Case Study

Managers of the Infrastructure (Transport) and Services for Telco Cloud and Edge Computing Communication Service Providers (CoSP)

Deploying SRv6-based Service Programming to Optimize Edge Networks

Tier 1 Communication Service Providers (CoSP) Uses Programmable networking software and P4-based hardware platforms combined with SRv6 to deliver a new path to network-wide deployment of scalable, sustainable VNFs and CNFs

At a Glance

- Tier 1 North American CoSP uses programmability of a new generation of network processing hardware, P4-based software and the SRv6 protocol making possible evolved network solutions that ease the deployment of new network services and distribution of network infrastructure closer to the edge. This infrastructure enables efficient scaling of both existing infrastructure and new cloud capabilities.
- A growing ecosystem of equipment and solution vendors is already providing SRv6 network solutions for traffic engineering, L2/L3 VPN services and service programming (service insertion and service chaining).

Business Challenge

A tier-1 North American communications service provider (CoSP) has been undergoing a dramatic evolution of its network infrastructure to accommodate the move to 5G. The company's key goals are to lower latency, to simplify service programming and to enable decentralization of core network services. A major barrier to reaching these goals is that most mobile networks are based on a centralized model where base stations are backhauled to a few major central offices. This makes it particularly challenging to support the latency requirements specified in service level agreements for many 5G use cases.

To address service programmability, existing network solutions such as multi-protocol label switching (MPLS) could be used but would require that all transport and service nodes in the network be MPLS aware, limiting the vendors and network silicon that could be used. Fortunately, an alternative has become available: the adoption and deployment of segment routing over IPv6 (SRv6).

Redefining the Edge: Multi-access Edge Computing

Edge computing is gaining traction because it enables CoSPs to offer new ultra-reliable low latency communications (URLLC) services made possible by 5G. Rather than ship packets from the mobile phone all the way to a data center, it processes information closer to the origination point. As a result, CoSPs can better support lower latency applications and additionally reduce backbone network traffic.

To take advantage of these edge network features, network operators need to evolve their mobile edge to enable Segment Routing over IPv6 (SRv6). This is because SRv6 does not have the limitations of older MPLS networks in terms of numbers of segments supported. The greater flexibility of SRv6 not only enables traffic engineering but can also program multiple services within the network independently. Ripping out and replacing the existing equipment is not an option.

Delving Deeper into Segment Routing with SRv6

SRv6 is a next-generation internet protocol that leverages the source routing paradigm. It supports easy-to-use network programming and flexible, agile extensions. The protocol creates a network domain with predefined network segments that can be set up within an IPv6 network wherever segment-based traffic steering is desired. An SRv6 domain is composed of three types of nodes:

- The ingress head-end node encodes SRv6 header data into each packet.
- The transit node routes a packet based on the information in that added segment routing header (SRH). It shifts the header to the next ID in the segment.
- The egress end node removes the SRH and forwards the packet using the packet's original protocol.

SRv6 enables CoSPs to deploy networking services on top of new and existing infrastructure. It reduces the number of required protocol types; creates greater extensibility and programmability; and supports diversified new network services. In addition, the standard provides high levels of reliability and works with cloud services.

CoSPs also require strong security. Securing their networks is a challenge due to the complexity of integrating the access, core, and cloud network elements. DDoS and firewalls are typically a first line of defense, but next-generation solutions are needed to guard the edge.

Scalability problems need to be addressed as well. The decentralization of the network imposes the need to scale these services locally but deploying traditional load-balancers at every edge site isn't economical.

Wanting to explore the complexity and benefits of the SRv6 solution in its own network the tier 1 CoSP searched for companies with experience in programmable networks that could provide the services they needed at the edge, and meet the needs outlined above. They turned to Lanner Electronics Inc. and NoviFlow, both of which are Intel® Network Builders ecosystem members, to develop an SRv6-based multiaccess edge network (MEC) solution for service delivery.

The companies collaborated with Intel and other key industry players to assemble an all-in-one solution that could provide storage, programmable networking, and compute, one with the agility and performance required for today's complex, high-performance network edge.

Building an SRv6 Service Engine for the Edge

For the CoSP's lab, Lanner delivered the HTCA-6600, a scalable platform with multiple compute blades together with dual P4 programmable switching blades, that with software from NoviFlow provided load balancing and service chaining for security and other network services.

The turnkey system includes the following components:

The Lanner HTCA-6600 multi access edge computing (MEC) platform is a scalable, all-in-one, zero-touch configuration, software-defined MEC switch-server system. The HTCA-6600 delivers compute and storage resources along with network switching in a robust, secure, and scalable edge system. The platform features up to six dual-CPU blades supporting up to 12 Intel® Xeon® Scalable processor CPUs with a potential of up to 336 physical cores.

These CPUs are built for cloud-optimized edge networks. They deliver edge application, service, and control plane processing, high-performance packet processing and signal processing for both legacy and cloud workloads.

The HTCA-6600's networking blades make use of Intel® Tofino™ Intelligent Fabric Processors (IFPs) for Ethernet switching. Intel Tofino supports the P4 programming language, an open source programming language for defining packet forwarding in a wide range of networking systems. Intel Tofino IFPs were chosen for the edge because they can support even the most dynamic workloads. The programmability of the Intel Tofino series makes it possible 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 NoviFabric™ and NoviAnalytics™

On the software side, NoviFlow's NoviFabric software provides key network capabilities directly in the SRv6 Service Engine's Tofino-based programmable network fabric. This enables network-wide scalability for service chaining, SRv6 proxy, and load-balancing capabilities with visual analytics.

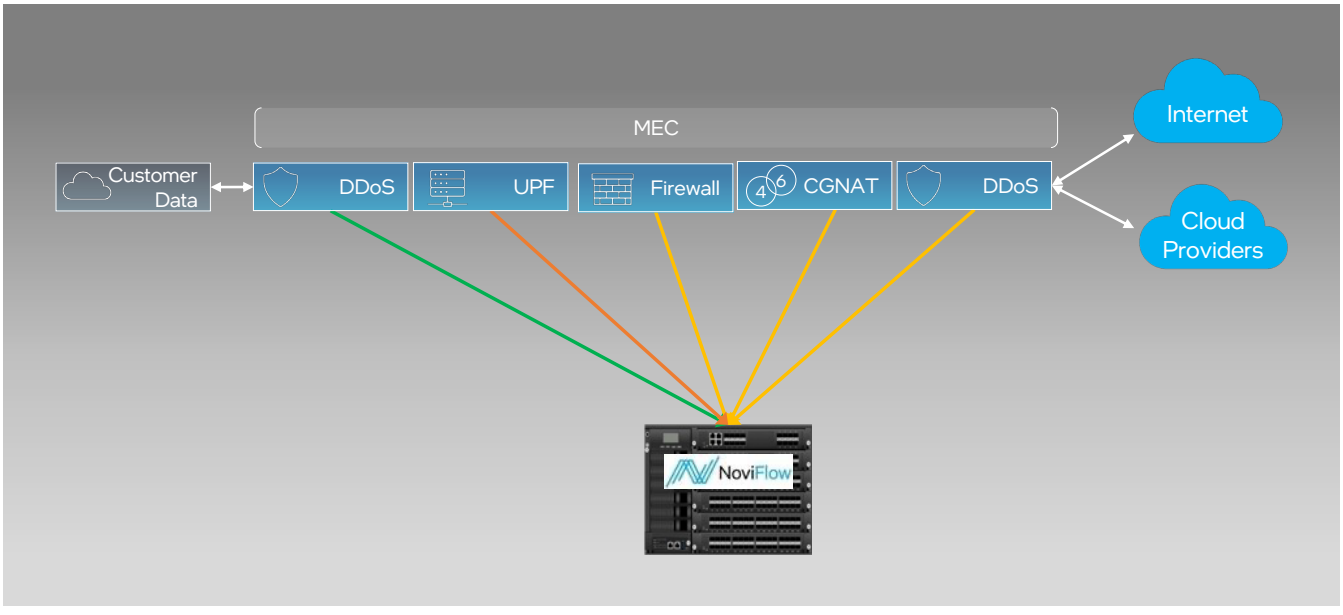

Figure 1. Block diagram showing the type of NoviFlow network functions that can be processed in the data plane of the HTCA-6600.

NoviFlow NoviFabric is a load balancing and service chaining application directly in the network forwarding plane. It allows the applications on the compute blades to scale across multiple virtual machines and across server blades for high performance. This design provides lower latency and frees up CPU compute cycles for other processing chores by transferring networking intensive tasks such as filtering and redirection to the Tofino network processors. The SRv6 Service Engine allows new services to be seamlessly inserted into the network by providing an SRv6 proxy function.

NoviFlow's NoviAnalytics software provides an end-to-end view of the HTCA-6600 platform for hardware and software ongoing performance. It collects information from the hardware, network, OS, and applications and produces a unified health view. All the collected information is stored in a time series database on the platform to enable high-fidelity troubleshooting. The "Green-to-go" visual interface simplifies the operations, administration, and management (OAM) of the HTCA-6600 platform, reduces onsite support costs, and implements preventative maintenance to minimize expensive truck rolls.

Results

The full Lanner/NoviFlow hardware and software solution was first deployed in a field trial at a Tier-1 CoSP in November 2021. The objectives of the trial were to demonstrate how SRv6 can be used to reduce service cost and latency at the edge, create a network that supports future services, scales on demand, and simplifies service deployment and network management.

A further benefit also became evident: because the turnkey system was delivered pre-integrated, the CoSP did not have to run cabling among its servers, load balancers, and switches. They were all connected in the backplane. In addition:

- The new solution seamlessly integrated with the existing network infrastructure.
- The carrier can set policies via software-based APIs, eliminating the need to re-wire and re-configure hardware whenever updates are needed.
- The solution automatically balances workloads, avoiding any potential bottlenecks and improving system performance.

The programmability of the Lanner/NoviFlow solution not only delivers an ideal platform on which to deploy SRv6 in the immediate term but also provides the tools needed to quickly evolve the network as business needs evolve.

Conclusion

Networks have become the foundation for modern business and personal communication. At the turn of the millennium, telecommunications networks underwent a paradigm shift in design, one that ushered in a raft of new services and changes in how business was conducted, and consumers were entertained. With the advent of 5G, virtualization, disaggregation, and cloud native services, there is a need for an evolved network infrastructure, this time focusing on bringing services closer to the end-users.

The move to edge computing and SRv6 promises such a transformation, one that reduces latencies, enables easier service insertion and elastic scalability, and that is built to evolve as CoSP's business evolves. The solution delivered by Lanner, NoviFlow, and Intel has clearly demonstrated the viability of SRv6 as a vehicle for network evolution. The companies helped one CoSP to be at the forefront of the SRv6 movement, and with this growing ecosystem of mature SRv6 offerings is well positioned to deliver new network services wherever and however they emerge.

Learn More

Lanner HTCA-6600

Intel® Xeon® Scalable processors

Intel® Tofino™

Intel® Network Builders

NoviFlow NoviFabric™

NoviFlow NoviAnalytics™

0223/LV/H09/PDF ♻Please Recycle 354429-001US