interchangeably throughout this document. More precisely:

- . Edge: may include multiple devices (virtual or physical);
- . endpoint: refers to a WAN port of device located in the edge;
- . C-PE: provider-owned edge, e.g. for SECURE-EVPN's PE-based BGP/MPLS VPN, where PE is the edge node;
- . CPE: device located in enterprise premises.

2. Conventions used in this document

Cloud DC: Third party Data Centers that usually host applications and workload owned by different organizations or tenants.

Controller: Used interchangeably with Overlay controller to manage overlay path creation/deletion and monitor the path conditions between sites.

CPE-Based VPN: Virtual Private Network designed and deployed from CPEs. This is to differentiate from most commonly used PE-based VPNs a la [RFC 4364](#).

OnPrem: On Premises data centers and branch offices

SDWAN: Software Defined Wide Area Network, "SDWAN" refers to the solutions of pooling WAN bandwidth from multiple underlay networks to get better WAN bandwidth

management, visibility & control. When the underlay is a private network, traffic may be forwarded without any additional encryption; when the underlay networks are public, such as the Internet, some traffic needs to be encrypted when passing through (depending on user-provided policies).

3. Gap Analysis for Accessing Cloud Resources

Because of the ephemeral property of the selected Cloud DCs for specific workloads/Apps, an enterprise or its network service provider may not have direct physical connections to the Cloud DCs that are optimal for hosting the enterprise's specific workloads/Apps. Under those circumstances, an overlay network design can be an option to interconnect the enterprise's on-premises data centers & branch offices to its desired Cloud DCs.

However, overlay paths established over the public Internet can have unpredictable performance, especially over long distances and across operators' domains. Therefore, it is highly desirable to minimize the distance or the number of segments that traffic had to be forwarded over the public Internet.

The Metro Ethernet Forum's Cloud Service Architecture [MEF-Cloud] also describes a use case of network operators using Overlay paths over an LTE network or the public Internet for the last mile access where the VPN service providers cannot always provide the required physical infrastructure.

In these scenarios, some overlay edge nodes may not be directly attached to the PEs that participate to the delivery and the operation of the enterprise's VPN.

When using an overlay network to connect the enterprise's sites to the workloads hosted in Cloud DCs, the existing C-PEs at enterprise's sites have to be upgraded to connect to the said overlay network. If the workloads hosted in Cloud DCs need to be connected to many sites, the upgrade process can be very expensive.

[Net2Cloud-Problem] describes a hybrid network approach that extends the existing MPLS-based VPNs to the Cloud DC Workloads over the access paths that are not under the VPN provider's control. To make it work properly, a small number of the PEs of the BGP/MPLS VPN can be designated to connect to the remote workloads via secure IPsec tunnels. Those designated PEs are shown as fPE (floating PE or smart PE) in Figure 3. Once the secure IPsec tunnels are established, the workloads hosted in Cloud DCs can be reached by the enterprise's VPN without upgrading all of the enterprise's CPEs. The only CPE that needs to connect to the overlay network would be a virtualized CPE instantiated within the cloud DC.

[3.1.](#) Multiple PEs connecting to virtual CPEs in Cloud DCs

To extend BGP/MPLS VPNs to virtual CPEs in Cloud DCs, it is necessary to establish secure tunnels (such as IPsec tunnels) between the PEs and the vCPEs.

Even though a set of PEs can be manually selected for a specific cloud data center, there are no standard protocols for those PEs to interact with the vCPEs instantiated in the third party cloud data centers over unsecure networks. The interaction includes exchanging performance, route information, etc..

When there is more than one PE available for use (as there should be for resiliency purposes or because of the need to support multiple cloud DCs geographically scattered), it is not straightforward to designate an egress PE to remote vCPEs based on applications. It might not be possible for PEs to recognize all applications because too much traffic traversing the PEs.

When there are multiple floating PEs that have established IPsec tunnels with a remote CPE, the remote CPE can forward outbound traffic to the optimal PE, which in turn forwards traffic to egress PEs to reach the final destinations. However, it is not straightforward for the ingress PE to select which egress PEs to send traffic. For example, in Figure 1:

- fPE-1 is the optimal PE for communication between App-1 <-> Host-a due to latency, pricing or other criteria.
- fPE-2 is the optimal PE for communication between App-1 <-> Host-b.

[3.2.](#) Access Control for workloads in the Cloud DCs

There is widespread diffusion of access policy for Cloud Resource, some of which is not easy for verification and validation. Because there are multiple parties involved in accessing Cloud Resources, policy enforcement points are not easily visible for policy refinement, monitoring, and testing.

The current state of the art for specifying access policies for Cloud Resources could be improved by having automated and reliable tools to map the user-friendly (natural language) rules into machine readable policies and to provide interfaces for enterprises to self-manage policy enforcement points for their own workloads.

[3.3. NAT Traversal](#)

Cloud DCs that only assign private IPv4 addresses to the instantiated workloads assume that traffic to/from the workload usually needs to traverse NATs.

An overlay edge node can solicit a STUN (Session Traversal of UDP Through Network Address Translation, [[RFC3489](#)]) Server to get the information about the NAT property, the public IP addresses and port numbers so that such information can be communicated to the relevant peers.

[3.4. BGP between PEs and remote CPEs via Internet](#)

Even though an EBGP (external BGP) Multi-Hop design can be used to connect peers that are not directly connected to each other, there are still some issues about extending BGP from MPLS VPN PEs to remote CPEs via any access path (e.g., Internet).

The path between the remote CPEs and VPN PEs that maintain VPN routes can traverse untrusted segments.

EBGP Multi-hop design requires configuration on both peers, either manually or via NETCONF from a controller. To use EBGP between a PE and remote CPEs, the PE has to be manually configured with the "next-hop" set to the IP address of the CPEs. When remote CPEs, especially remote virtualized CPEs are dynamically instantiated or removed, the configuration of Multi-Hop EBGP on the PE has to be changed accordingly.

Egress peering engineering (EPE) is not sufficient. Running BGP on virtualized CPEs in Cloud DCs requires GRE tunnels to be established first, which requires the remote CPEs to support address and key management capabilities. [RFC 7024](#) (Virtual Hub & Spoke) and Hierarchical VPN do not support the required properties.

Also, there is a need for a mechanism to automatically trigger configuration changes on PEs when remote CPEs' are instantiated or moved (leading to an IP address change) or deleted.

EBGP Multi-hop design does not include a security mechanism by default. The PE and remote CPEs need secure communication channels when connecting via the public Internet.

Remote CPEs, if instantiated in Cloud DCs might have to traverse NATs to reach PEs. It is not clear how BGP can be used between devices located beyond the NAT and the devices located behind the NAT. It is not clear how to configure the Next Hop on the PEs to reach private IPv4 addresses.

3.5. Multicast traffic from/to the remote edges

Among the multiple floating PEs that are reachable from a remote CPE, multicast traffic sent by the remote CPE towards the MPLS VPN can be forwarded back to the remote CPE due to the PE receiving the multicast packets forwarding the multicast/broadcast frame to other PEs that in turn send to all attached CPEs. This process may cause traffic loops.

This problem can be solved by selecting one floating PE as the CPE's Designated Forwarder, similar to TRILL's Appointed Forwarders [[RFC6325](#)].

BGP/MPLS VPNs do not have features like TRILL's Appointed Forwarders.

4. Gap Analysis of Overlay Edge Node's WAN Port Management

Very often the Hybrid Cloud DCs are interconnected by overlay networks that arch over many different types of networks, such as VPN, public Internet, wireless and wired infrastructures, etc. Sometimes the enterprises' VPN providers do not have direct access to the Cloud DCs that host some specific applications or workloads

operated by the enterprise.

Under those circumstances, the overlay network's edge nodes can have WAN ports facing networks provided by different ISPs, some of these networks may not be trustable, some others can be trusted like VPNs (to some extent), etc.

If all WAN ports of an edge node are facing an untrusted network, then all sensitive data to/from this edge node have to be encrypted, usually by means of IPsec tunnels which can be terminated at the WAN port address, at the edge node's loopback address if the loopback address is routable in the wide area network, or even at the ingress ports of the edge node.

If an edge node has some WAN ports facing trusted networks and others facing untrusted networks, sensitive data can be forwarded through ports facing the trusted networks natively (i.e., without encryption) and forwarded through ports facing untrusted networks assuming encryption. To achieve this flexibility of sending traffic either encrypted or not encrypted depending on egress WAN ports, it is necessary to have the IPsec tunnels terminated at the WAN ports facing the untrusted networks.

In order to establish peer-wise secure encrypted communications among those WAN ports of two edge nodes, it is necessary for the edge nodes (peers) to be informed of the WAN port properties.

Some of those overlay networks (such as some deployed SDWAN networks) use the modified NHRP protocol [[RFC2332](#)] to register WAN ports of the edge nodes with their Controller (or NHRP server), which then maps a private VPN address to a public IP address of the destination node/port. DSVPN [[DSVPN](#)] or DMVPN [[DMVPN](#)] are used to establish tunnels between WAN ports of SDWAN edge nodes.

NHRP was originally intended for ATM address resolution, and as a result, it misses many attributes that are necessary for dynamic endpoint C-PE registration to the controller, such as:

- Interworking with the MPLS VPN control plane. An overlay edge can have some ports facing the MPLS VPN network over which packets can be forwarded without any encryption and some ports facing the public Internet over which sensitive traffic needs to be

encrypted.

- Scalability: NHRP/DSVPN/DMVPN work fine with small numbers of edge nodes. When a network has more than 100 nodes, these protocols do not scale well.
- NHRP does not have the IPsec attributes, which are needed for peers to build Security Associations over the public Internet.
- NHRP messages do not have any field to encode the C-PE supported encapsulation types, such as IPsec-GRE or IPsec-VxLAN.
- NHRP messages do not have any field to encode C-PE Location identifiers, such as Site Identifier, System ID, and/or Port ID.
- NHRP messages do not have any field to describe the gateway(s) to which the C-PE is attached. When a C-PE is instantiated in a Cloud DC, it is desirable for the C-PE's owner to be informed about how and where the C-PE is attached.
- NHRP messages do not have any field to describe C-PE's NAT properties if the C-PE is using private IPv4 addresses, such as the NAT type, Private address, Public address, Private port, Public port, etc.

[BGP-SDWAN-PORT] describes how to use BGP to distribute SDWAN edge properties to peers. SDWAN is an overlay network with specific properties, such as application-based forwarding, augmented transport, and user specified policies. There is a need to extend the protocol to register WAN port properties of an edge node to the overlay controller, which then propagates the information to other overlay edge nodes that are authenticated and authorized to communicate with them.

[5. Aggregating VPN paths and Internet paths](#)

Most likely, enterprises (especially the largest ones) already have their C-PEs interconnected by VPN service providers, based upon VPN techniques like EVPN, L2VPN, or L3VPN, and which can be lead to PE-based or CPE-based VPN service designs. The commonly used PE-based BGP/MPLS VPNs have C-PEs directly attached to PEs, the communication between C-PEs and PEs is considered as secure as they are connected by direct physical links albeit there could be routes leaking or unauthorized routes being injected. MP-BGP can be used to learn & distribute routes among C-PEs, but sometimes routes among C-PEs are statically configured on the C-PEs.

For enterprises already interconnected by VPNs, if there are short term high traffic volume that can't justify increasing the VPNs capacity, it is desirable for the CPE to aggregate the bandwidth that pertains to VPN paths and Internet paths by adding ports that connect the CPE to the public Internet. Under this scenario, which is referred to as the Overlay scenario throughout this document, it is necessary for the C-PEs to manage and communicate with the controller on how traffic is distributed among multiple heterogeneous underlay networks, and also to manage secure tunnels over untrusted networks.

When using NHRP for WAN port registration purposes, C-PEs need to run two separate control planes: EVPN&BGP for CPE-based VPNs, and NHRP & DSVPN/DMVPN for ports connected to the Internet. Two separate control planes not only add complexity to C-PEs, but also increase operational costs.

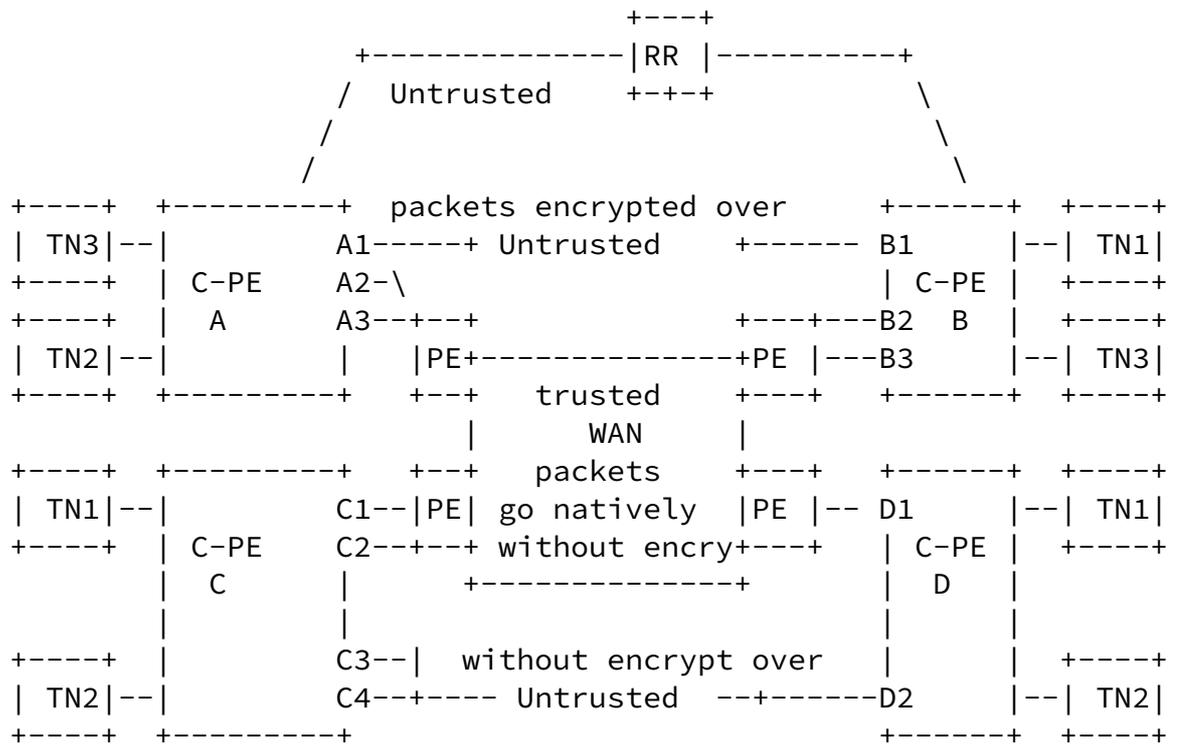


Figure 2: CPEs interconnected by VPN paths and Internet Paths

5.1. Control Plane for Overlay over Heterogeneous Networks

As described in [[BGP-SDWAN-Usage](#)], the Control Plane for Overlay network over heterogeneous networks has three distinct properties:

Dunbar, et al.

Expires November 1, 2020

[Page 12]

Internet-Draft

Net2Cloud Gap Analysis

May 2020

- WAN Port Property registration to the Overlay Controller.
 - o To inform the Overlay controller and authorized peers of the WAN port properties of the Edge nodes. When the WAN ports are assigned private IPv4 addresses, this step can register the type of NAT that translates these addresses into public ones.
- Controller-facilitated IPsec SA management and NAT information distribution
 - o The Overlay controller facilitates and manages the IPsec configuration and peer authentication for all IPsec tunnels terminated at the edge nodes.
- Establishing and Managing the topology and reachability for services attached to the client ports of overlay edge nodes.
 - o This is for the overlay layer's route distribution, so that a C-PE can populate its overlay routing table with entries that identify the next hop for reaching a specific route/service attached to remote nodes. [[SECURE-EVPN](#)] describes EVPN and other options.

5.2. Using BGP UPDATE Messages

5.2.1. Lacking SD-WAN Segments Identifier

There could be multiple SD-WAN networks with their edge nodes exchanging BGP UPDATE messages with the BGP RR. The multiple SD-WAN networks could have common underlay networks. Therefore, it is very important to have an identifier to differentiate BGP UPDATE messages belonging to different SD-WAN networks (or sometimes called SD-WAN Segmentations). Today's BGP doesn't have this feature yet, unless there are multiple BGP instances and their corresponding RRs.

5.2.2. Missing attributes in Tunnel-Encap

[Tunnel-Encap] describes the BGP UPDATE Tunnel Path Attribute that

advertises endpoints' tunnel encapsulation capabilities for the respective attached client routes encoded in the MP-NLRI Path Attribute. The receivers of the BGP UPDATE can use any of the supported encapsulations encoded in the Tunnel Path Attribute for the routes encoded in the MP-NLRI Path Attribute.

Here are some of the issues raised by the use of [\[Tunnel-Encap\]](#) to distribute Edge WAN port properties:

- [\[Tunnel-Encap\]](#) doesn't have the encoding to describe the NAT information for WAN ports that are assigned private IPv4 addresses yet. The NAT information needs to be propagated to the trusted peers such as the virtual C-PEs instantiated in public Cloud DCs via Controllers.
- The mechanism defined in [\[Tunnel-Encap\]](#) does not facilitate the exchange of IPsec SA-specific parameters independently from advertising the attached clients' routes, even after adding a new IPsec tunnel type.
[\[Tunnel-Encap\]](#) requires all tunnels updates to be associated with routes. There can be many client routes associated with an IPsec tunnel established between two C-PEs' WAN ports; the corresponding destination prefixes (as announced by the aforementioned routes) may also be reached through the VPN underlay without any encryption.

The establishment of an IPsec tunnel can fail, e.g., because the two endpoints support different encryption algorithms. Multiple negotiation messages that carry the IPsec SA parameters between two end-points may be exchanged. This is why it is cleaner to separate the establishment of an IPsec SA association between two end-points from the policies enforced to map routes to a specific IPsec SA.

If C-PEs need to establish a WAN Port-based IPsec SA, the information encoded in the Tunnel Path Attribute should only apply to the WAN ports and should be independent from the clients' routes.

In addition, the Overlay IPsec SA Tunnel is very likely to be established before clients' routes are attached.

- When an overlay network spans across large geographic regions (such as countries or continents), one C-PE in one region may not even be aware of remote CPEs in other regions that it needs to communicate. Therefore, the distribution of the Overlay Edge WAN ports information need to be restricted to the authorized peers.

[5.3. SECURE-L3VPN/EVPN](#)

[SECURE-L3VPN] describes a method to enrich BGP/MPLS VPN [[RFC4364](#)] capabilities to allow some PEs to connect to other PEs via public networks. [[SECURE-L3VPN](#)] introduces the concept of Red Interface & Black Interface used by PEs, where the RED interfaces are used to forward traffic into the VPN, and the Black Interfaces are used between WAN ports through which only IPsec-formatted packets are forwarded to the Internet or to any other backbone network, thereby eliminating the need for MPLS transport in the backbone.

[SECURE-L3VPN] assumes PEs use MPLS over IPsec when sending traffic through the Black Interfaces.

[SECURE-EVPN] describes a solution where point-to-multipoint BGP signaling is used in the control plane for the Scenario #1 described in [[BGP-SDWAN-Usage](#)]. It relies upon a BGP cluster design to facilitate the key and policy exchange among PE devices to create private pair-wise IPsec Security Associations without IKEv2 point-to-point signaling or any other direct peer-to-peer session establishment messages.

Both [[SECURE-L3VPN](#)] and [[SECURE-EVPN](#)] are useful, however, they both miss the aspects of aggregating VPN and Internet underlays. In summary:

- Both documents assume that an IPsec tunnel is associated with client traffic. Regardless of which WAN ports the traffic egress from the edge, the client traffic associated with IPsec is always encrypted. Within the context of an overlay architecture that relies upon minimizing resource used for encryption, traffic sent from an edge node can be encrypted once and forwarded through a WAN port towards an untrusted network, but can also remain

unencrypted and be forwarded at different times through a WAN port to the BGP/MPLS VPN.

- The [SECURE-L3VPN] assumes that a CPE "registers" with the RR. However, it does not say how. It assumes that the remote CPEs are pre-configured with the IPsec SA manually. For overlay networks to connect Hybrid Cloud DCs, Zero Touch Provisioning is expected. Manual configuration is not an option.

- The [SECURE-L3VPN] assumes that C-PEs and RRs are connected via an IPsec tunnel. For management channel, TLS/DTLS is more economical than IPsec. The following assumption made by [SECURE-L3VPN] can be difficult to meet in the environment where zero touch provisioning is expected:
 - A CPE must also be provisioned with whatever additional information is needed in order to set up an IPsec SA with each of the red RRs
- IPsec requires periodic refreshment of the keys. The draft does not provide any information about how to synchronize the refreshment among multiple nodes.
- IPsec usually sends configuration parameters to two endpoints only and lets these endpoints negotiate the key. The [SECURE-L3VPN] assumes that the RR is responsible for creating/managing the key for all endpoints. When one endpoint is compromised, all other connections may be impacted.

5.4. Preventing attacks from Internet-facing ports

When C-PEs have Internet-facing ports, additional security risks are raised.

To mitigate security risks, in addition to requiring Anti-DDoS features on C-PEs, it is necessary for C-PEs to support means to determine whether traffic sent by remote peers is legitimate to prevent spoofing attacks, in particular.

6. Gap Summary

Here is the summary of the technical gaps discussed in this document:

- For Accessing Cloud Resources
 - a) When a remote vCPE can be reached by multiple PEs of one provider VPN network, it is not straightforward to designate which egress PE to the remote vCPE based on applications
 - b) Need automated and reliable tools to map the user-friendly (natural language) access rules into machine readable

Dunbar, et al.

Expires November 1, 2020

[Page 16]

Internet-Draft

Net2Cloud Gap Analysis

May 2020

- polices and to provide interfaces for enterprises to self-manage policy enforcement points for their own workloads.
 - c) NAT Traversal. An enterprise's network controller needs to be informed of the NAT properties for its workloads in Cloud DCs. If the workloads are attached to the enterprise's own vCPEs instantiated in the Cloud DCs, the task can be achieved.
 - d) The multicast traffic to/from remote vCPE needs a feature like Appointed Forwarder specified by TRILL to prevent multicast data frames from looping around.
 - e) BGP between PEs and remote CPEs via untrusted networks.
 - f) Traffic Path Management
- Overlay Edge Node's WAN Port Management: BGP UPDATE propagate client's routes information, but don't distinguish network facing ports.
 - Aggregating VPN paths and Internet paths
 - a) Control Plane for Overlay over Heterogeneous Networks is not clear.
 - b) BGP UPDATE Messages missing properties:
 - Lacking SD-WAN Segments Identifier
 - Missing attributes in Tunnel-Encap
 - c) SECURE-L3VPN/EVPN is not enough
 - d) Missing clear methods in preventing attacks from Internet-facing ports

7. Manageability Considerations

Zero touch provisioning of overlay networks to interconnect Hybrid Clouds is highly desired. It is necessary for a newly powered up edge node to establish a secure connection (by means of TLS, DTLS, etc.) with its controller.

8. Security Considerations

Cloud Services are built upon shared infrastructures, therefore not secure by nature.

Dunbar, et al.

Expires November 1, 2020

[Page 17]

Internet-Draft

Net2Cloud Gap Analysis

May 2020

Secure user identity management, authentication, and access control mechanisms are important. Developing appropriate security measurements can enhance the confidence needed by enterprises to fully take advantage of Cloud Services.

9. IANA Considerations

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

10. References

10.1. Normative References

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

10.2. Informative References

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

[RFC5521] P. Mohapatra, E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", April 2009.

[[BGP-SDWAN-PORT](#)] L. Dunbar, et al, "Subsequent Address Family

Indicator for SDWAN Ports", [draft-dunbar-idr-sdwan-port-safi-00](#), Work-in-progress, March 2019.

[BGP-SDWAN-Usage] L. Dunbar, et al, "Framework of Using BGP for SDWAN Overlay Networks", [draft-dunbar-idr-sdwan-framework-00](#), work-in-progress, Feb 2019.

[[Tunnel-Encap](#)] E. Rosen, et al, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-10](#), July 2018.

Dunbar, et al. Expires November 1, 2020 [Page 18]

Internet-Draft Net2Cloud Gap Analysis May 2020

[SECURE-EVPN] A. Sajassi, et al, [draft-sajassi-bess-secure-evpn-01](#), work in progress, March 2019.

[SECURE-L3VPN] E. Rosen, "Provide Secure Layer L3VPNs over Public Infrastructure", [draft-rosen-bess-secure-l3vpn-00](#), work-in-progress, July 2018

[DMVPN] Dynamic Multi-point VPN:
<https://www.cisco.com/c/en/us/products/security/dynamic-multipoint-vpn-dmvpn/index.html>

[DSVPN] Dynamic Smart VPN:
<http://forum.huawei.com/enterprise/en/thread-390771-1-1.html>

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

[Net2Cloud-Problem] L. Dunbar and A. Malis, "Seamless Interconnect Underlay to Cloud Overlay Problem Statement", [draft-dm-net2cloud-problem-statement-02](#), June 2018

11. Acknowledgments

Acknowledgements to John Drake for his review and contributions. Many thanks to John Scudder for stimulating the clarification

discussion on the Tunnel-Encap draft so that our gap analysis can be more accurate.

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

Dunbar, et al.

Expires November 1, 2020

[Page 19]

Internet-Draft

Net2Cloud Gap Analysis

May 2020

Authors' Addresses

Linda Dunbar
Futurewei
Email: ldunbar@futurewei.com

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

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

Dunbar, et al.

Expires November 1, 2020

[Page 20]