Network Working Group
Internet Draft
Intended status: Informational
Expires: October 20, 2023

L. Dunbar
Futurewei
A. Malis
Malis Consulting
C. Jacquenet
Orange
M. Toy
Verizon
K. Majumdar
Microsoft

April 24, 2023

Dynamic Networks to Hybrid Cloud DCs: Problem Statement and Mitigation Practices

draft-ietf-rtgwg-net2cloud-problem-statement-26

Abstract

This document describes the network-related problems enterprises face at the moment of writing this specification when interconnecting their branch offices with dynamic workloads in third-party data centers (a.k.a. Cloud DCs) and some mitigation practices. There can be many problems associated with connecting to or among Cloud DCs; the Net2Cloud problem statements are mainly for enterprises that already have traditional VPN services and are interested in leveraging those networks (instead of altogether abandoning them). Other problems are out of the scope of this document.

This document also describes the mitigation practices for getting around the identified problems.

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."

This Internet-Draft will expire on October 24, 2023.

Copyright Notice

Table of Contents

1. Introduction

With the advent of widely available third-party cloud DCs and services in diverse geographic locations and the advancement of tools for monitoring and predicting application behaviors, it is very attractive for enterprises to instantiate applications and workloads in locations that are geographically closest to their end-users. Such proximity can improve end-to-end latency and overall user experience. Conversely, an enterprise can easily shutdown applications and workloads whenever end-users are in motion (thereby modifying the networking connection of subsequently relocated applications and workloads).

Key characteristics of Cloud Services are on-demand, scalable, highly available, and usage-based billing. Cloud Services, such as, compute, storage, network functions (most likely virtual), third party managed applications, etc. are usually hosted and managed by third party Cloud Operators. Examples of Cloud network functions are: Virtual Firewall services, Virtual private network services, Virtual PBX services including voice and video conferencing systems, etc. Cloud Data Center (DC) is shared infrastructure that hosts the Cloud Services to many customers.

2. Definition of terms

Cloud DC: Third party 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 or more sites.

Heterogeneous Cloud: applications and workloads split among Cloud DCs owned or managed by different operators.

Hybrid Clouds: Hybrid Clouds refers to an enterprise using its own on-premises DCs in addition to Cloud services provided by one or more cloud operators. (e.g. AWS, Azure, Google, Salesforce, SAP, etc).

VPC: Virtual Private Cloud is a virtual network dedicated to one client account. It is logically isolated from other virtual networks in a Cloud DC. Each client can launch his/her desired resources, such as compute, storage, or network functions into his/her VPC. At the moment of of writing this specification, most Cloud operators' VPCs only support private addresses, some support IPv4 only, others support IPv4/IPv6 dual stack.

3. Issues and Mitigation Methods of Connecting to Cloud DCs

There are many problems associated with connecting to Cloud DCs, many of which are out of the IETF scope. This section is to identify some of the high-level problems that can be addressed by IETF, especially by Routing area. Other problems are out of the scope of this document. By no means has this section covered all problems for connecting to Hybrid Cloud Services, e.g., difficulty in managing cloud spending is not discussed here.

3.1. Increased BGP errors and Mitigation Methods

Traditional network service providers usually have prior negotiated peering policies with their BGP peers over fixed interfaces. Cloud GWs need to peer with a larger variety of parties, via private circuits or IPsec over public internet. Many of those peering parties may not be traditional network service providers. Their BGP configurations practices might not be consistent, and some are done by less experienced personnel. All those can contribute to increased BGP peering errors such as capability mismatch, unwanted route leaks, missing Keepalives, and errors causing BGP ceases. Capability mismatch can cause BGP sessions not established properly.
If a BGP speaker receives from its peer a capability that it does not itself support or recognize, it must ignore that capability and the BGP session must not be terminated per RFC5492. When receiving a BGP UPDATE with a malformed attribute, the revised BGP error

handling procedure [RFC7606] should be followed instead of session resetting.

Many Cloud DCs don't support multi hop eBGP peering with external devices. To get around this limitation, it is necessary for enterprises GWs to establish IPsec tunnels to the Cloud GWs to form IP adjacency.

Some Cloud DC eBGP peering only supports limited number of routes from external entities. To get around this limitation, on-premises DCs need to set up default routes to be exchanged with the Cloud DC eBGP peers. When inbound routes exceed the maximum routes threshold for a peer, the current common practice is generating out of band alerts (e.g., Syslog) via management system to the peer, or terminating the BGP session (with cease notification messages [RFC 4486] being sent).

3.2. Site failures and Methods to Minimize Impacts

Failures within a site include (but not limited to) a site capacity degradation or entire site going down. The reasons for these capacity degradations or failures can include: a) fiber cut for links connecting to the site or among pods within the site, b) cooling failures, c) insufficient backup power, d) cyber threat attacks, e) too many changes outside of the maintenance window, or other errors. Fiber-cut is not uncommon within a Cloud site or between sites.

As described in RFC7938, Cloud DC BGP might not have an IGP to route around link/node failures within its domain.

When a site failure happens, the Cloud DC GW visible to clients is running fine; therefore, the site failure is not detectable by the Clients using Bidirectional Forwarding Detection (BFD).

When a site capacity degrades or goes to zero, there are massive numbers of routes needing to be changed.

The large number of routes switching over to another site can also cause overloading that triggers more failures.

In addition, the routes (IP addresses) in a Cloud DC cannot be aggregated nicely, triggering very large number of BGP UPDATE messages when a failure occurs.

It might be more effective to do mass reroute, similar to EVPN [RFC7432] defined mass withdraw mechanism to signal a large number of routes being changed to remote PE nodes as quickly as possible.

3.3. Optimal Paths to Cloud DC locations

Many applications have multiple instances instantiated in different Cloud DCs. A commonly deployed solution has DNS server(s) responding to an FQDN (Fully Qualified Domain Name) inquiry with an IP address of the closest or lowest cost DC that can reach the instance. Here are some problems associated with DNS-based solutions:

- Dependent on client behavior
  - Misbehaving client can cache results indefinitely.
  - Client may not receive service even though there are servers available in other Cloud DCs because the failing IP address is still cached in the DNS resolver and has not expired yet.
- No inherent leverage of proximity information present in the network (routing) layer, resulting in loss of performance.
- Inflexible traffic control:
  The Local DNS resolver becomes the unit of traffic management. This requires DNS to receive periodical update of the network condition, which is difficult.

To address the problems listed above, ANYCAST addresses can be utilized so that network proximity and conditions can be inherently considered in optimal path selection.

3.4. Network Issues for 5G Edge Clouds and Mitigation Methods

The 5G edge clouds [3GPP-5G-Edge] may host edge computing servers (virtual or physical) for Ultra-low latency services that must be near the UEs (User equipment). Those edge computing applications have low latency connections to the UEs and regular connections to backend servers or databases in other locations.

The low latency services traffic to/from the edge Clouds are transported through the 5G Local Data Networks (LDN) and 5G UPFs to the UEs. The LDN's ingress routers, directly connected to the User Plane Functions (UPF), might be co-located with 5G Core functions in the edge Cloud data centers. The 5G Core functions include Radio

Control Functions, Session Management Functions (SMF), Access Mobility Functions (AMF), User Plane Functions (UPF), and others.

Here are some network problems with connecting the services in the 5G Edge Cloud DCs:

1) The difference in routing distances to multiple server instances in different edge Clouds is relatively small. Therefore, the edge Cloud with the shortest routing distance might not be the best in providing the overall latency.
2) Capacity status at the Edge Cloud DC might play a bigger role for E2E performance.
3) Source (UEs) can ingress from different LDN Ingress routers due to mobility.

To get around those problems, the ingress routers can incorporate the destination site's capabilities with the routing distance in computing the optimal paths.

3.5. DNS Practices for Hybrid Workloads

DNS name resolution is essential for on-premises and cloud-based resources. For customers with hybrid workloads, which include on-premises and cloud-based resources, extra steps are necessary to configure DNS to work seamlessly across both environments.

Cloud operators have their own DNS to resolve resources within their Cloud DCs and to well-known public domains. Cloud's DNS can be configured to forward queries to customer managed authoritative DNS servers hosted on-premises and to respond to DNS queries forwarded by on-premises DNS servers.

For enterprises utilizing Cloud services by different cloud operators, it is necessary to establish policies and rules on how/where to forward DNS queries. When applications in one Cloud need to communicate with applications hosted in another Cloud, DNS queries from one Cloud DC could be forwarded to the enterprises' on-premises DNS, which in turn be forwarded to the DNS service in another Cloud. Configuration can be complex depending on the application communication patterns.

However, even with carefully managed policies and configurations, collisions can still occur. If an organization uses an internal name like .internal and then want your services to be available via or within some other cloud provider which also uses .internal, then

collisions might occur. Therefore, it is better to use the global domain name even when an organization does not make all its namespace globally resolvable. An organization's globally unique DNS can include subdomains that cannot be resolved outside certain restricted paths, zones that resolve differently based on the origin of the query, and zones that resolve the same globally for all queries from any source.

Globally unique names do not equate to globally resolvable names or even global names that resolve the same way from every perspective. Globally unique names can prevent any possibility of collisions at present or in the future, and they make DNSSEC trust manageable. Consider using a registered and fully qualified domain name (FQDN) from global DNS as the root for enterprise and other internal namespaces.

3.6. NAT Practice for Accessing Cloud Services

Cloud resources, such as VM instances, are usually assigned private IP addresses. By configuration, some private subnets can have the NAT function to reach out to external networks, and some private subnets are internal to Cloud only.

Different Cloud operators support different levels of NAT functions. For example, AWS NAT Gateway does not currently support connections towards, or from VPC Endpoints, VPN, AWS Direct Connect, or VPC Peering [AWS-NAT]. AWS Direct Connect/VPN/VPC Peering does not currently support any NAT functionality.

Google's Cloud NAT [Google-NAT] allows Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters to connect to the Internet. Cloud NAT implements outbound NAT in conjunction with a default route to allow instances to reach the Internet. It does not implement inbound NAT. Hosts outside the VPC network can only respond to established connections initiated by instances inside the Google Cloud; they cannot initiate new connections to Cloud instances via NAT.

For enterprises with applications running in different Cloud DCs, proper configuration of NAT must be performed in Cloud DCs and their on-premises DC.

3.7. Cloud Discovery Practices

One of the concerns of using Cloud services is not being aware of where the resource is located, especially that Cloud operators can

move application instances from one place to another. When applications in Cloud communicate with on-premises applications, it may not be clear where the Cloud applications are located or to which VPCs they belong.

Being able to detect Cloud services location can help on-premises gateways (routers) to switch the services to a more optimal site when the current cloud site encounters failures or degradation. A significant difference is that cloud discovery uses the cloud vendor's API to extract data on your cloud services rather than the direct access used in scanning your on-premises infrastructure.

For enterprises that instantiate virtual routers in Cloud DCs, metadata can be attached (e.g., GENEVE header or IPv6 optional header) to indicate Geo-location of the Cloud DCs.

4. Dynamic Interconnecting Enterprise Sites with Cloud DCs

For many enterprises with established provide VPNs (e.g., private circuits, MPLS-based L2VPN/L3VPN) interconnecting branch offices & on-premises data centers, connecting to Cloud services will be a mix of different types of networks. When an enterprise's existing VPN service providers do not have direct connections to the desired cloud DCs that the enterprise prefers to use, the enterprise faces additional infrastructure and operational costs to utilize the Cloud services.

This section describes practices to connect to Cloud services.

4.1. Sites to Cloud DC

Most Cloud operators offer some type of network gateway through which an enterprise can reach their workloads hosted in the Cloud DCs. For example, AWS (Amazon Web Services) offers the following options to reach workloads in AWS Cloud DCs [AWS-Cloud-WAN]:

- AWS Internet gateway allows communication between instances in AWS VPC and the internet.
- AWS Virtual gateway (vGW) where IPsec tunnels [RFC6071] are established between an enterprise's own gateway and AWS vGW, so

that the communications between those gateways can be secured from the underlay (which might be the public Internet).

- AWS Direct Connect, which allows enterprises to purchase direct connect from network service providers to get a private leased line interconnecting the enterprises gateway(s) and the AWS Direct Connect routers. In addition, an AWS Transit Gateway can be used to interconnect multiple VPCs in different Availability Zones. AWS Transit Gateway acts as a hub that controls how traffic is forwarded among all the connected networks which act like spokes.

Microsoft Azure's Virtual WAN [Azure-SD-WAN] allows extension of a private network to any of the Microsoft cloud services, including Azure and Office365. ExpressRoute is configured using Layer 3 routing. Customers can opt for redundancy by provisioning dual links from their location to two Microsoft Enterprise edge routers (MSEEs) located within a third-party ExpressRoute peering location. The BGP routing protocol is then setup over WAN links to provide redundancy to the cloud. This redundancy is maintained from the peering data center into Microsoft's cloud network.

Google's Cloud Dedicated Interconnect offers similar network connectivity options as AWS and Microsoft. One distinct difference, however, is that Google's service allows customers access to the entire global cloud network by default. It does this by connecting the on-premises network with the Google Cloud using BGP and Google Cloud Routers to provide optimal paths to the different regions of the global cloud infrastructure.

Figure 1 below shows an example of some of a tenant's workloads that are accessible via a virtual router connected by AWS Internet Gateway; some are accessible via AWS vGW, and others are accessible via AWS Direct Connect.

Different types of access require different level of security functions. Sometimes it is not visible to end customers which type of network access is used for a specific application instance. To get better visibility, separate virtual routers (e.g., vR1 & vR2) can be deployed to differentiate traffic to/from different cloud GWs. It is important for some enterprises to be able to observe the specific behaviors when connected by different connections.

Customer Gateway can be customer owned router or ports physically connected to AWS Direct Connect GW.

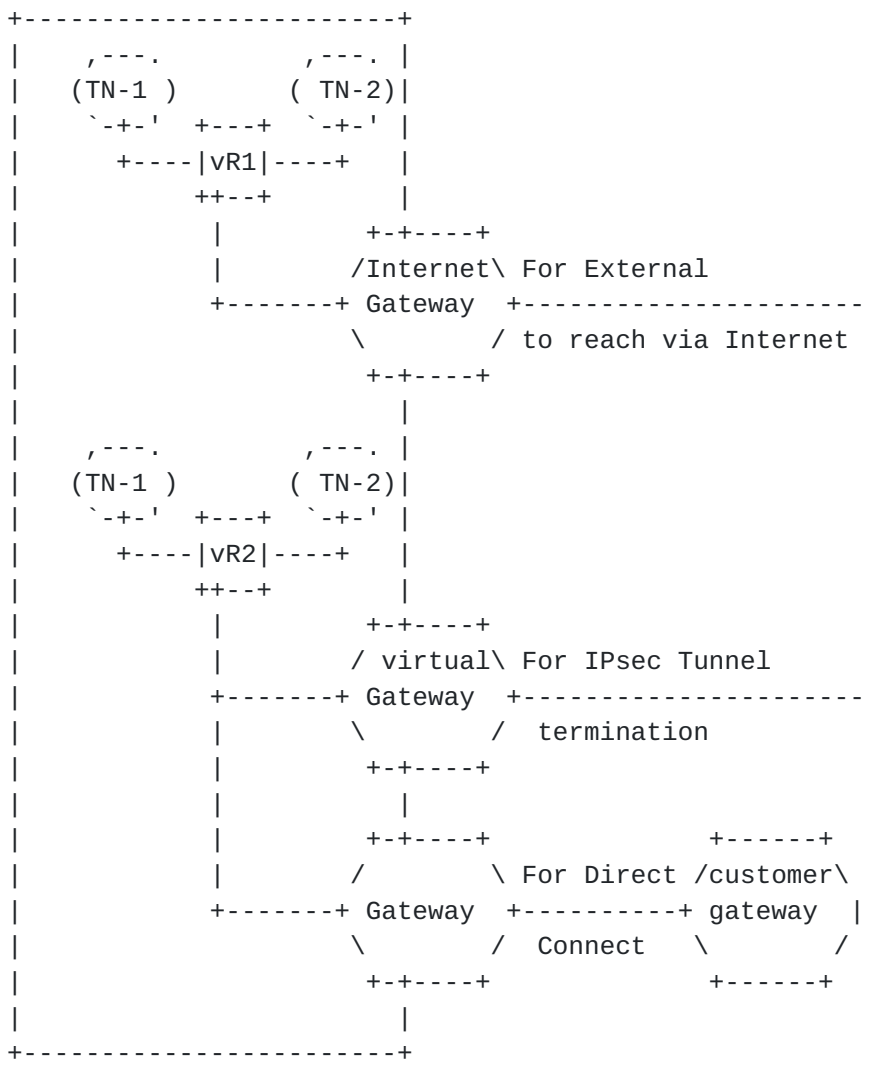

Figure 1: Examples of Multiple Cloud DC connections.

4.2. Inter-Cloud Connection

The connectivity options to Cloud DCs described in the previous section are for reaching Cloud providers' DCs, but not between cloud DCs. When applications in AWS Cloud need to communicate with applications in Azure, today's practice requires a third-party gateway (physical or virtual) to interconnect the AWS's Layer 2 DirectConnect path with Azure's Layer 3 ExpressRoute.

Enterprises can also instantiate their own virtual routers in different Cloud DCs and administer IPsec tunnels among them, which by itself is not a trivial task. For example, open-source VPN software such as strongSwan can be leveraged to create an IPSec connection to the Azure gateway. The StrongSwan instance within AWS

not only can connect to Azure but can also be used to facilitate traffic to other nodes within the AWS VPC by configuring forwarding and using appropriate routing rules for the VPC.

Most Cloud operators, such as AWS VPC or Azure VNET, use non-globally routable CIDR from private IPv4 address ranges as specified by RFC1918. To establish IPsec tunnel between two Cloud DCs, it is necessary to exchange Public routable addresses for applications in different Cloud DCs.

In summary, here are some approaches, available to interconnect workloads among different Cloud DCs:

a) Utilize Cloud DC provided inter/intra-cloud connectivity services (e.g., AWS Transit Gateway) to connect workloads instantiated in multiple VPCs. Such services are provided with the cloud gateway to connect to external networks (e.g., AWS DirectConnect Gateway).

b) Hairpin all traffic through the customer gateway, meaning all workloads are directly connected to the customer gateway, so that communications among workloads within one Cloud DC must traverse through the customer gateway.

c) Establish direct tunnels among different VPCs (AWS' Virtual Private Clouds) and VNET (Azure's Virtual Networks) via client's own virtual routers instantiated within Cloud DCs. NHRP (Next Hop Resolution Protocol) [RFC2735] based multi-point techniques can be used to establish direct Multi-point-to-Point or multi-point-to multi-point tunnels among those client's own virtual routers.

Approach a) usually does not work if Cloud DCs are owned and managed by different Cloud providers.

Approach b) creates additional transmission delay plus incurring cost when exiting Cloud DCs.

For the Approach c), NHRP [RFC2735] is used for spoke nodes to register their IP addresses & WAN ports with the hub node. The IETF ION (Internetworking over NBMA (non-broadcast multiple access)) WG standardized NHRP for connection oriented NBMA network (such as ATM) network address resolution more than two decades ago.

There are many differences between virtual routers in Public Cloud DCs and the nodes in an NBMA network. NHRP cannot be used for

registering virtual routers in Cloud DCs unless an extension of such protocols is developed for that purpose, e.g., taking NAT or dynamic addresses into consideration. Therefore, existing NHRP based VPN technique cannot be used directly for connecting workloads in hybrid Cloud DCs.

4.3. Extending Private VPNs to Hybrid Cloud DCs

Traditional private VPNs, including private circuits or MPLS based L2/L3 VPNs, have been widely deployed as an effective way to support businesses and organizations that require network performance and reliability. L2/L3 VPN shifts the burden of managing a VPN service from enterprises to service providers. The CPEs attached to a private VPN are simpler and less expensive because they do not need to manage routes to remote sites; they pass all outbound traffic to the private VPN PEs to which the CPEs are attached (albeit multi-homing scenarios require more processing logic on CPEs). Private VPN has addressed the problems of scale, availability, and fast recovery from network faults, and incorporated traffic-engineering capabilities.

However, an enterprise's private VPN's PE (Provider Edge) nodes might not have the direct connections to the third-party cloud DCs needed by the enterprise to provide easy access to its end users. When the user base changes, the enterprise might want to migrate its workloads/applications to a new cloud DC location closest to the new user base. The existing private VPN provider might not have circuits at the new location. Deploying PEs routers at new locations takes a long time (weeks if not months), which defeats one of the benefits of Clouds' geographically diverse locations allowing workloads to be as close to their end-users as possible.

When the private VPN network can't reach the desired Cloud DCs, IPsec tunnels can be used to dynamically connect the private VPN PEs with the desired Cloud DCs GWs. As the private VPNs provide more secure and higher quality services, choosing a PE closest to the Cloud GW for the IPsec tunnel is desirable to minimize the IPsec tunnel distance over the public Internet.

In order to support Explicit Congestion Notification (ECN) [RFC3168] usage by private VPN traffic, the PEs that establish the IPsec

tunnels with the Cloud GW need to comply with the ECN behavior specified by RFC6040 [RFC6040].

An enterprise can connect to multiple Cloud DC locations and establish different BGP peers with Cloud GW routers at different locations. As multiple Cloud DCs are interconnected by the Cloud provider's own internal network, its topology and routing policies are not transparent or even visible to the enterprise customer's on-prem routers. One Cloud GW BGP session might advertise all of the prefixes of the enterprise's VPC, regardless of which Cloud DC a given prefix resides, which can cause improper optimal path selection for on-prem routers. To get around this problem, virtual routers in Cloud DCs can be used to attach metadata (e.g., in the GENEVE header or IPv6 optional header) to indicate the Geo-location of the Cloud DC, the delay measurement, or other relevant data.

5. Methods to Scale IPsec tunnels to Cloud DCs

As described in Section 4.3, IPsec tunnels can be used to dynamically establish connection between private VPN PEs with Cloud GW. Enterprises can also instantiate virtual routers within Cloud DCs to connect to their on-premises devices via IPsec tunnels.

As described in [Int-tunnels], IPsec tunnels can introduce MTU problems. This document assumes that endpoints manage the appropriate MTU sizes, therefore, not requiring VPN PEs to perform the fragmentation when encapsulating user payloads in the IPsec packets.

5.1. Scaling Issues with IPsec Tunnels

IPsec tunnels are a very convenient solution for an enterprise with limited locations to reach a Cloud DC.

However, for a medium-to-large enterprise with multiple sites and data centers to connect to multiple cloud DCs, there are N*N number of IPsec tunnels among Cloud DC gateways and all those sites. Each of those IPsec Tunnels requires pair-wise periodic key refreshment. For a company with hundreds or thousands of locations, managing hundreds (or even thousands) of IPsec tunnels can be very processing intensive. That is why many Cloud operators only allow a limited number of (IPsec) tunnels & bandwidth to each customer.

To scale the IPsec key management, a solution like group encryption where a single IPsec SA is necessary at the GW can be considered. But the drawback of the group encryption is higher security risk of the key distribution and maintenance of a key server.

5.2. Poor performance when overlay public internet

IPsec encap & decap are very processing intensive, which can degrade router performance. NAT also adds to the performance burden.

When enterprise CPEs or gateways are far away from cloud DC gateways or across country/continent boundaries, performance of IPsec tunnels 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.

Many cloud providers can replicate workloads in different available zones. An App instantiated in a cloud DC closest to clients may have to cooperate with another App (or its mirror image) in another region or database server(s) in the on-premises DC. This kind of coordination requires predicable networking behavior/performance among those locations.

6. Requirements for Dynamic Cloud Data Center VPNs

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:

- Global reachability from different geographical zones, thereby facilitating the proximity of applications as a function of the end users' location, to improve latency.
- Elasticity: prompt connection to newly instantiated applications at Cloud DCs when usages increase and prompt release of connection after applications at locations being removed when demands change.
- Scalable policy management: apply the appropriate polices to the newly instantiated application instances at any Cloud DC locations.
- The solution should allow enterprises to take advantage of the current state-of-the-art private VPN technologies, including

the traditional circuit-based, MPLS-based VPNs, or IPsec-based VPNs (or any combination thereof) that run over the public Internet.

- The solution should not require an enterprise to upgrade all their existing CPEs.
- The solution should support scalable IPsec key management among all nodes involved in DC interconnect schemes.
- The solution needs to support easy and fast, on-the-fly, 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.

7. Security Considerations

The security issues in terms of networking to clouds include:

- Service instances in Cloud DCs are connected to users (enterprises) via Public IP ports which are exposed to the following security risks:

  a) Potential DDoS attack to the ports facing the untrusted network (e.g., the public internet), which may propagate to the cloud edge resources. To mitigate such security risk, it is necessary for the ports facing internet to enable Anti-DDoS features.

  b) Potential risk of augmenting the attack surface with inter-Cloud DC connection by means of identity spoofing, man-in-the-middle, eavesdropping or DDoS attacks. One example of mitigating such attacks is using DTLS to authenticate and encrypt MPLS-in-UDP encapsulation (RFC 7510).

- Potential attacks from service instances within the cloud. For example, data breaches, compromised credentials, and broken authentication, hacked interfaces and APIs, account hijacking.

- When IPsec tunnels established from enterprise on-premises CPEs are terminated at the Cloud DC gateway where the workloads or applications are hosted, traffic to/from an enterprise's workload can be 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, IPsec tunnels may go all the way to the workload (servers, or VMs) within the DC.

Many Cloud operators offer monitoring services for data stored in Clouds, such as AWS CloudTrail, Azure Monitor, and many third-party monitoring tools to improve visibility to data stored in Clouds.

Solution drafts resulting from this work will address security concerns inherent to the solution(s), including both protocol aspects and the importance (for example) of securing workloads in cloud DCs and the use of secure interconnection mechanisms.

8. IANA Considerations

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

9. References

9.1. Normative References

[RFC2735] B. Fox, et al "NHRP Support for Virtual Private networks". Dec. 1999.

[RFC4271] Y. Rekhter, et al "BGP-4". Jan 2006

[RFC5492] J. Scudder, R. Chandra "Capabilities Advertisement with BGP-4". Feb 2009

[RFC6040] B. Briscoe, "Tunnelling of Explicit Congestion Notification", RFC6040, Nov 2010.

[RFC7606] E. Chen, et al "Revised Error Handling for BGP UPDATE Messages". Aug 2015.

[RFC7938] P. Lapukkov, et al "Use of BGP for Routing in Large-Scale Data Centers", Aug. 2016

9.2. 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.

[3GPP-5G-Edge] 3GPP TS 23.548 v18.1.1, "5G System Enhancements for Edge Computing", April 2023.

[AWS-NAT] NAT gateways - Amazon Virtual Private Cloud.

[AWS-Cloud-WAN] Introducing AWS Cloud WAN (Preview) | Networking & Content Delivery (amazon.com).

[Azure-NAT] What is Azure Virtual Network NAT? | Microsoft Learn

[Azure-SD-WAN] Architecture: Virtual WAN and SD-WAN connectivity - Azure Virtual WAN | Microsoft Learn.

[Google-NAT] Cloud NAT overview | Google Cloud.

[Int-tunnels] J. Touch and W Townsley, "IP Tunnels in the Internet Architecture", draft-ietf-intarea-tunnels-13.txt, March, 2023.

10. Acknowledgments

Many thanks to Adrian Farrel, Alia Atlas, Chris Bowers, Paul Vixie, Paul Ebersman, Timothy Morizot, Ignas Bagdonas, Donald Eastlake, Michael Huang, Liu Yuan Jiao, Katherine Zhao, and Jim Guichard for the discussion and contributions.

Authors' Addresses

Linda Dunbar
Futurewei
Email: Linda.Dunbar@futurewei.com

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

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

Mehmet Toy
Verizon
One Verizon Way
Basking Ridge, NJ 07920
Email: mehmet.toy@verizon.com

Kausik Majumdar
Microsoft Azure
kmajumdar@microsoft.com