

SOLUTION BRIEF

Telecommunications Network Optimization

Accelerate SRv6 Processing



Technologies from HCL and Intel help optimize your network for modern service delivery



The Need for Network Optimization

Facing intense competition with each new industry player and disruptive business model, communications service providers (CoSPs) are seeking ways to differentiate themselves and enhance customer experiences in a fast-evolving telecommunication (telco) market—all while keeping costs under control. Exponential traffic growth and constant pressure to add more services and subscribers can tax legacy infrastructure and require that a CoSP have an optimized and simplified network. Many CoSPs have deployed network functions virtualization (NFV) in an effort to optimize their networks. However, due to the influx of new subscribers and growing data, many CPU cycles are consumed in the routing of traffic. This leaves fewer compute resources available for the actual containerized network functions (CNFs) and virtualized network functions (VNFs) that CoSPs want to support, resulting in suboptimal performance and the need for more hosts.

To help overcome these challenges, CoSPs are turning to technologies such as hardware acceleration and segment routing over IPv6 (SRv6). SRv6 helps address the requirements of NFV and software-defined networking (SDN) architecture. It provides a unified solution for networking programmability, service function chaining (SFC), protocols simplification, traffic engineering, and mobile and fixed network convergence. With segment routing, the number of protocols to be implemented in the network can be reduced, which can lead to lower operating expenses (OpEx). At the same time, network programmability is natively supported in the segment routing foundation, which can then seamlessly support NFV environments. SRv6 simplifies NFV implementation by allowing SDN, service chaining, and tunneling at the same time.

Overcoming Software Bottlenecks

Despite the gains provided by SRv6, bottlenecks remain in software-based approaches. A solution from Intel and HCL overcomes these bottlenecks and achieves up to 3x savings in cores through offloading low-level SRv6 processing to the Intel® Field-Programmable Gate Array Programmable Acceleration Card (Intel® FPGA PAC) N3000.^{1*} The card is reprogrammable and delivers the flexibility that CoSPs need to support new networking workloads. By taking advantage of the plugin-based framework of vector packet processing (VPP) and offloading CPU-intensive operations to the Intel FPGA PAC N3000, HCL has built an optimized architecture that enhances network throughput and predictability while reducing latency.

* Refer to <http://software.intel.com/en-us/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

The solution frees CPU cores and cycles so that cores once used for networking infrastructure can instead be dedicated to vital CNF workloads running on that infrastructure. The solution's small footprint can help reduce power and cooling costs, and it is available for both VNF-based environments through VPP support and CNF-based environments (and Kubernetes) through Contiv-VPP support.

The HCL solution based on the Intel FPGA PAC N3000 supports the SRv6 endpoint behaviors listed below, all of which enable SFC, L2VPN, and L3VPN:

- Static proxy (End.AS)
- Dynamic proxy (End.AD)
- Decapsulation and cross-connect (End.DX)
- Decapsulation and specific table lookup (End.DT)

By offloading CPU-intensive functions of segment routing to the Intel FPGA PAC N3000, the solution frees up CPU cores. As shown in Figure 2, four cores in a hardware-assisted solution can deliver comparable performance to 12 cores with software-based SRv6—a 3x savings in CPU cores.^{1,*}

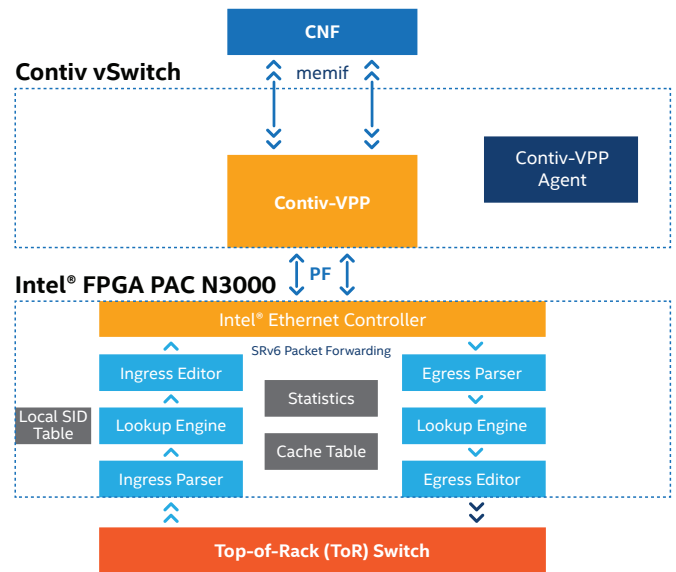


Figure1. Architecture of HCL SRv6 acceleration solution




Figure 2. Results of HCL performance testing based on a dynamic proxy (End.AD) use case; note that, with six cores, the Intel and HCL solution provides better throughput than 14 cores with a software-only solution

Key Features and Roadmap

The SRv6 acceleration solution currently supports the features and functionality shown in Table 1, with additional features on the near horizon.

* Refer to <http://software.intel.com/en-us/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

Table 1. Current and planned features and functionality for the SRv6 acceleration solution

Current Features	Roadmap
<ul style="list-style-type: none"> SRv6 endpoint behaviors offloaded for acceleration: <ul style="list-style-type: none"> Dynamic (AD) and static proxy (AS) support Decapsulation support: DX and DT SFC Kubernetes platform deployment through Contiv-VPP support Supported port configuration: 2 x 25 Gb over two quad small form-factor pluggables (QSFPs) or 1 x 4 x 25 Gb over a single QSFP QSFP28 compatible Workload scalability 	<ul style="list-style-type: none"> SRv6 endpoint for L2, virtual LAN (VLAN), and media access control (MAC) masquerading SRv6 transit for insertion and reduction, encapsulation and reduction, and applied and red L2 frames Scaling SRv6 local SID flow to 25K–1M Port mirroring: flow, VLAN, Generic Route Encapsulation (GRE), Multiprotocol Label Switching (MPLS) One-touch router: L2 switching and L3 routing, network address translation (NAT), Access Control Lists (ACL), overlays SRv6: GRE, MPLS-GRE, NSH-GRE MPLS, and Virtual Extensible LAN (VXLAN)

An Optimized Network Is Closer Than You Think

The HCL and Intel SRv6 acceleration solution helps optimize CoSP networks for improved performance. SRv6 acceleration on the Intel FPGA PAC N3000 helps increase overall throughput, allowing for more predictable performance and freeing up CPU cores and cycles for other VNFs and workloads. The solution can help make more compute resources available for additional revenue-generating services and functions, and it can help CoSPs reduce their carbon footprint while enabling cloud SFC of VNFs.

To see if hardware-based SRv6 acceleration is right for you, contact HCL at nfvi_acceleration@hcl.com

To learn more about the Intel FPGA PAC N3000, visit intel.com/pacn3000.

About HCL

HCL Technologies is a \$9.7 billion, next-generation global technology company that helps enterprises reimagine their businesses for the digital age. Our technology products, services, and engineering are built on four decades of innovation, with a world-renowned management philosophy, a strong culture of invention and risk-taking, and a relentless focus on customer relationships.

We offer an integrated portfolio of products, solutions, services, and IP through our Mode 1-2-3 strategy, built around Digital, IoT, Cloud, Automation, Cybersecurity, Analytics, Infrastructure Management and Engineering Services, amongst others. With a worldwide network of R&D, innovation labs and delivery centers, and 149,000+ 'Ideapreneurs' working in 45 countries, HCL serves leading enterprises across key industries, including 250 of the Fortune 500 and 650 of the Global 2000.

Contact for more details: NFVI_acceleration@hcl.com

Follow us on Twitter: [http://twitter.com/hclers](https://twitter.com/hclers)

Visit our Website: <http://www.hcltech.com/engineering-services/>

