Enterprise Networking with Secure Acess Service Edge

New Role of the Internet amid COVID-19 Crisis

Zhaobo Zhang
Silicon Valley Network Technology Lab.
Futurewei Technologies Inc.
Santa Clara, CA, USA
e-mail: zzhang1@futurewei.com

***Abstract*—Since the coronavirus (COVID-19) outbreak, people are forced to stay home worldwide to prevent the spread. The Internet has never been more essential to support people's daily lives. Organizations have been scrambling to establish large-scale access for remote workforce. The requirements of large numbers of secure and reliable connections from home and mobile devices bring new challenges to enterprise networking. The Secure Access Service Edge (SASE) architecture converges network and security into cloud-based services, which provides a fast, scalable and secure way for employees to connect remotely. This paper presents the evolution of Internet and enterprise networking with the rise of cloud computing, and the recent industry progress on the SASE-based solutions. Direct cloud access and converged secure edge are the future directions of enterprise networking.**

Keywords-Enterprise Networking; Enterprise Mobility; Secure Access Service Edge; Endpoint Security; Point-of-Presences; SD-WAN.

I. INTRODUCTION

Since the coronavirus disease 2019 (COVID-19) outbreak first identified in December 2019 in China, about 26 million cases have been confirmed from 188 countries in eight months [1]. This pandemic created unprecedented disruptions to human society. Social distancing and shelter-in-place orders were announced to prevent the spread. With most people staying at home, the internet inevitably becomes the most important channel to connect people and deliver critical services, like telemedicine and videoconferencing.

The average daily in-home data usage in the U.S. soared 38 percent in mid-March as the coronavirus pandemic started, compared to the same time in March 2019 [2]. The monthly average of data consumption in the first quarter of 2020 per subscriber has increased to 402.5 GB from 273.5 GB during the same time last year, a 47 percent increase [3]. Increased data usage is mainly from social media platforms, streaming platforms, and online gaming. Under the strain of unprecedented usage, the European Union even urged streaming platforms, like Netflix and YouTube, to stop showing video in high definition to prevent the internet from breaking down. Fortunately, the Internet held up to the surge and continuously support the heavy in-home data usage. Increased disruptions across provider networks have been observed, but those are mainly caused by increased traffic engineering activity, rather than traffic congestion [4].

Together with Internet Service Providers (ISPs), Cloud Providers (CPs) provide significant support in terms of Internet traffic delivery during the pandemic, especially for many ubiquitous digital services, making them effectively an extension of the Internet. For example, in order to support video conferencing, Oracle Cloud transferred 7 PB of data and supported 300 million meeting participants daily for Zoom, in April 2020 [5].

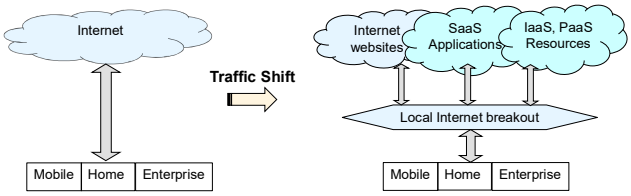

Figure 1. Internet Traffic Shift in the Cloud Era.

Before the pandemic, the Internet had been reshaped with the rise of cloud computing. Its role got enriched from delivering contents to delivering computing resources. Since Amazon first announced delivering computing and storage resource over Internet in 2006, the cloud-based technologies including Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-service (SaaS) have made great strides, and the Internet traffic has been shifted, illustrated in Figure 1. The global network infrastructure has been greatly improved as well, with the rapid growth of CPs private backbone.

This paper aims to provide an overview of the changes of to the Internet and enterprise networks with the rise of cloud computing, particularly under the crisis of COVID-19. In Section II, the private backbone networks of CPs are first introduced. In Section III, the transformation of enterprise Wide Area Network (WAN) is discussed. In Section IV, SASE, a network and security converged solution is explained. Conclusions are presented at the end.

II. CLOUD PROVIDERS GLOBAL BACKBONE

With the wide adoption of cloud computing, Cloud providers expand their global backbone network at an incredible speed. Improving the availability and efficiency of WAN is central to their ability to provide services in a fast, reliable and cost-effective manner. In Figure 2, a simplified private backbone network is illustrated. Cloud providers build data centers and Point-of-Presences (PoPs) globally to increase accessibility, and connect them with optical fiber,

 ISBN: 978-1-61208-822-8

subsea cables and super high bandwidth switches. Therefore, instead of using public Internet, customers at different locations can connect to the nearest PoPs, and then communicate within CPs' private network to reduce latency.

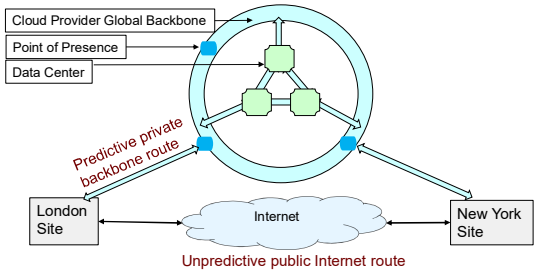

Figure 2. Cloud Providers Global Backbone Networks.

For example, Microsoft owns and runs one the largest WAN backbones in the world. By now, it has built more than 60 regions (data centers), more than 170 global network PoPs, and 130,000 miles of optical fiber with elastic optical network [6][7]. Bandwidth between data centers is up to 1.6 PB/s [6]. Using private subsea cable systems, network latency and performance predictability across continents have been greatly improved, compared to using public Internet [8]. Google Cloud and Amazon Web Services have taken similar strategies to expand their global reach. More details on cloud performance comparison can be found in [8].

As COVID-19 lockdowns went into effect, network traffic and the usage of cloud services significantly increased. According to Microsoft's statistics, the Azure WAN usage immediately jumped by 60 percent after the US lockdown, and 94 percent growth was seen in connections to the Azure VPN service [9]. Traffic into Azure networks that connect data center regions to each other and to the outside world spiked by 40 times [9]. To meet the incredible surge in demand, workloads from the hot regions were shifted to regions that had more capacity. More gateways and front ends were added to scale out services. Duplicate and non-essential features in the services were minimized to reduce the resource usage. Meanwhile, cloud providers rapidly increased network and data center capacity, opened new sites, signed new contracts with ISPs, and even leased extra subsea cables. Network bandwidth got greatly increased within a few months. Although some logistical issues were caused by the pandemic, the expansion of cloud infrastructure has not been slowing down. A total of 26 hyperscale data centers opened in the first half of 2020, and 176 new data centers are in the pipeline [10].

III. ENTERPRISE WAN TRANSFORMATION

For more than a decade, Multi-Protocol Label Switching (MPLS) was the standard approach for building a corporate WAN since 2000. However, due to its high costs and limited agility, a wave of transition to Software-Defined WAN (SD-WAN) started around 2013. Early SD-WAN deployments reuse the MPLS networks and broadband Internet connections at enterprise branches. With dynamic path selection and application visibility, SD-WAN improves the network availability and manageability, and lower the costs. However, the unpredictability of the Internet limits the growth of early SD-WAN.

With the mainstream adoption of cloud, more and more enterprise applications move to the SaaS model, e.g., Office 365, Workday, Service Now. According to the Enterprise Strategy Group's 2019 report, 60 percent of organizations will use SaaS applications for greater than half of their business needs over the next two years [11]. Therefore, the enterprise traffic to headquarters shifts to the public cloud. Direct cloud access drives the recent changes of enterprise WAN, which is considered as SD-WAN 2.0. To facilitate the changes, CPs directly connect their network with ISPs networks in the Internet Exchange Points (IXPs), or provide ISPs the connections to their nearest PoPs. Customers can purchase premium connections to public clouds from ISPs. For large enterprises, enterprise network appliances can even be placed at the same colocation center with CPs. A dedicated link up to 100 Gbps can be established directly between enterprise network and CPs network, and multiple links can be configured if needed.

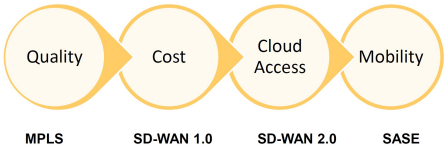

Figure 3. Drivers of Enterprise WAN Transformation at Different Phases.

As of today, with a large number of employees working from home in this pandemic, the changes of enterprise WAN are driven by mobility. Previously, remote employees used VPN technologies to securely connect to the headquarters. However, since enterprise applications and data are now located among various cloud platforms, it is not efficient to backhaul all the traffic to headquarters and then redirect to the Internet. In addition, the VPN technologies present security challenges in a cloud environment. With VPN connections, users are given broad access to an entire flat network, rather than to only the applications that are needed for their jobs. To respond to this challenge, industry moves from VPNs to Zero Trust (ZT). Zero Trust security includes two core concepts: 1) provide and continuously authorize access to resources based on identity, instead of location. 2) enforce the principle of least privilege, i.e., only grant users (or services) access to assets they specifically need, and nothing more [12]. ZT security now has been widely accepted by both users and vendors.

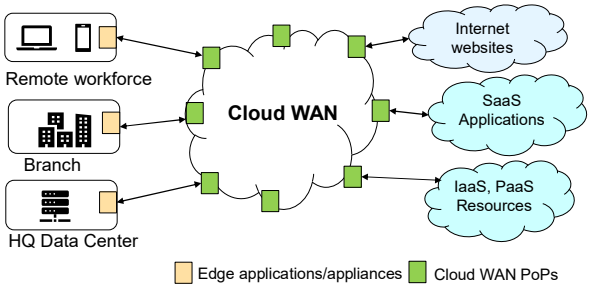

Figure 4. A Cloud-native Indentity-based Enterprise WAN.

Driven by direct cloud access and mobility, enterprise WAN architecture evolves to a cloud-native, identity-based, network and security converged architecture, illustrated in Figure 4. A global distributed architecture provides fast cloud access. ZT framework guarantees secure connections from any devices and anywhere. Cloud-native architecture provides dynamic service orchestration and centralized network and security management. Its key characteristics include:

- Zero trust framework
- Cloud-native architecture
- Globally distributed network
- Security-as-a-service at cloud edge

There are many names for this unified architecture. One of them is SASE, the convergence of network services and security services, coined by Gartner in July 2019 [13]. The SASE architecture and SASE-centric solutions have been enriched by multiple vendors in the past year. The capabilities, benefits and main components of SASE-based enterprise WAN are elaborated in Section IV.

IV. SECURE ACCESS SERVICE EDGE

Due to the uncertainties brought by the pandemic, business agility and scalability are crucial for enterprise networks, and the dynamic secure access requirements put significant pressure on existing networks. Therefore, a cloud-native solution with converged security services becomes a necessity. SASE is such a distributed and cloud-native platform that connects all edges to one logic network and delivers network services and security services as needed. Its main four characteristics are shown in Figure 5.

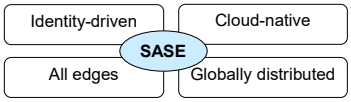

Figure 5. Four Characteristics of SASE.

Compared to previous SD-WAN solutions, SASE emphasize a converged design to streamline and optimize security solutions, including Zero Trust Network Access (ZTNA), Firewall-as-a-service (FWaaS), Secure Web Gateway (SWG), Cloud Access Security Broker (CASB), Web Application and API Protection (WAAP), Remote Browser Isolation (RBI), Intrusion Prevention System (IPS), Data Loss Prevention (DLP), etc. It is not required to implement all these functions at once. Gartner categorized SASE capabilities to three levels, illustrated in Figure 6. Additionally, the capabilities are not achieved by a simple service function chaining. Parallel processing should be adopted to minimize processing latency.

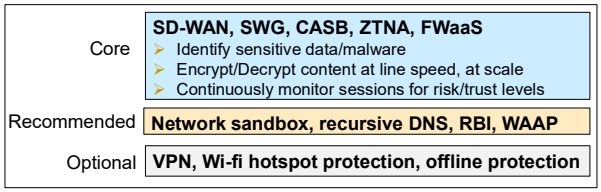

Figure 6. SASE capabilities by Gartner [13].

Some key benefits of SASE include greater business agility, dynamic access control for different user groups, fine-grained visibility of user activities, end-to-end traffic observability, reliable network performance, increased operational efficiency, reduced management complexity and total cost of ownership, etc. Those benefits are mainly from enterprise IT's perspective. As for the employees, better user experience and network performance should be seen. In order to implement SASE-centric networks, three main components are listed as following.

A. *Edge Connector*

Edge connectors are responsible for securely driving traffic from edges to the cloud network for processing. Different edges may have different forms of connectors. For example, branches could have an IPsec-enabled plug-and-play device to connect to the local Internet; personal computers and cellphones may use an identity-based security application or a web-based agentless terminal; data centers could adopt high performance routers to guarantee bandwidth.

B. *Point-of-Presences (PoPs)*

To provide low-latency access to users, devices and cloud services anywhere, enterprises need a worldwide fabric of PoPs and peering relationship [13]. The PoPs comprise a backbone network interconnecting multiple provider networks with SLA-backed IP connections, offering better performance and resiliency than the traditional unpredictive Internet. Converged network services and security services running on each PoP steer network traffic, inspect security, and monitor underlying networks, etc. Services are designed in a cloud-native way. Benefit from scalability and elasticity of the cloud, services can be easily deployed and scaled out among all PoPs.

Since most processing is done on the PoPs, connecting any type of edges, e.g., branches, cloud resources, mobile devices, IoT devices, becomes simple and consistent. The only requirement is a simple connector to establish an encrypted tunnel across an Internet connection to the nearest PoP. By co-locating PoPs and cloud IXPs in the same physical data centers, cloud resources are directly available in the enterprise networks. Minimal latency is achieved without deploying additional software or hardware.

C. *Cloud-based Orchestration and Management Center*

In contrast to previous point solutions, enterprise IT can manage their networks and security through one unified operation center, e.g., enabling connections to new branches, authorizing new users, prioritizing office applications, modifying firewall policies, and upgrading services. Granular access at both network level and application level can be programmed for each individual, and consistent network policies will be deployed to each PoP. In addition, standardized metrics, logs, and tracing are built into cloud-native services, and end-to-end observability is available in the operation center for performance analysis and fast troubleshooting.

 ISBN: 978-1-61208-822-8

V. CONCLUSIONS

Cloud providers have expanded global backbone networks rapidly, which play an important role in delivering Internet traffic. ISPs and CPs collaborate closely to improve the Internet performance and resiliency. The Internet and cloud are inseparable and synergize each other today.

COVID-19 is a catalyst, accelerating the business transition to a more flexible model. As enterprise applications move to the cloud and employees move remotely, enterprise networks have no longer been in a constraint perimeter. A trend of SASE-based WAN composed of edge connectors and global PoPs emerges to better support the needs of direct cloud access and secure access from any devices anywhere. The SASE architecture provides good guidance, but the convergence design of network and security services at the cloud edge is still ongoing work.

REFERENCES

[1] E. Dong, H. Du, and L. Gardner, “An interactive web-based dashboard to track COVID-19 in real time,” in the Lancet Infectius Disease, vol. 20, no. 5, pp. 533-534, May 2020.

[2] Statista, “Coronavirus: impact on online usage in the U.S.”, https://www.statista.com/statistics/1106863/COVID-19-daily-in-home-data-usage-change-us-2020/, [retrieved Sept. 2020].

[3] OpenVault, “Boradband Insights Report of Q1 2020”, https://openvault.com/complimentary-report-Q120/, [retrieved Sept. 2020].

[4] ThousandEyes, “Internet Performance Report”, 2020, https://www.thousandeyes.com/resources/internet-performance-report-COVID-19-impact/, [retrieved Sept. 2020].

[5] Oracle, “Facing unprecedented growth, Zoom turns to Oracle Cloud Infrastructure to support millions of users”, https://www.oracle.com/customers/zoom.html/, [retrieved Sept. 2020].

[6] Microsoft Azure, “Azure global network”, https://azure.microsoft.com/en-us/global-infrastructure/global-network/, [retrieved Sept. 2020].

[7] M. Filer, et al., "Elastic optical networking in the microsoft cloud," in IEEE/OSA Journal of Optical Communications and Networking, vol. 8, no. 7, pp. A45-A54, 2016.

[8] ThousandEyes, “Cloud Performance Benchmark”, 2019, https://www.thousandeyes.com/resources/cloud-performance-benchmark-report-november-2019/, [retrieved Sept. 2020].

[9] Microsoft Azure, “Azure responds to COVID-19” , https://azure.microsoft.com/en-us/blog/azure-responds-to-covid19/ , [retrieved Oct. 2020].

[10] Synergy Research Group, “Hyperscale Data Center Count”, https://www.srgresearch.com/articles/hyperscale-data-center-count-reaches-541-mid-2020-another-176-pipeline, [retrieved Oct. 2020].

[11] Enterprise Strategy Group, “The Rise of Direct Internet Access”, 2019, https://security.umbrella.com/esg-report-rise-of-dia/, [retrieved Sept. 2020].

[12] National Institue of Standards and Technology (NIST), “Zero Trust Architecture”, NIST Special Publication, 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207-draft2.pdf/, [retrieved Sept. 2020].

[13] Gartner, “The Future of Network Security is in the Cloud”, https://www.valtix.com/uploads/secure-access-service-edge-gartner.pdf/, [retrieved Oct. 2020].

 ISBN: 978-1-61208-822-8