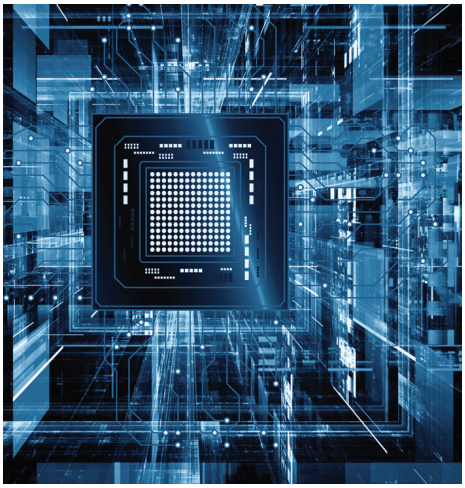

Cloud Networks: Enhancing Performance and Resiliency

Stefano Secci, *Université Pierre et Marie Curie*
San Murugesan, *BRITE Professional Services*

New and ever more sophisticated cloud applications pose key challenges for improving cloud network performance and resiliency.

Cloud networks—networks connecting cloud services and linking datacenter networks—are key elements for successfully embracing the cloud and for fostering deployment of new applications that demand higher cloud service performance. Without functional and resilient cloud networks, the services they're designed to deliver are rendered suboptimal, or even useless.

However, legacy approaches to cloud network provisioning often treat the network as a dumb pipeline—overprovisioning it in advance with enough idle bandwidth to allow for the fastest possible connection and rapid growth in usage and demand by new clients and applications. This has led to a centralization of services at large, monolithic datacenters—an approach justified by economies of scale, perhaps, but essentially driven by hosted legacy services, with Web HTTP services accounting for the majority.

TRADITIONAL DATACENTERS

Such large datacenters certainly have advantages: they are relatively easy to operate and maintain, allowing staff to implement new capacity with minimal disruption and respond to outages quickly. But there are also practical limits to their size.

Increasing demands for electricity and overall datacenter eco-sustainability create basic physical constraints. Just as important, though, are service-level agreement (SLA) requirements for highly reliable and consistently available cloud services. Disaster resiliency, for one, has become a critically important consideration in choosing cloud applications and cloud service providers—particularly in the aftermath of 2012's Hurricane Sandy, which resulted in multiday network downtime for many northeastern US Web services, including IEEE's.

Such concerns call for a resilient cloud network that distributes services across multiple sites geographically. Besides achieving redundancy goals, datacenter distribution can help to improve performance and facilitate deployment of emerging cloud applications such as seamless remote desktop, online networked gaming, and augmented reality.

RESILIENT, GEOGRAPHICALLY DISTRIBUTED DATACENTERS

Power outages are the most common failure affecting high-density datacenters, followed by accidental intra-datacenter link disconnection.[1] In addition, external datacenter transit-provision networks can affect cloud access reliability. So, to support the near 100 percent reliability required by increasingly common reliability-specific SLAs associated with infrastructure as a service (IaaS), the current trend is to deploy IaaS virtual machines across multiple, geographically distant sites.

However, distributing services across distant sites—or cloud datacenters—as depicted in Figure 1, is challenging from a networking perspective. The ecosystem of network communication protocols used in the wide-area segment external to the datacenter is extremely heterogeneous, and difficult to manage from the edges. Further, each IaaS requires at least one and perhaps multiple virtual networks. Within a virtual network, virtual machines are volatile: attachment points can

 Published by the IEEE Computer Society

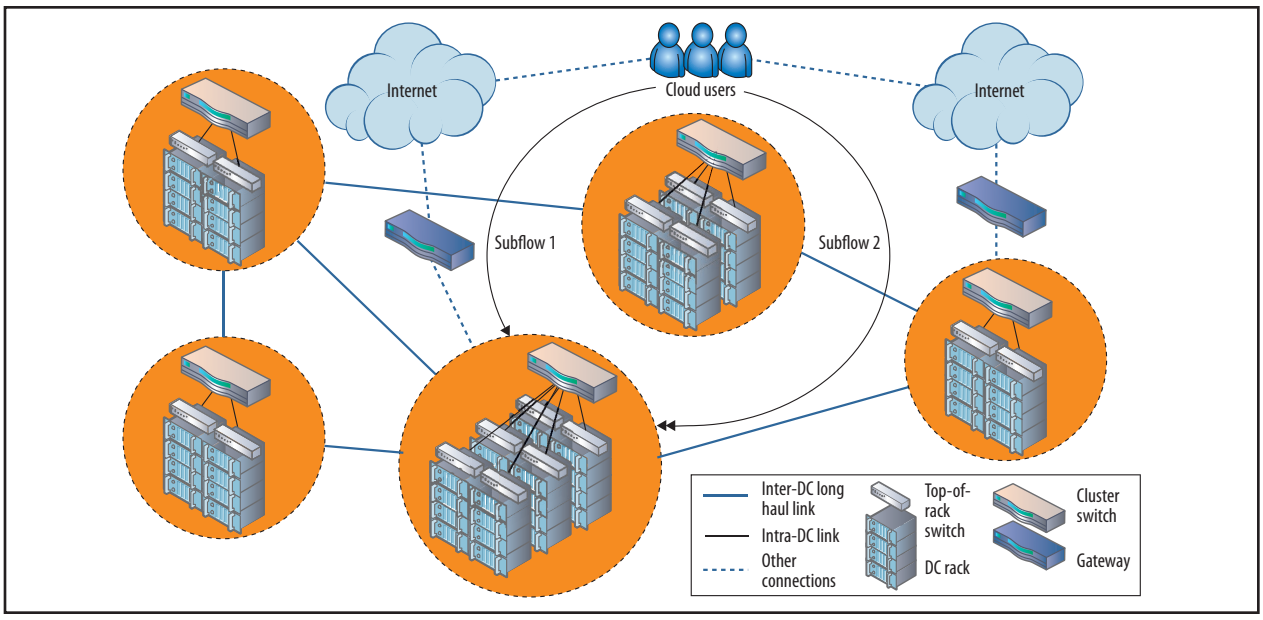

Figure 1. Sample network of geographically distributed datacenters (DCs). Distributing services over distant cloud sites is challenging in such networks.

vary according to network linkage and node states, and as a function of idle computing capacity. For these reasons, a cloud orchestrator must map multiple virtual networks on the physical infrastructure.[2] Moreover, virtual machine mobility—an increasingly important networking element—requires adequate support to guarantee seamless migrations and IP continuity.[3]

These and various other cloud datacenter characteristics impose novel requirements in designing cloud network overlay protocols.

CLOUD NETWORK OVERLAY PROTOCOLS

Optimizing datacenter operations across multiple sites calls for cloud network overlay protocols—network protocols that can isolate traffic and perform traffic routing for different IaaS networks below the application layer. The focus here is to specify encapsulation solutions for an advanced control plane that support virtual machine mobility and virtual network management. To date, several cloud network overlay protocols have been specified, and many datacenters and virtualization products already use them.

Protocols available at the Ethernet level include Shortest Path Bridging (SPB) and Transparent Interconnection of Lots of Links (TRILL); at the IP level is the Locator/Identifier Separation Protocol (LISP); and protocols on hybrid Ethernet-over-IP encapsulations include Virtual Extensible LAN (VXLAN), Network Virtualization Using Generic Routing Encapsulation (NVGRE), and Stateless Transport Tunneling (STT).

Table 1 briefly summarizes these six current cloud network overlay protocols. As the table suggests, they all natively support multipath forwarding and load balancing, at least to some extent. The most promising protocols, however, are also incrementally deployable, support virtual network segmentation, natively pass through IP networks, and easily cross firewalls.

AUGMENTING CLOUD ACCESS

Interfaces that provide diverse path management are key among cloud network offerings. Appropriately exploited, cloud network overlay features can be particularly useful for traffic routing over geographically distributed datacenters. For example, multipath communication capabilities brought to a user device using Multipath TCP—currently available for most mobile device operating systems—can significantly augment cloud access by splitting traffic over different connection subflows.[4] Cloud networks can provide the interfaces necessary for routing connection subflows over different paths.

As a practical matter, a geographically distributed cloud network needs, in total, three physical datacenter facilities—two located in one metropolitan area and a third located in another—to satisfy even those SLAs requiring very high reliability. Despite this, cloud service providers that offer advanced IaaS are increasingly adding to the number of sites where they make computing resources available—with roughly 10 percent of providers now having more than three datacenter sites. As the graph in Figure 2 shows, among cloud providers registered

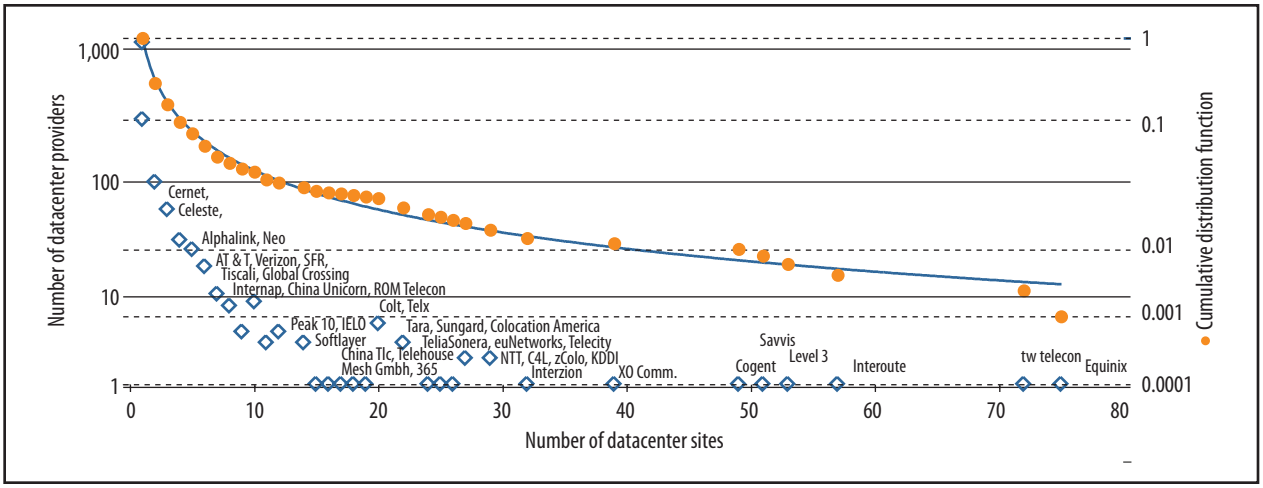

Figure 2. Distribution of number of cloud service providers as a function of the number of datacenter sites they offer. While having three sites is considered entirely adequate, roughly 10 percent of providers registered with Data Center Map (www.datacentermap.com) have four or more sites.

Table 1. Summary of cloud network overlay protocols.*

	Cloud network overlay protocol					
Cloud network overlay feature	**SPB**	**TRILL**	**LISP**	**VXLAN**	**NVGRE**	**STT**
Encapsulation	Ethernet over Ethernet	Ethernet over Ethernet	IP over IP	Ethernet over IP	Ethernet over IP	Ethernet over TCP/IP
Inter-datacenter link	Ethernet	Ethernet	IP	IP	IP	IP
Intra-datacenter link	Ethernet	Ethernet	IP	IP	IP	IP
User device integration	None	None	Yes	None	None	None
Virtual network segmentation	Yes	Limited	Yes	Yes	Yes	Yes
Firewall friendliness	Very high	Very high	High	High	Low	Very low
Incremental deployability	Low	High	Very high	High	Low	Low
Multipath and load balancing	Native	Native	Native	Partial	Partial	Partial
Multicast	Native	Native	Ongoing	Native	Partial	Partial

*Shortest Path Bridging (SPB), Transparent Interconnection of Lots of Links (TRILL), Locator/Identifier Separator Protocol (LISP), Virtual Extensible LAN (VXLAN), Virtualization Using Generic Routing Encapsulation (NVGRE), and Stateless Transport Tunneling (STT).

with Data Center Map (www.datacentermap.com), 9 percent have between 4 and 22 sites, with 1 percent having 23 or more sites.

Apart from increasing capacity and reliability[5] and reducing energy and real estate overhead, greater geographic distribution among sites allows easier cloud connection for small- and medium-enterprise (SME) customers. A large number of SMEs and public organizations now migrating to the cloud require very low latency (below a few dozen ms) combined with very high reliability (less than an hour of downtime each year).

Moreover, cloud applications such as seamless remote desktop, online networked gaming, telemedicine, augmented reality, and mobile computation offloading[6] may require still higher cloud access performance in terms of jitter and bandwidth, and, more importantly, access latency and round-trip time. For example, efficient remote desktop requires access latency below at least 50 ms to get barely adequate quality of experience, a constraint almost impossible to satisfy for US-based users accessing EU-based datacenters, and vice versa. Online networked gaming pushes this basic

requirement below at least 30 ms. And, for adoption at any real scale, augmented reality applications such as those used by camera-equipped glasses requiring real-time image and/or voice recognition will almost certainly push cloud access latency down to no more than a few ms.

In addition, network function virtualization (NFV) for particular node types—firewall, deep-packet inspection, router, cellular base-station control, and set-top boxes, for example—which is currently under evaluation or in early adoption by many network providers, requires the minimum possible latency between access equipment and virtualization servers. Technological convergence among ISPs and geographically distributed datacenter providers is imminent; with NFV, potentially any network node function can be hosted as a virtual machine, then migrated and resized according to demand and network state.

PERVASIVE IAAS SERVICES

Geographic distribution of cloud facilities for improved performance and resilience has already resulted in the advanced IaaS services and cloud network overlay protocols we've outlined here. To support further advances, the next task for network operators and vendors is to increase cloud network decentralization—both expanding the distribution of cloud servers and "embedding" them into access and backhauling networks.

To this end, the concept gaining most momentum is the *cloudlet*.[7] A cloudlet can be defined as a "mini" datacenter within the access network such that the link between user and server is owned by a single network provider—that is, the same ISP or local cloud network facility—for particular purposes, such as hostile environments that require tactical solutions.[8] Cloudlets' major benefits are the ubiquity and pervasiveness they provide for computing services, as well as their access performance in terms of latency and reliability—the latter because the user–server network segment is relatively limited in length. Should a user move too far away from IaaS services, adaptive virtual machine migration across the access network is possible with adequate support from cloud network overlay and storage protocols.

Reliable, high-performance IaaS services can be delivered only by decentralizing cloud networks, which can increase the path diversity in cloud access, augment user's quality of experience via optimized multipath communications, and facilitate pervasive IaaS virtual machine deployment down to the mobile and access network levels. These trends stimulate novel server and network virtualization solutions—the former still in their nascency and especially challenging due to the inherent distribution of protocols and nodes in a physical network of nodes.

The convergence of cloud and network technologies calls for close collaboration among cloud service providers, network providers, and industry associations as well as further cutting-edge research, education, and training in all these areas.

References

1. *National Survey on Datacenter Outages*, research report, Ponemon Institute, 30 Sept. 2010; www.ponemon.org/library/national-survey-on-data-center-outages.
2. M.F. Bari, "Datacenter Network Virtualization: A Survey," *IEEE Comm. Surveys & Tutorials*, vol. 15, no. 2, 2013, pp. 909–928.
3. P. Raad et al., "Achieving Sub-second Downtimes in Large-Scale Virtual Machine Migrations with LISP," *IEEE Trans. Network and Service Management*, vol. 11, no. 2, 2014, pp. 133–143.
4. M. Coudron et al., "Cross-Layer Cooperation to Boost Multipath TCP Performance in Cloud Networks," *Proc. 2nd IEEE Int'l Conf. Cloud Networking* (CloudNet 13), 2013, pp. 11–13.
5. R. de Souza Couto et al., "Network Design Requirements for Disaster Resilience in IaaS Clouds," *IEEE Comm. Magazine*, vol. 52, no. 10, 2014, forthcoming.
6. S. Abolfazli et al., "Cloud-Based Augmentation for Mobile Devices: Motivation, Taxonomies, and Open Challenges," *IEEE Comm. Surveys & Tutorials*, vol. 16, no. 1, 2014, pp. 337–368.
7. M. Satyanarayanan et al., "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, 2009, pp. 14–23.
8. M. Satyanarayanan et al., "The Role of Cloudlets in Hostile Environments," *IEEE Pervasive Computing*, vol. 12, no. 4, 2013, pp. 40–49.

Stefano Secci *is an associate professor of computer networking at Université Pierre et Marie Curie (UPMC), and is vice-chair of the joint IEEE Communications Society/IEEE Internet Society Internet Technical Committee. Contact him at Stefano.Secci@upmc.fr or via his website, http://lip6.fr/Stefano.Secci.*

San Murugesan, *Cloud Cover column editor, is director of BRITE Professional Services, adjunct professor at the University of Western Sydney, Australia, and editor in chief of* IT Professional. *He is also coeditor of the forthcoming* Encyclopedia of Cloud Computing. *Contact him at san@computer.org or via his website, http://tinyurl.com/sanbio.*