Network Programmability Pushes the Edge to Evolve

By:

The Service Provider Case

Another factor pushing for network evolution is the arrival of 5G networking. 5G places stringent demands on the forwarding plane and user plane, impacting both network infrastructure and the application space, and forcing a much tighter integration of networking, compute and storage resources. The massively multi-subscriber nature of carrier services further highlights the need for reducing latency, network flexibility and the ability to dynamically chain services. In carrier networks, edge characteristics are often seen as key limitations to the effective deployment of VMs, and the ability to scale services up and down in order to maximize price and performance. Additional requirements include the abilities to:

1. Cope with the communications needs of many different types of devices,
2. Ensure the network's ability to carry multiple different types of traffic with each having widely divergent characteristics, and
3. Efficiently deliver many different types of services with each imposing their own requirements on compute capacity, response time, volume of storage, and networking.

Due to the public nature of their offerings, service providers are even more concerned with the response to the ever-increasing threat of cybersecurity breaches and denial of service attacks. Traditional architectures were not designed to cope with such volumes and so are hard to scale. Rising complexity at the edge is thus becoming a limiting factor in protecting the network core, as well as increasing constraints on deployment and onboarding of new services. It also significantly augments traffic latencies and has become an impediment to service scalability.

When you look at the overall evolution of fixed and mobile networks, their user planes are essentially coming together. We have also seen a trend toward the disaggregation of network functions, where services are delivered on hardware that provides a software-agnostic platform instead of the traditional purpose-built appliances or edge devices.

On the left-hand side of Figure 1 (next page), you typically see radio-access devices, customer

premise equipment or multiple CPEs, as well as other IoT devices, which are connected via the aggregation side typically using lower capacity links between 1Gbps and 10 Gbps. The edge device thus becomes an aggregation point for these multiple feeds into the core network. Therefore, there are some key functionalities that need to be supported by the service provider edge devices, which either must act as virtual BNG or virtual broadband gateways or introduce some functionality like virtual EPC that is primarily software-defined.

Another need that we see at the service provider edge is for greater flexibility and improved programmability to support dynamic route flows in the platform. We also need to be able to bring compute resources directly to the edge, as we are doing some of the processing right at the service provider edge before passing the traffic on to the regional or core network infrastructure or data center.

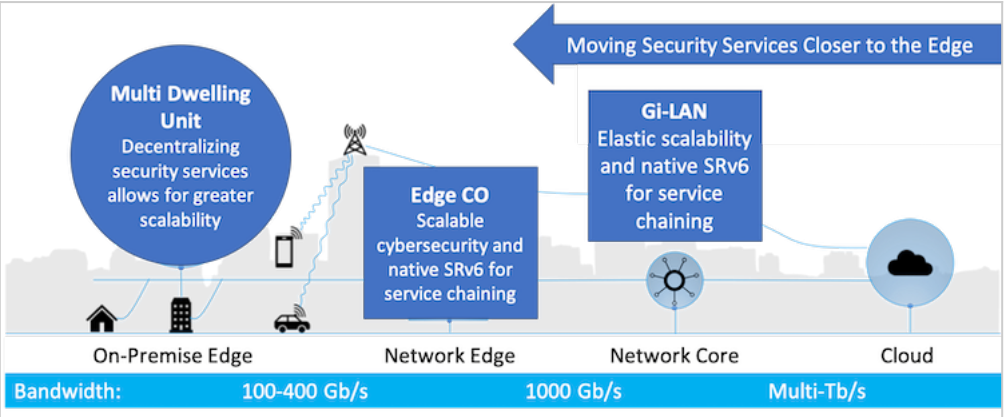

Figure 1: Moving security services closer to the edge
[click to enlarge]

Agility as a core platform design requirement

A network edge solution designed specifically to address the concerns listed above is needed. To meet these needs, this article proposes a fully integrated architecture that lowers service latency, improves security, provides an agile programmable networking platform, and supports service chaining at the edge to flexibly deploy multiple applications.

The following architecture describes such a solution bringing together hardware and software elements from NoviFlow, Lanner Electronics, and Fortinet. It fulfills the needs outlined using standard commercial off-the-shelf hardware to deliver an agile, customer defined infrastructure that can implement multiple virtual network functions on the same physical platform.

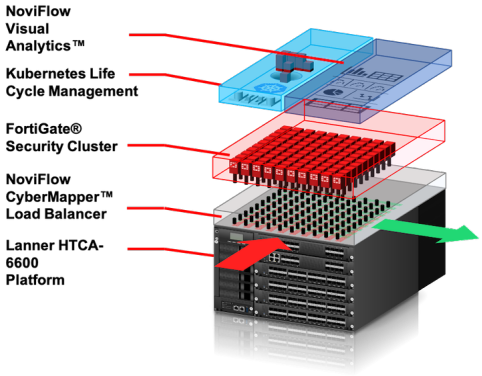

Platform management tools come pre-integrated to enable the remote management and provisioning of the platform without the need to go onsite to do cabling, wiring or troubleshooting. This removes the need for physical interactions with the edge or resources like the cost of a truck roll when problems occur.

The core idea of this solution is to fully leverage SDN infrastructure to implement directly in the network fabric key networking tasks such as filtering, load balancing and telemetry, and implementing other security functions as VMs, removing the need for expensive fixed-function appliances that physically constrain scalability. Delivering these capabilities on white box hardware makes for a more efficient, cost-effective solution.

In a service provider scenario, a multi-access edge solution usually needs to support a large number of lower-speed traffic ports. This solution brings with it a large number of physical networking ports, simplifying deployment. Each compute blade comes with its own internal NIC card, and the box integrates two 6.4 Tbps switches, eliminating the need for external cabling. The integrated switching makes it possible to do filtering decisions and redirection of traffic right at the network edge, before traffic even has a chance to consume core network resources or expose core resources to malware, viruses and cybersecurity exploits.

As the network fabric is managed via P4 and P4-runtime software control fabrics, the end result is open and programmable, meaning the entire platform can be repurposed via software to support multiple applications, multiple traffic handling functions, and multiple networking rolls with the same physical box.

Service chaining is also brought onboard to can pass traffic onto another dedicated compute node or even another hop at the network edge to do processing with the next available compute resources. Applications can reside wherever they can be most efficiently and economically deployed, and service providers are able to provide different compute resources at different locations based on changing needs at each location.

The implementation of SRv6 allows the pooling of services between the domains of a network. A server that has excess capacity can offer that capacity to another server that is running at full load elsewhere in the network. As a result, there is no longer a need to over-provision each access point; the cost of providing redundancy is reduced, as well as the cost of introducing new services, which often have initially very low traffic volumes.

Finally, the all-in-one package significantly simplifies the architecture of the edge, lowering costs and making the solution particularly advantageous to deploy in use case scenarios that are limited in space, power or remote locations.

Comparing the costs

There are big differences in price over a more traditional multi-box approach. The base software and hardware add up to an order of magnitude less than traditional solutions with equivalent compute, networking and storage capacity.

The NoviFlow/Lanner/Fortinet MEC Platform shown above delivers this functionality in a chassis that is 12U high, or roughly one quarter the footprint of the traditional solution that will take a full 44U high

rack to deploy.

Also, with this lower footprint come lower power consumption and air-conditioning costs. The end result is the same edge compute capacity, but at 33 percent of the total power costs. These are recurring savings on every power bill as long as the unit is in operation.

One final consideration: because all the elements of the unit are internally cabled, they can be managed entirely by software remotely, significantly reducing the complexity and cost of initial on-site installation, as well as severely reducing the need for truck rolls when new applications are deployed, when resources in the chassis need to be rebalanced, or when things go wrong.

Conclusion

According to NoviFlow President and CEO Dominique Jodoin: “What we see in the marketplace is that rapid evolution in network services and deployment models is driving the importance of a secure and programmable network edge for enterprises. But the economics to make this work commercially also require the ability to start small and in a modular fashion gracefully scale into the multi-Tbps.”

A unified platform provides significant performance advantages by integrating key network functions such as load balancing of VNFs, packet filtering, and flow redirection directly in the network fabric, eliminating the need for many physical appliances. By pre-integrating compute, storage and networking in a single chassis, the Lanner/NoviFlow/Fortinet MEC solution greatly reduces cabling cost and complexity.

In terms of installation, the all-in-one design is ready to deploy, with the SDN NOS, sophisticated analytics as well as key VNF enablers such as service chaining, and flexible networking that enable a wide variety of deployment scenarios, as well as eliminating the need for expensive truck rolls when reconfiguring or maintaining the device.

The built-in control software facilitates agile and transparent deployment of a wide variety of cybersecurity, vBNG and other VNFs from leading third-party vendors, supporting the most common use cases. Support for service chaining via SRv6 enables the pooling of resources and maximization of utilization of any service, anywhere in the network, at any time.

Finally, the flexible and programmable network fabric is software upgradable to support functions and protocols that don’t even exist yet, leveraging the powerful P4 language to extend the functionality of the fabric as customer needs and network usage evolves.

You can find out more about the NoviFlow/Lanner/Fortinet Network Edge solution at: https://noviflow.com/mecsolutions/