www.pipelinepub.com
Volume 19, Issue 5

A New Network Edge Platform for CSPs

By: Marc LeClerc

Along with the rapid growth of connected devices and applications, traffic patterns are becoming more fluid and harder to predict. In conventional architectures, network capabilities and applications are mostly isolated one from the other, with neither one knowing very much of what the other is doing or how its behavior might benefit or hinder the functions of the other.

To meet these needs, networking technology has come a long way in recent years, and communications service providers have had to rapidly adapt their network infrastructures, presenting them with both technical and business challenges that traditional network fabrics were never designed to meet. With the right strategies and solutions, CSPs can be well-positioned to succeed in the face of these challenges.

Service insertion and service composition

Both of these capabilities can be difficult to implement, as they require a deep understanding of the network and the services being used. This is made even more complex as many users access Internet and cloud-based applications from multiple different sources with expectations of a good user experience regardless of what network these services originate from. An economic concern for CSPs is how to transparently deploy services anywhere in the network where compute, storage, and latency requirements can be delivered most cost-effectively, or how to take advantage of underutilized resources as network usage patterns fluctuate in time or according to geographic factors such as time zones.

The need to elastically scale services

The fourth major challenge that communications service providers face is how to expand options for the monetization of new services and support new business models. With the rapid evolution of networking technology and data demand, service providers must find ways not only to deploy but also to monetize their services and support new business models to remain profitable. This can be difficult, as it requires access to network characteristics (such as bandwidth, response times, cybersecurity options, or multiple levels of customer service) and application capabilities (time or geographic awareness, access to messaging, and more). A good rule of thumb is that if it can be measured or monitored, it can be monetized. All in all, the recent network market and technological evolution have created major challenges for communications service providers, each presenting significant barriers to

successfully meeting and overcoming them. In response to these needs, a bewildering variety of new open technology solutions have been developed and are coming into play from all parts of the networking ecosystem. With all these options, a key question becomes where to start. What mix of hardware, software functionalities, and control mechanisms need to be provided? And because CSPs have significant investments in their current network infrastructure, can solutions be built that can integrate with and expand the in-use lifetime existing paid-for network assets?

The advantages of the network edge

From a hardware perspective, MEC platforms provide the flexible compute, storage and—critically—the networking capacity needed to deploy these functionalities as close as possible to service users, and to do so using economical, open, standards-based white box hardware. To make it possible to leverage virtual network functions (VNFs), virtual machines (VMs), and containers, the platform should have software that supports these standards and integrates with service providers' network orchestration software via standard application programing interfaces (APIs). A final consideration is the ability to elastically implement key network services such as load balancing, filtering, and packet brokering in software, along with the monitoring and control needed to change the behavior of these network capabilities in real time.

Optimization from anywhere in the network

This is where Segment Routing (SR) standards such as SRv6 come into play. SR makes it possible to easily redirect traffic transparently to users, and with a proxy function, even transparently to the service itself. For SR-unaware services, an SR service proxy would be needed for immediate deployment and to protect investments in the current service infrastructure. When implementing SR using SRv6, the benefits also include:

- Improved scalability: SRv6 allows for efficient scaling of networks, making them more resilient to traffic surges.
- Reduced latency: SRv6 reduces latency by making it easier to deploy services to the network edge and thus closer to users, as well as to route traffic network-wide more efficiently.
- Reduced costs: SRv6 reduces costs by making it possible to safely increase network resource utilization rates and reducing the need to overprovision for geographic or for redundancy purposes.
- Increased flexibility: SRv6 allows network operators to quickly and easily reconfigure networks in response to changing business needs.

services to the network edge. Another major benefit is the ability to transparently move services from private network functions (PNFs), to VNFs and cloud-native network functions (CNFs), and to support the transition to SR-aware applications while still supporting SR-unaware services.

Overall, CSPs are faced with many challenges to remain competitive in the marketplace and must work to ensure service visibility, optimize current network resources, and reduce their energy consumption and carbon footprints. With the right strategies and solutions, CSPs can be well-positioned to succeed in the face of these challenges.

Meeting the challenges

For the CSP's lab, we selected Lanner's HTCA-6600, an all-in-one MEC compute/storage/networking platform with NoviFlow's NoviFabric, NoviAnalytics, and NoviDashboard software that provides load balancing, service chaining, SRv6 service proxy, monitoring, and single-pane-of-glass management.

The programmability of the Intel Tofino makes it possible for NoviFlow's software to customize the data plane for various workloads, support new protocols, run networking programs on the switch fabric, and deliver detailed in-band network telemetry, creating real-time network visibility. NoviFlow's NoviFabric™ software provides key network capabilities directly in the Tofino-basedprogrammable network fabric including network-wide scalability for service chaining, SRv6 proxy, and load-balancing capabilities with visual analytics. This design provides lower latency and frees up CPU compute cycles for other processing chores and allows new services to be seamlessly inserted into the network by providing the SRv6 Proxy function.

Delivering services at the network edge

Taking a modular, open approach to building a MEC solution specifically architected for delivering services at the network edge is not only possible but with current technology and commercially available hardware and software it can be deployed now. And combined with SRv6 support, it becomes the vehicle to spread the agility and performance advantages of the network edge network-wide, delivering accelerated service insertion, more flexible service composition, elastic scaling of services to multiple terabits, savings in operating costs, better security, and faster problem resolution.