Lanner

Introducing the SRv6 Service Engine

Using SRv6 and Programmable Networks to Reduce Service Costs at the Network Edge

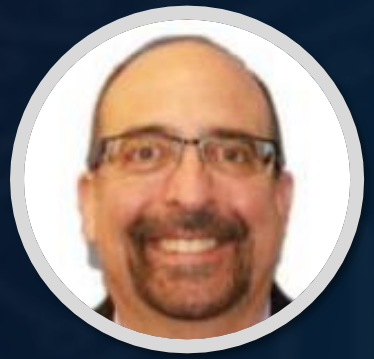

John DeMay
BXD Business Development

Sven Freudenfeld
CTO Telecom ABU

Jesper Eriksson
VP Product Management

July 12th 2022

What is SRv6?

- IETF Standard
- The SRv6 protocol creates a network domain with predefined segments:
  - A segment is an instruction for how a node will process the packet
  - The instructions are extensible
- A SRv6 domain is composed of three types of nodes:
  - The **headend** node classifies each packet, encapsulates the packet with an IPv6 header and encodes an ordered list of segments (instructions) into the packet using an IPv6 extension header, the Segment Routing Header (SRH). A list of segments is called a policy
  - The **transit** node routes or processes the packet based on the information in the SRH
  - The **egress** node removes the outer IPv6 header and the SRH and forwards the packet using the packet's original protocol
- SRv6 is used for traffic engineering, L2/L3 VPN services and service chaining
- SRv6 may be deployed as an overlay on an existing IPv6 network

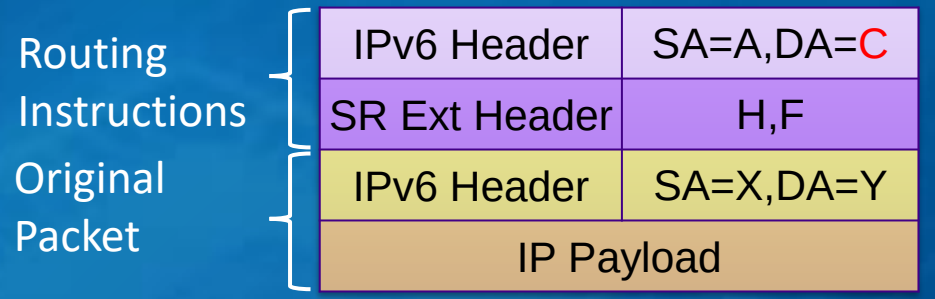

Why SRv6?

Reduces costs of delivering network services: Service Programming - chains of *network services are invoked by policy rather than hardwired into specific locations of the network*

Better scaling of the network: *Supports more nodes and services in a domain compared to a traditional approach*

Simplifies the network: *Reduces the number of protocols needed*

Easy to deploy: *Supported by most router vendors on existing platforms. May be deployed as an overlay on an existing IPv6 network*

Supports existing IP/MPLS network functionality: *IPv6 forwarding, traffic engineering, L2/L3 VPN services, HA*

Extendable to new network functionality: *SR Proxy, service chaining, 5G UPF*

NoviFlow Lanner

What do the carriers and vendors say?

"Segment Routing is fundamental for today's reality, which requires on-demand services as well as exponential bandwidth growth. Streamlining the IP protocols stack in order to provide a simplified service assurance support model for day-to-day is the key benefit."

Bell Canada's Technical Fellow Daniel Voyer

"SoftBank keeps focusing on improving service quality and enhancing the reliability and agility of networks while reducing costs…"

SoftBank Corp's Senior Vice President, Keiichi Makizono

SRv6 "dramatically simplifies" the routing architecture in comparison to legacy MPLS and VLAN designs and makes it easier to "automate and deploy services wherever you want to deploy them in the network.

EVP and GM of Cisco's Mass-Scale Infrastructure Group, Jonthan Davisdon

NoviFlow Lanner

• SRv6 Service Engine – Problems Solved at Bell*

Build for the future – evolve network to make it easier to deploy new network services
Traffic Engineering | Service Chaining | Monitoring & Monetization

Moving Infrastructure to the Edge for Low latency 5G Applications
Lower Cost | Vendor Independence | Better Customer Experience

Virtualize Infrastructure – Flexibly Insert services where they are needed in an instant
Dynamically scale | Software Driven| Deploy on Whiteboxes

Efficient scaling – Leverage existing infra or new cloud capabilities
SRv6 Proxy | Inline Network Services| Cloud Scaling

Application Acceleration – Accelerate NFV deployments with Intel® Technologies
SR-IOV Network Acceleration | Intel® Tofino™ IFP | General Purpose Processing

* As presented with Daniel Bernier, Bell Canada at MPLS World Congress

SRv6 Service Engine Components

NoviFabric SRv6 Service Chaining Proxy

- Software controller running on Intel® Xeon® servers
- Manages network services running on Intel Tofino:

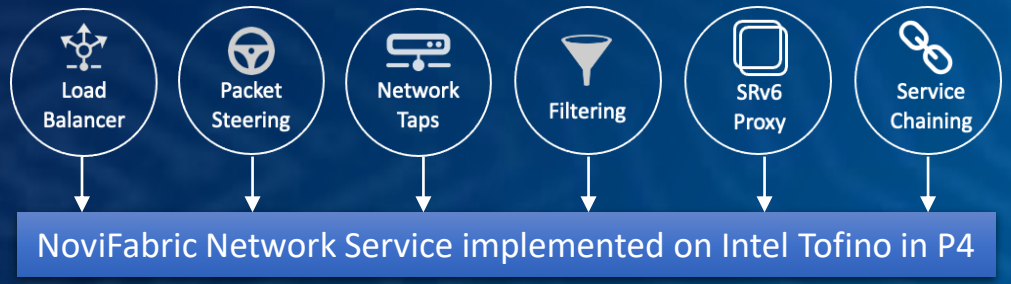

NoviAnalytics and NoviDashboard

- Measure performance of the applications running on the compute blades
- Collect information from the network, Hardware Platform Manager and syslogs to visualize health and status of platform

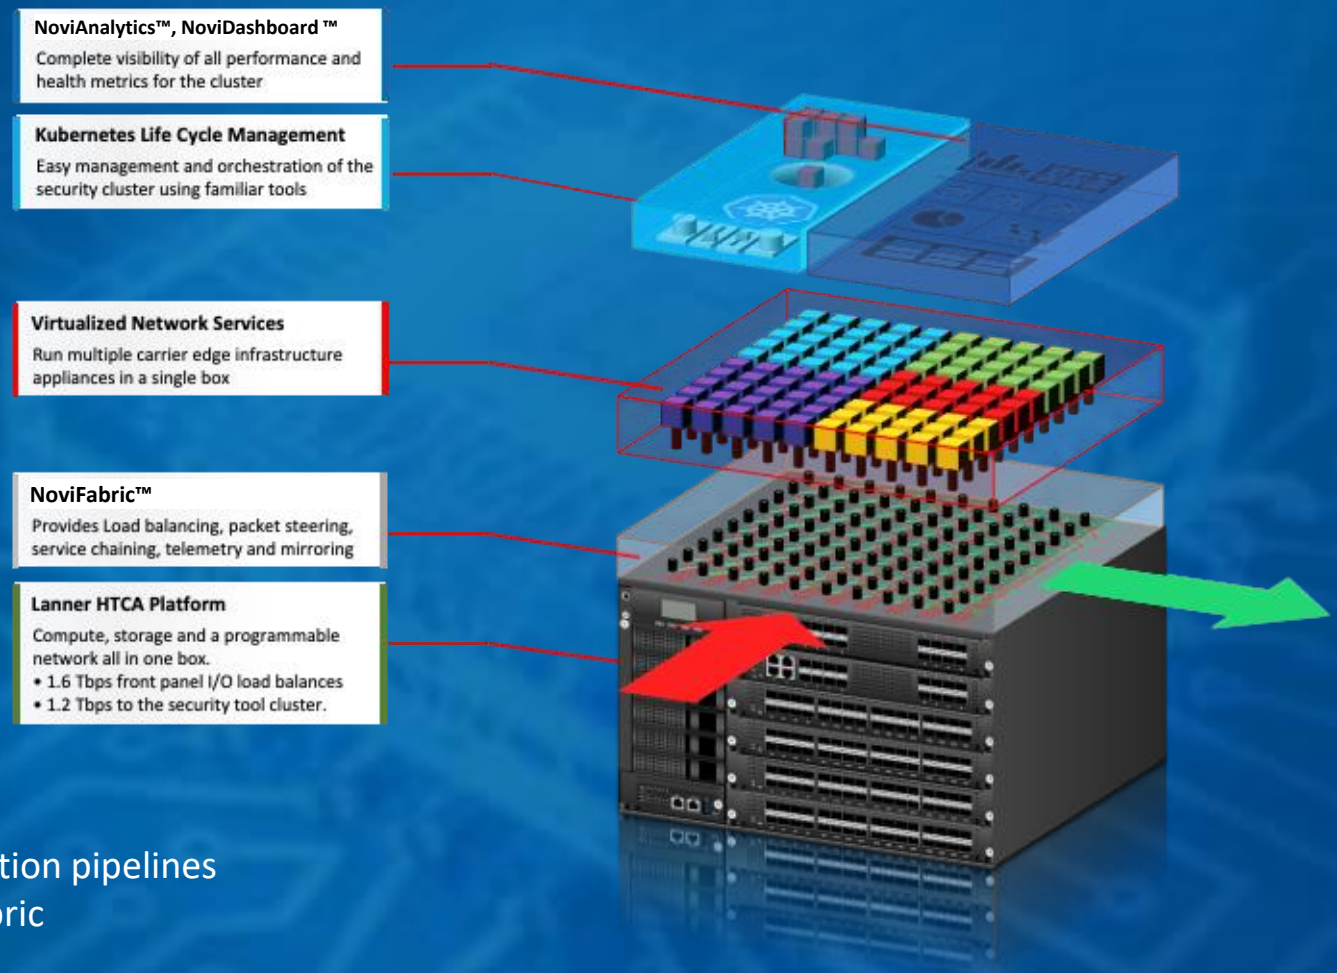

Lanner HTCA-6600 Platform with Intel Xeon and Tofino

- All-in-one Compute, Networking and Storage
- Carrier Grade NOS architected specifically to enable match-action pipelines
- Leverages multi-Terabps P4 Tofino programmable network fabric

NoviFlow Lanner

Key Take-Aways

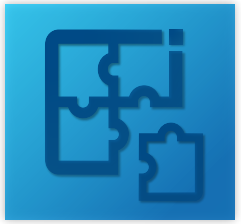

Networks are moving towards software

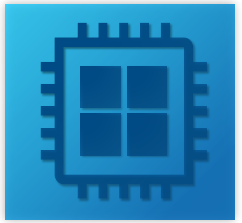

Edge is transforming and automating every industry

Intel Connectivity Group was purposefully created to deliver fully-programmable networks

Intel has the right vision, technology, to partner with customers and help build next-generation programmable networks

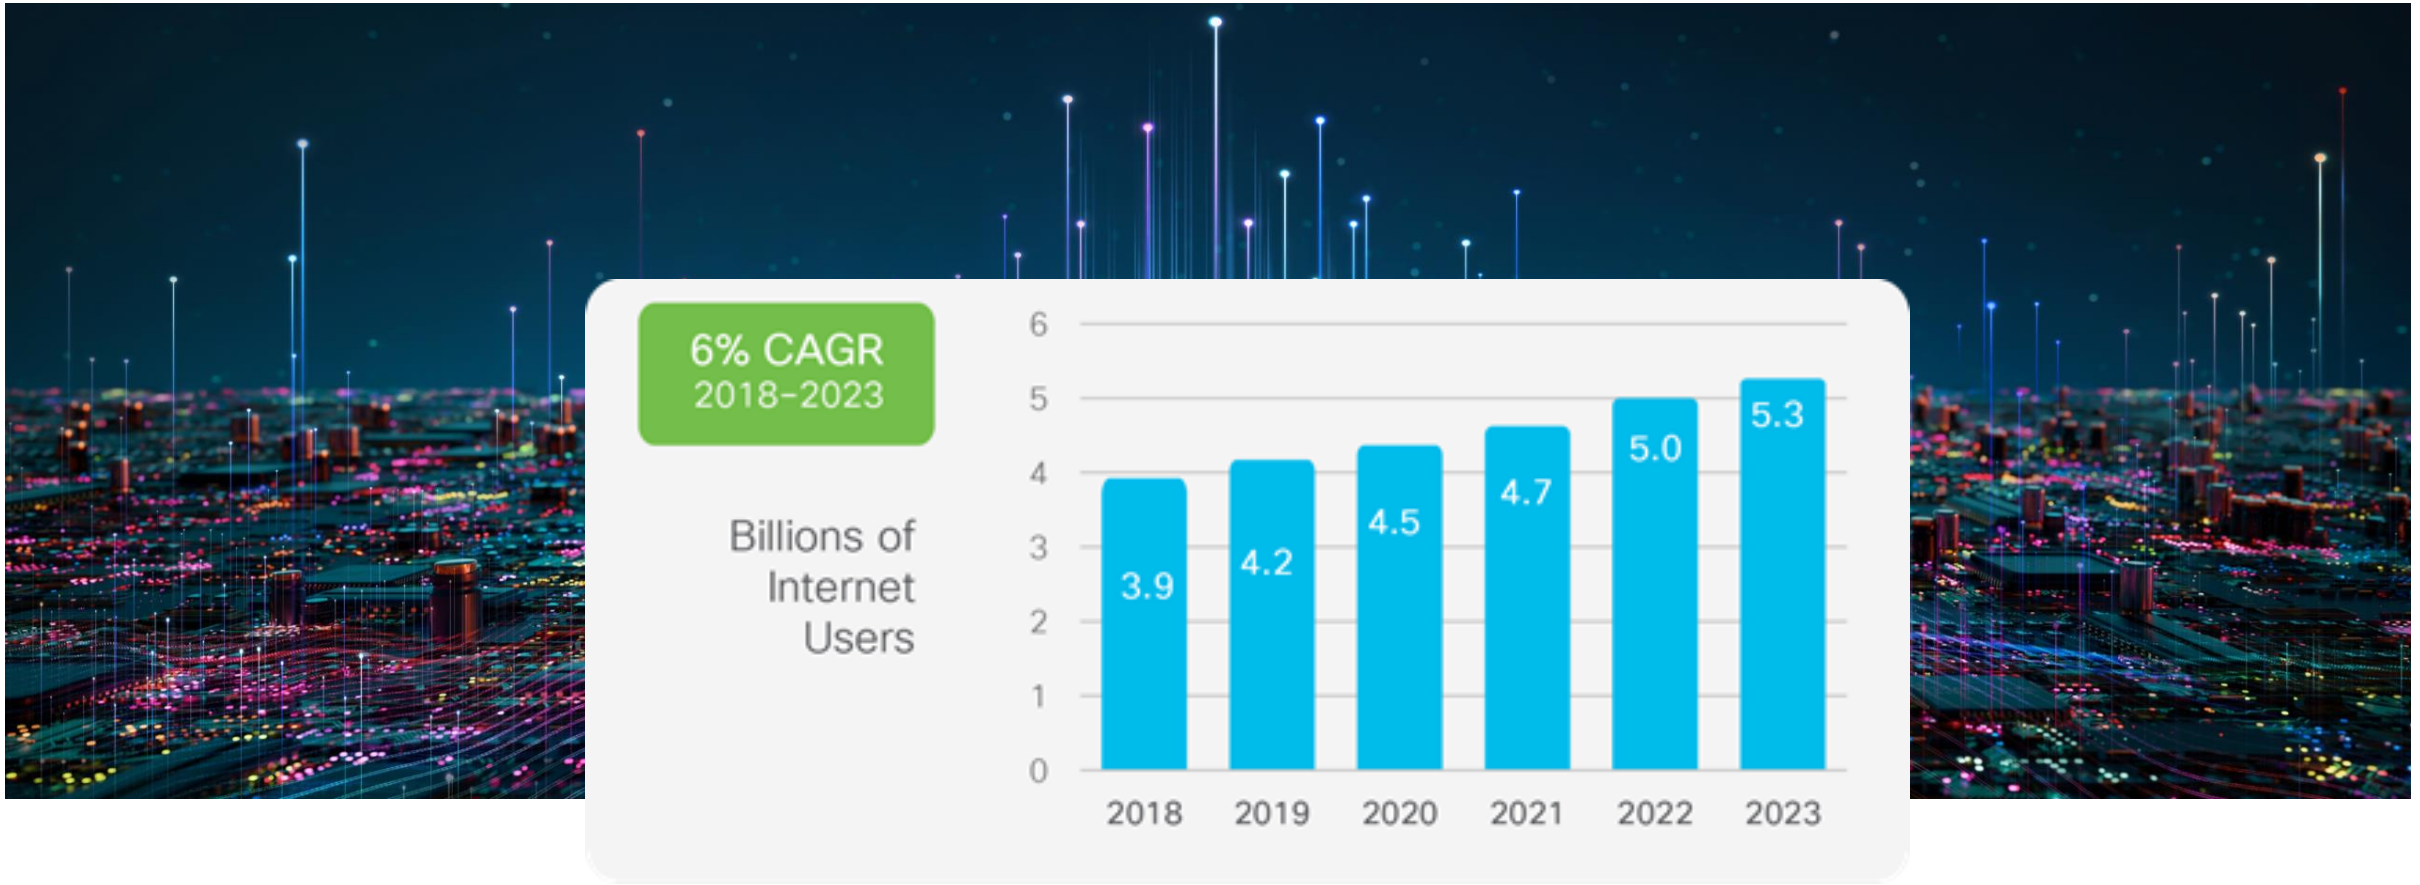

Global Internet user growth

Source: Cisco Annual Internet Report, 2018-2023

Seismic Shifts

Proliferation of Programmable Networks

+

Arrival of the Intelligent Edge

5G Networking Challenges

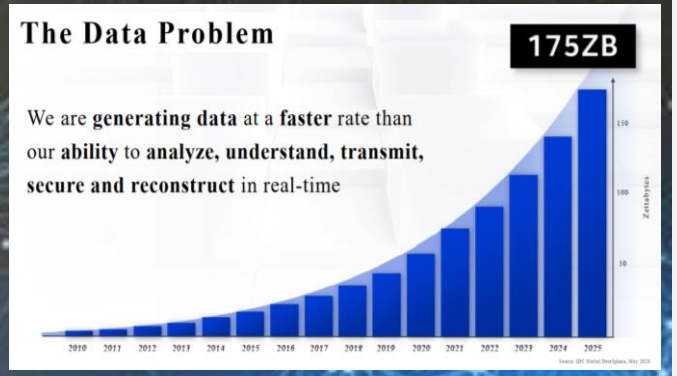

- CoSP's, CSP's are moving to Cloud-native architectures with container-based processing, orchestration and automation
- New services rollouts, complex network architectures & deployments and coverage
- Operational challenges of root causing network slowdowns
- End-to-end Security
- Increasing network CAPEX and OPEX investment for service providers

For 5G services, networks need to get smarter while increasing bandwidth

VISION: Intel® Intelligent Fabric

End-to-End Co-Optimizations

Photonics Integration

Resilience, Optimization, Adaptability

OPTICAL MODULES

High-bandwidth connectivity at 100G, 400G and beyond

ETHERNET SWITCH

P4-programmable scale-out fabric with uncompromising performance

ETHERNET IPUs and NETWORK ADAPTERS

Programmable infrastructure acceleration for demanding data movement

CPUs & xPUs

Fabric-enabled endpoints aligned to accelerators & software pipelines

RESILIENCY AND OPTMIZATION

Improved Density, Power and Cost

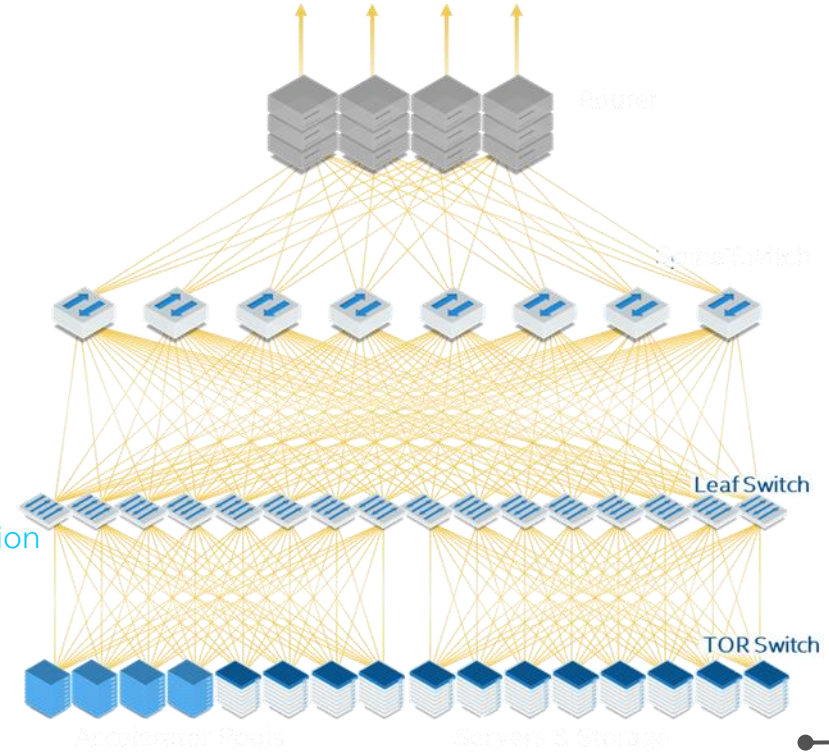

Intel Portfolio

intel. Si PHOTONICS

intel. TOFINO

intel. IPU

intel. ETHERNET

intel. XEON

intel. AGILEX

intel. eASIC

habana

Xe

Industry Standards

P4 ipdk Infrastructure Programmer Development Kit SONiC SPDK eBPF DPDK

Intel® Intelligent Fabric Key Benefit Vectors

INTELLIGENCE

- Fully Customizable P4-Programmable Pipeline
- Intelligent Packet Processing for Accelerating AI/ML Workloads
- Expandable Table and Buffer Sizes with Intel® FPGAs

PERFORMANCE

- 6.4/12.8/25.6 Tbps Total Throughput
- 112G/56G SerDes for High Speed and Easy Migration
- High-speed Intel® Silicon Photonics
- Power-optimized Hyperscaler Use Cases for Intel® Tofino™ Intelligent Fabric Processors

VISIBILITY & CONTROL

- Enhanced Congestion Control
- Identify Delays or Hotspots with Real-time In-band Network Telemetry (INT)
- Analyze Packet Flows with Intel® Deep Insight Network Analytics Software
- Increase INT Data Available with Intel® IPUs and Ethernet Network Adapters

XFG 2022 Switch Roadmap

P4 + ipdk Infrastructure Programmer Development Kit

25G SerDes	56G SerDes	56G & 112G SerDes
intel TOFINO	intel TOFINO 2	intel TOFINO 3
1.2, 1.8, 2.0, 3.2, 6.4 Tbps P4 Smart switch 100G MAC	4.8, 6.4, 8.0, 12.8 Tbps P4 Smart switch 400G MAC	6.4, 8.0, 12.8, 16.0, 25.6 Tbps P4 Smart switch 400G MAC Program-variable power

FPGA aside to Tofino (XL tables)

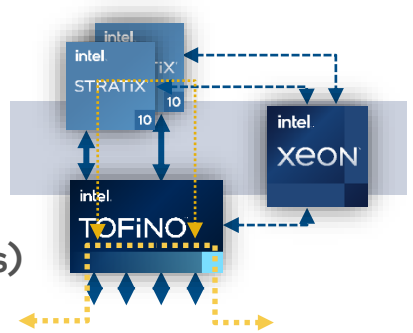

At any generation, complement Tofino with FPGAs, CPUs & IPUs to enable 10-100x increase in table and buffer capacity

- Telco Networks
- Cloud Gateway (LB, NAT, FW, VxLAN)
- Broadband Access and 5G gateway
- Carrier-grade security and NFV
- Network Packet Broker
- AI/ML acceleration
- Server Network Appliance

Programmability, Performance, and Power Efficiency at the Same Time

Intel® delivers full programmability without compromise on performance or power consumption when compared to fixed-function alternatives.

Intel® Intelligent Fabric Processors (IFPs):

Improved performance each generation

Higher performance

Improved power optimization each generation

Lower power consumption

For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.

Introduction to Segment Routing IPv6 (SRv6)

Defines packet processing in the network as a program in 128-bit segments

Forwards packets based on native IPv6 technology

Next-generation IP bearer protocol that combines Segment Routing (SR) and IPv6

Backed up many hardware implementations and open-source community

Why SRv6?

- SRv6 is SDN solution for Source Routing. Provides powerful programming ability, and flexible solution for traffic engineering

- Easy to deploy, as it uses IPv6 forwarding plane and does not require complex MPLS protocols or HW support for forwarding

Where do you insert services?

- Service Insertion via SRTE
- Service Path insertion via
  - Route Coloring
  - PCE Control
  - SDN Control (gRPC)

- Service Proxy can support
  - PNFs, VNFs, CNFs
- Service Proxy performs PSP and forward towards next SR node

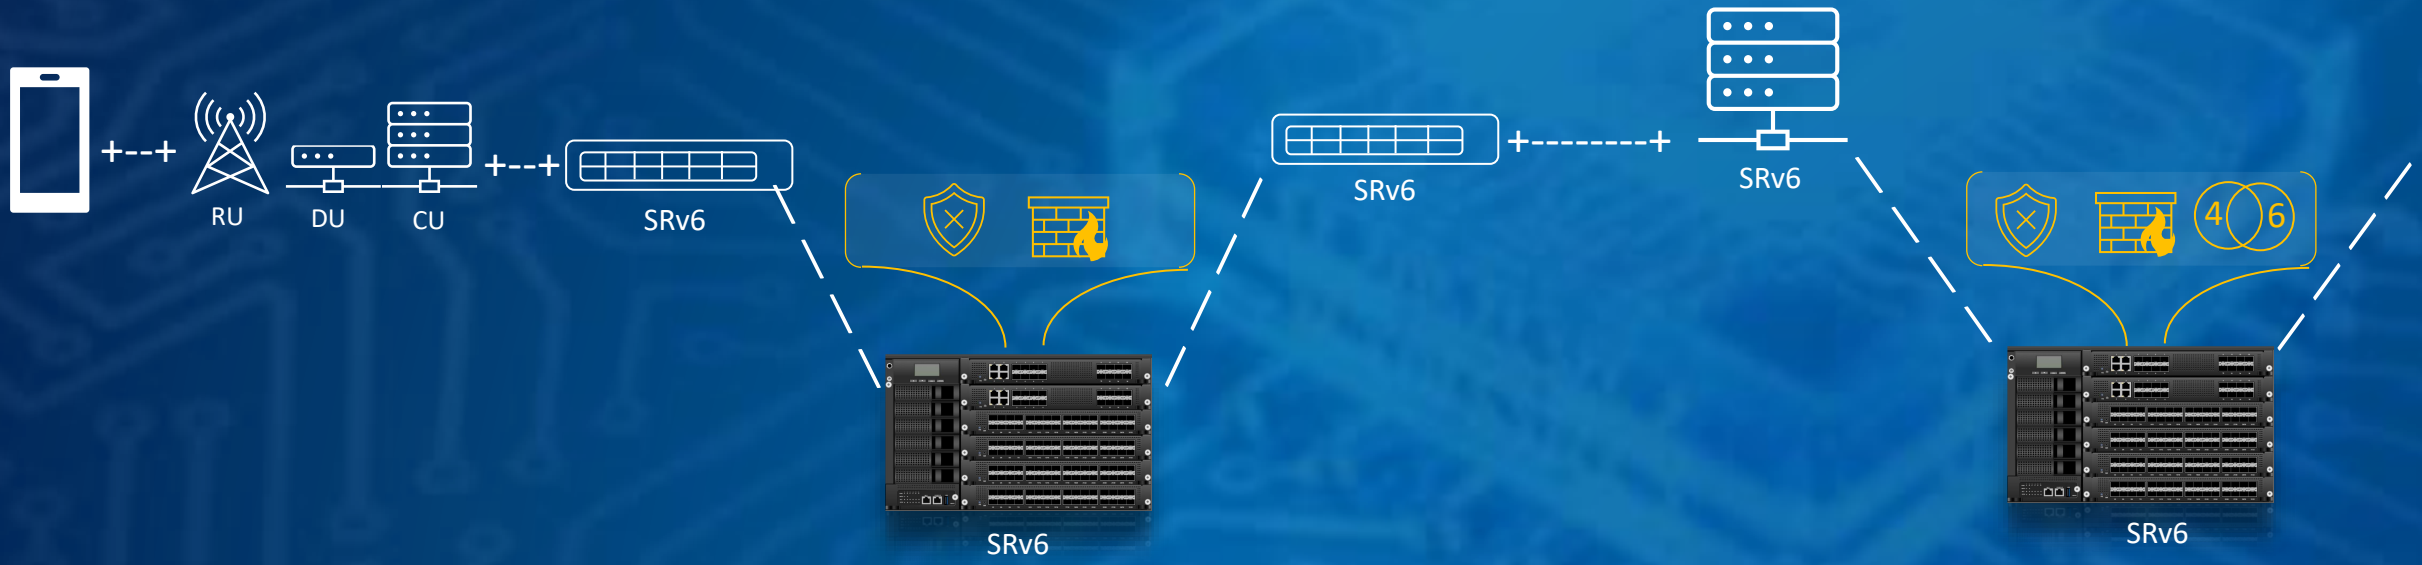

NoviFlow Lanner

Seamless Service Insertion via SRv6

It doesn't matter where the edge lives
Any Service! Anywhere!

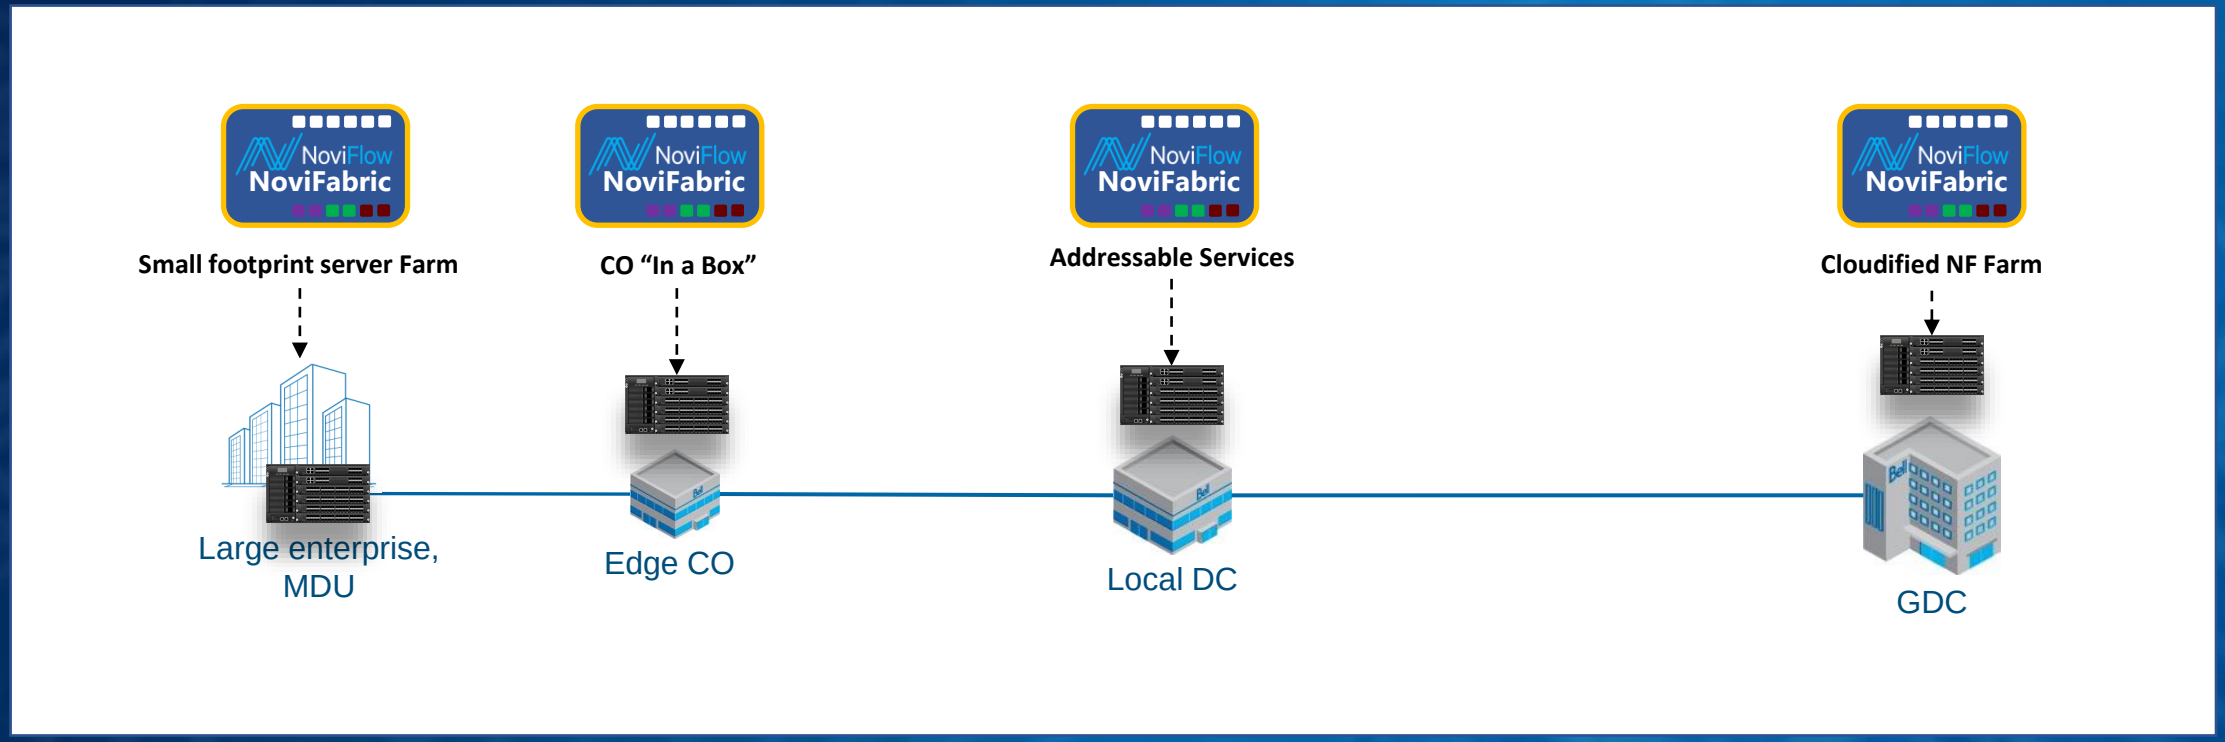

Can also be deployed in classic bump-in-the-wire mode

NoviFlow Lanner

NoviFabric – SRv6 Service Proxy

Reduce Service Complexity

- Acts as an orchestrator and enforcement point for a Service policy.
- Invokes, in order, the services listed in the service policy
- Support for a mix of different service types
- Each service can be HW or NFV
  - Integrated load balancing enabling service scalability
- **Shipping: GA since 2021!**

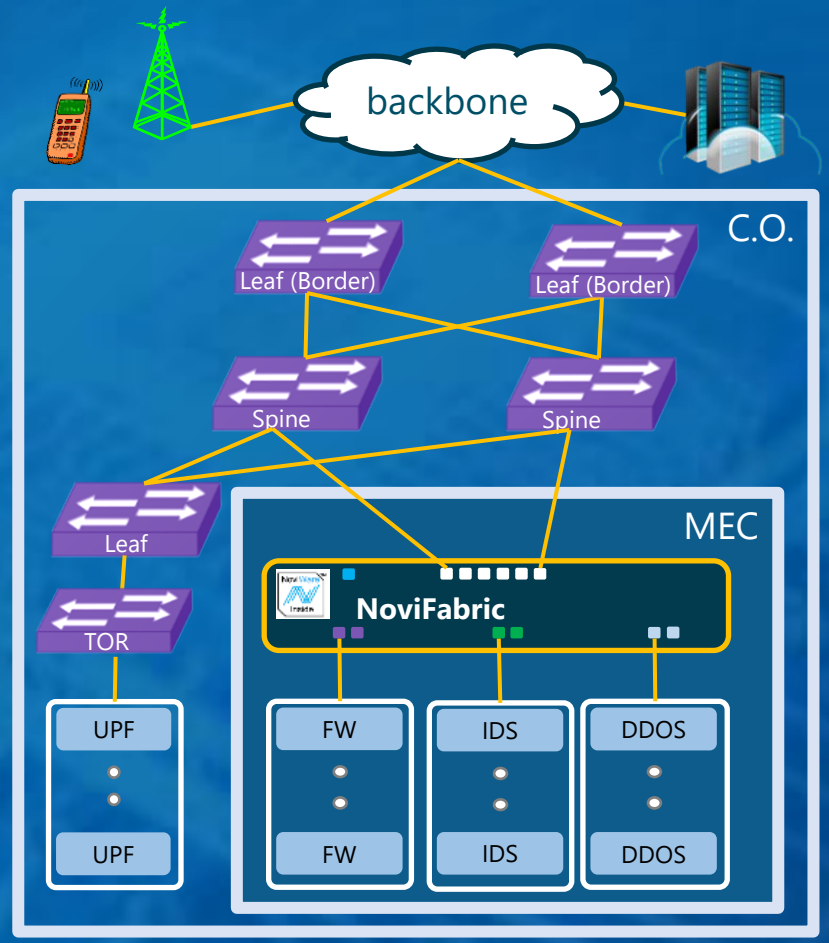

NoviFlow Lanner

NoviFabric - SRv6 Service Proxy

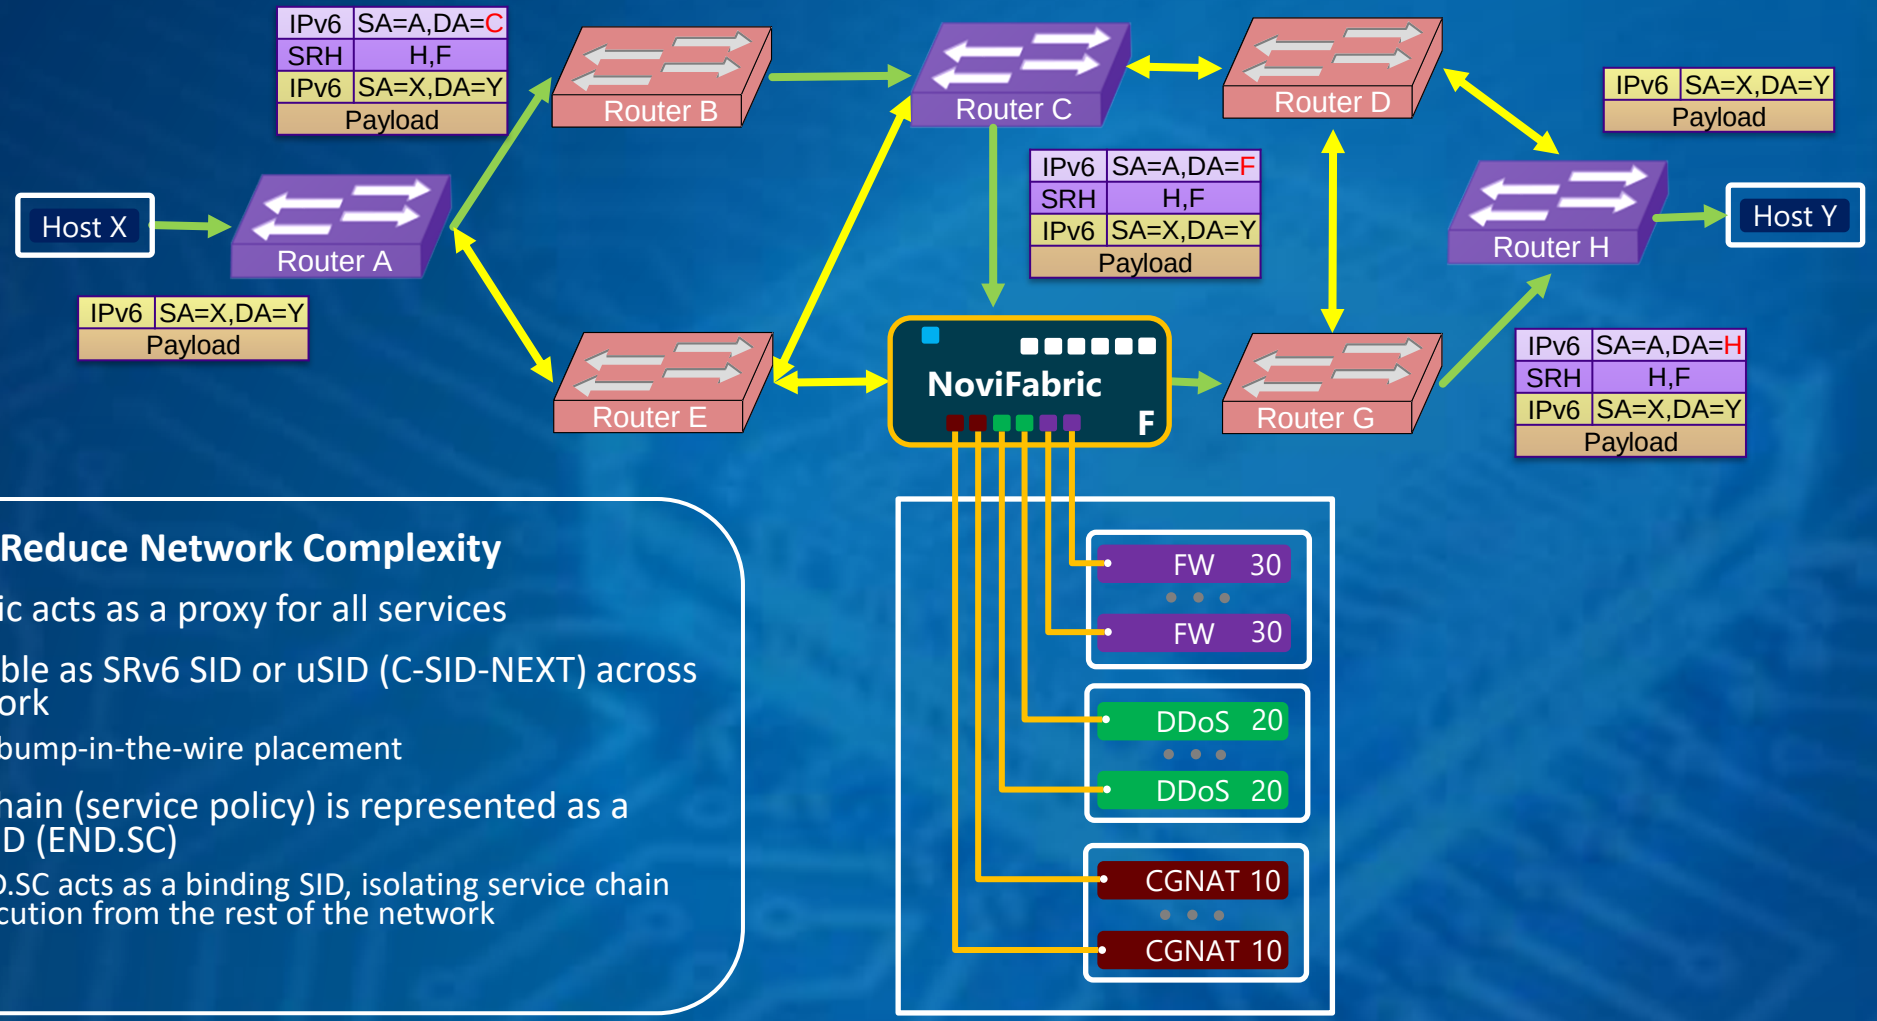

Reduce Network Complexity

- NoviFabric acts as a proxy for all services
- Addressable as SRv6 SID or uSID (C-SID-NEXT) across the network
  - No bump-in-the-wire placement
- Service chain (service policy) is represented as a unique SID (END.SC)
  - END.SC acts as a binding SID, isolating service chain execution from the rest of the network

NoviFlow Lanner

SRv6 Service Proxy - Service Types

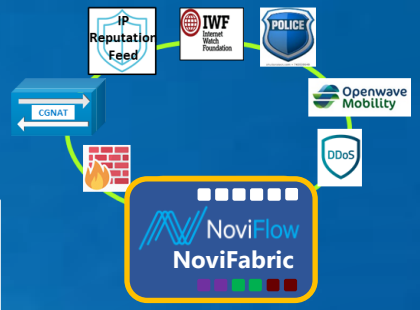

- Type I – **SR aware** – The service is a SR node (has its own SID).
  - The packet with the SRv6 header is transmitted to and returned by the service
  - The service executes the SR process on the SR header and its own service behavior on the encapsulated packet
- Type II – **SR passthrough** – The service is not an SR node
  - The packet with the SRv6 header is transmitted and returned by the service
  - The outer SRv6 header is left unchanged by the service
  - The service executes its own service behavior on the encapsulated packet
- Type III – **SR unaware** – The service is not an SR node
  - The SRv6 header is removed before the packet is transmitted to the service
  - Once the packet is returned, the SRv6 header is reinserted

NoviFlow Lanner

Carrier-grade NFVI Ready Platform

HTCA (Hybrid Telecommunications Computing Architecture)

- 2U to 6U appliance
- Integration of control plane and data plane
- High-speed switching
- NEBS ready, high availability design

NoviFlow Lanner

SRv6 Service Engine

NoviFlow software + Lanner HTCA 6600 Platform

Programmable

Goes beyond legacy silicon and legacy protocols

Carrier Edge needs programmable networks not based on legacy protocols

- Adaptable to support new features
- Software defined for automation

Intel® Tofino™ IFP

Better Visibility

Tofino provides programmable telemetry in silicon

Provide visibility into the performance of platform and VNFs

- In-band Network Telemetry (INT) for VNF performance
- Monitor the operational performance of the network

Scalable Architectures

Collapse multiple appliances on Tofino

Scalable network, compute and applications

- NoviFlow uses the network to scale across multiple virtual machines, blades or platforms

Coming Soon! Next Gen Intel® Xeon® Processors

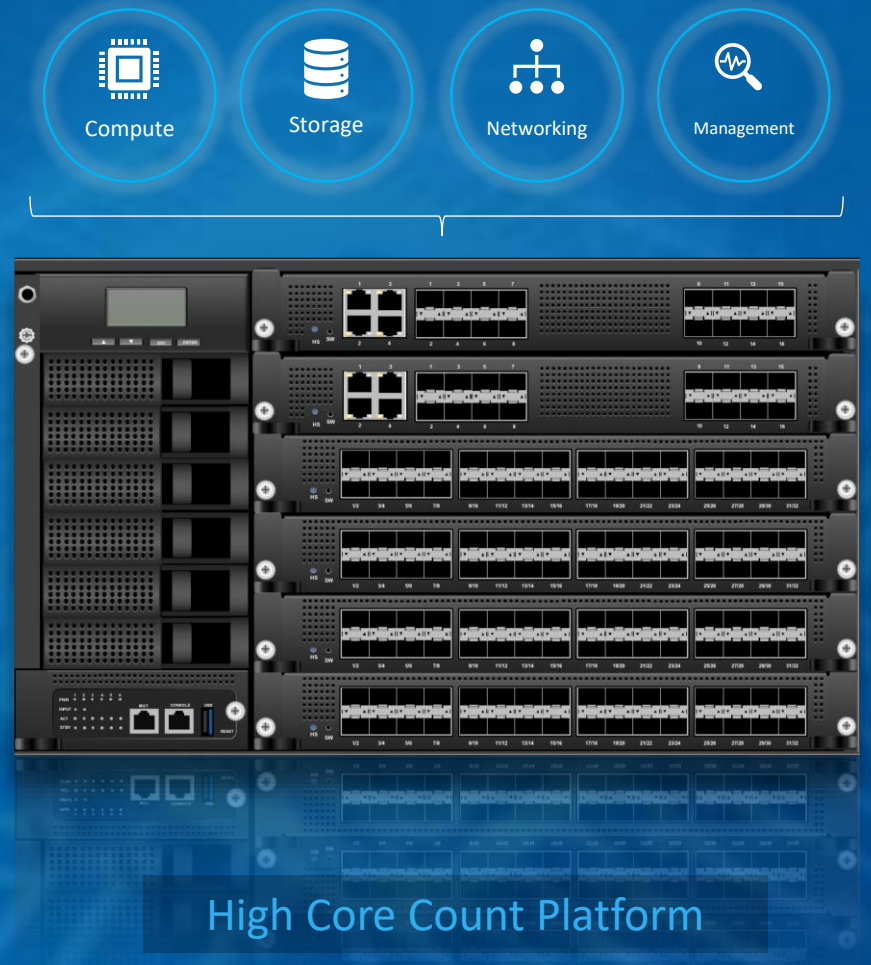

NoviFlow Lanner

SRv6 Service Engine

NoviFlow software + Lanner HTCA e400 Platform

Programmable

Goes beyond legacy silicon and legacy protocols

Carrier Edge needs programmable networks not based on legacy protocols

- Adaptable to support new features
- Software defined for automation

Intel® Tofino™ IFP

Better Visibility

Tofino provides programmable telemetry in silicon

Provide visibility into the performance of platform and VNFs

- In-band Network Telemetry (INT) for VNF performance
- Monitor the operational performance of the network

Scalable Architectures

Collapse multiple appliances on Tofino

Scalable network, compute and applications

- NoviFlow uses the network to scale across multiple virtual machines, blades or platforms

Next Gen Intel® Xeon® Processors

Highly Configurable Platform
PCIe Slot for Intel® IPU/FPGA Accelerators

NoviFlow Lanner

Hardware Offload Features – Direct NUMA Access, QuickAssist and SR-IOV: HTCA-6600

- The Intel Tofino IFP has dual dedicated 100G Ethernet to each NUMA node
  - One connection to each switch
  - Total of 400Gbps bandwidth to each server
- The HTCA Server Blades each have dual QuickAssist accelerators, one for each NUMA node
  - Uses Intel C627 for 100G offload for each QAT (200G per blade). QAT is embedded in the Platform Controller Hub (PCH)
- Using Intel e810 Ethernet adapters, SR-IOV provides network acceleration when used in a virtual environment.

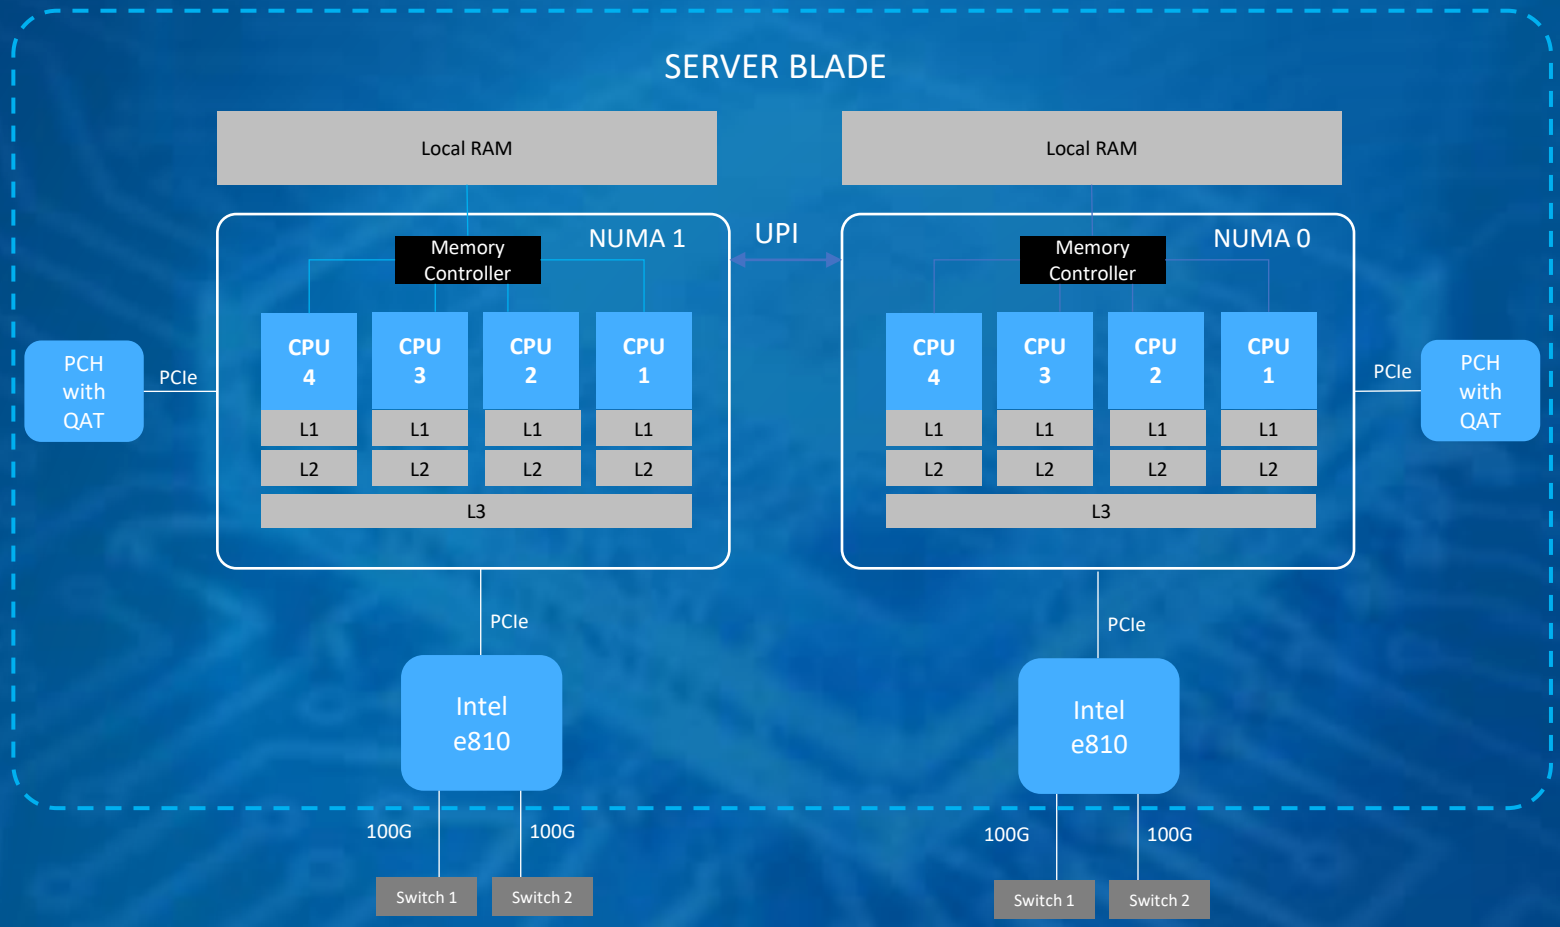

Cost Efficiency Comparison

Traditional Architecture

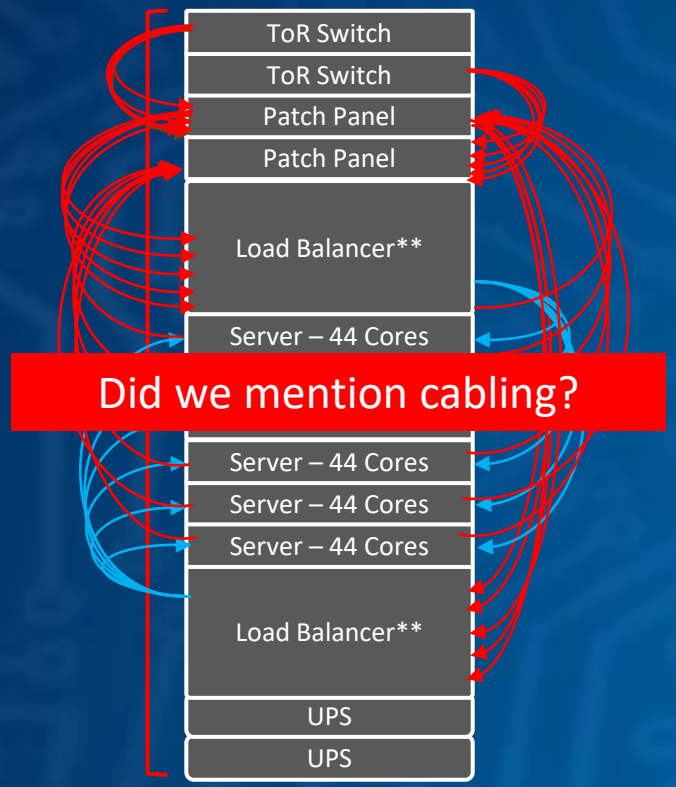

~$2.1 M USD*

Total Cost

NoviEdge MEC Platform

Lanner/NoviFlow

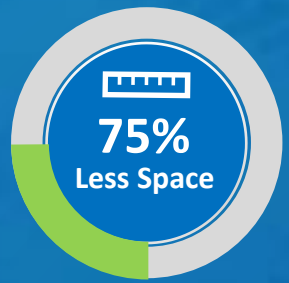

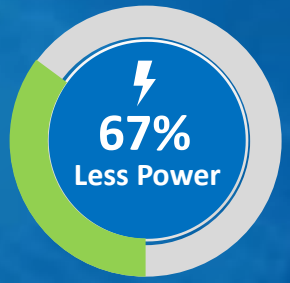

~$236 K USD*

Total Cost

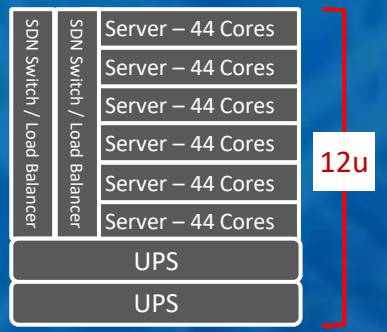

* List Prices ** “Sandwiched” configuration, single point of failure

NoviFlow Lanner

Scalable vRAN using HTCA-e400

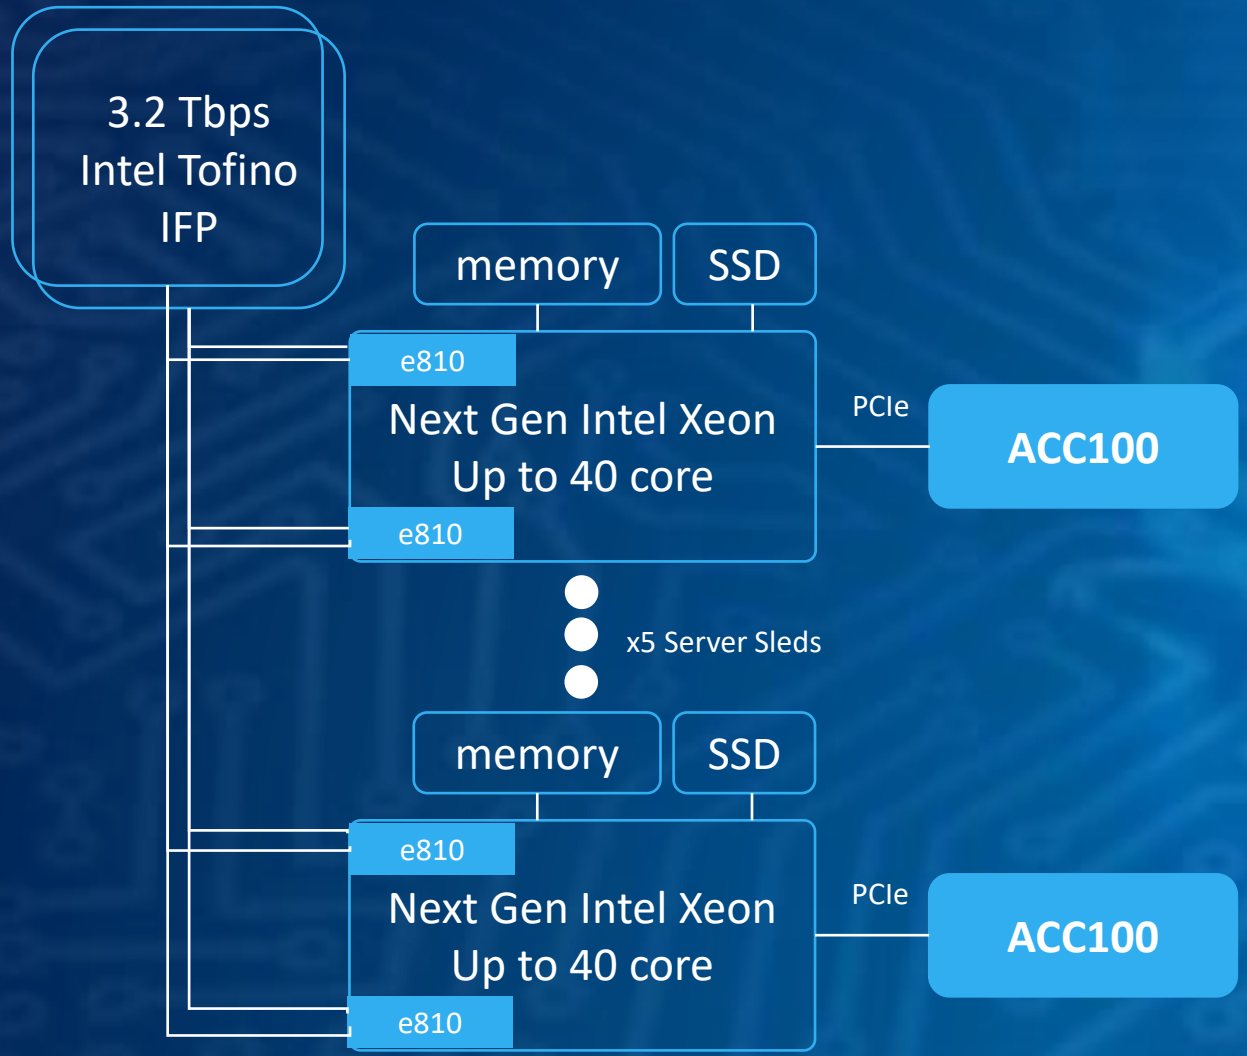

- Modular Design allows you to customize for vRAN acceleration
- SRv6 Service Engine allows you to loadbalance vRAN over 5 server sleds

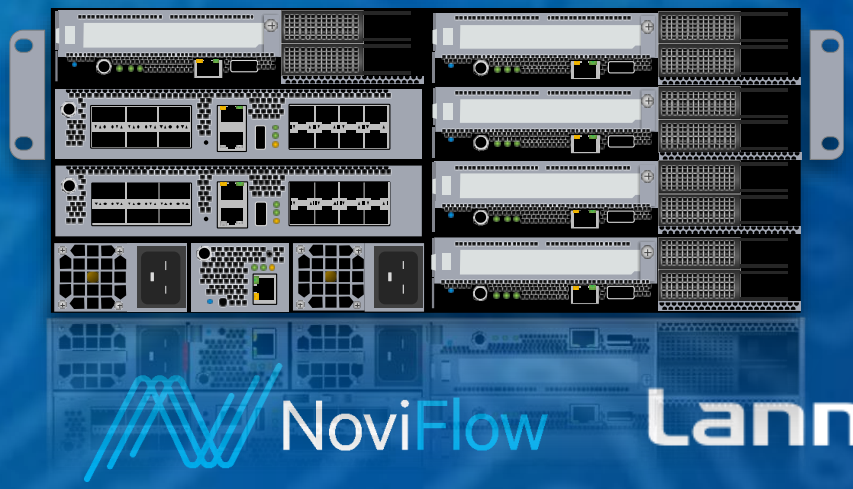

Security Proxy Use Case using HTCA-6600

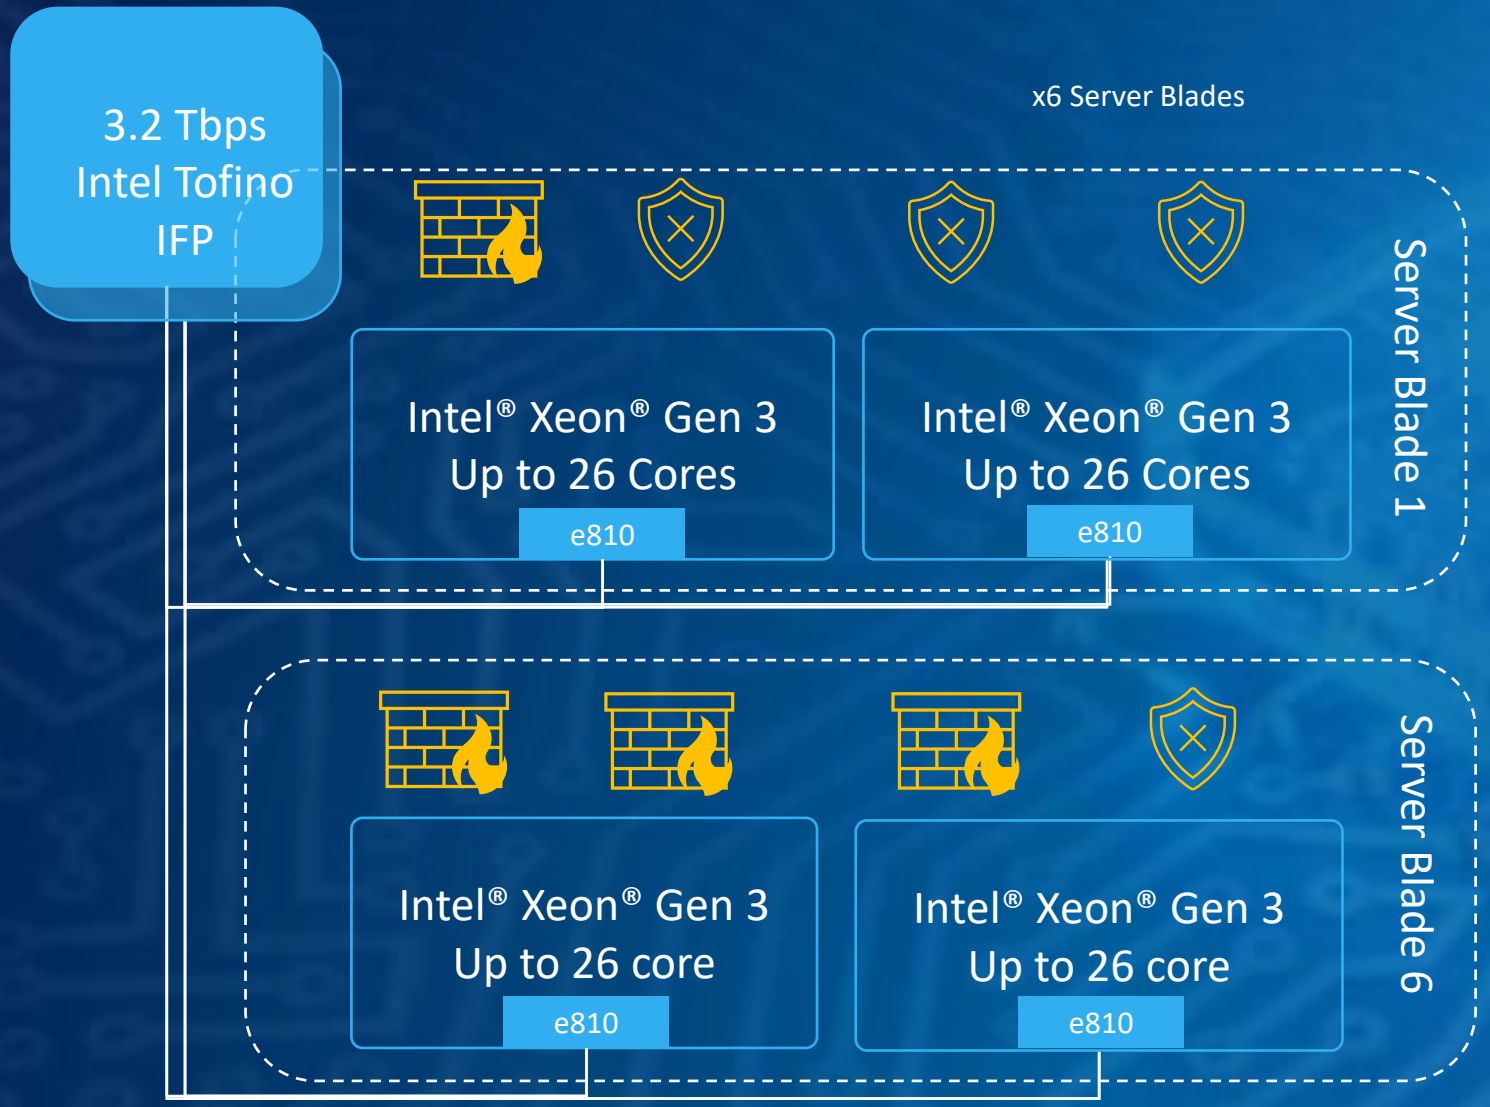

- Multiple virtualized services can be load-balanced, service chained and supports non-SRv6 aware applications using the SRv6 Service Engine

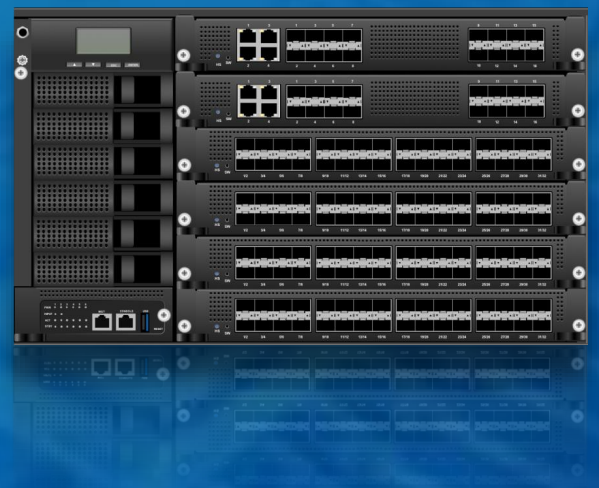

NoviFlow Lanner

Why SRv6 Service Engine Now?

- SRv6 Architecture and Network Programming make new solutions possible
  - Extensible Protocol Architecture to adapt to varying use cases
  - Massive Scale comes in various sizes enabling integration to various targets
  - Elastic scaling of services over entire networks

- Open, Programmable Dataplane (P4, etc.)
  - Easily implement new behavior from concept to line rate hardware implementation
  - Quickly create, fix, recode HW pipeline to adapt to demand
  - Deliver network services where they are needed, from central office to network edge to on-premise

- Great partner ecosystem (Intel, Lanner, NoviFlow)
  - Built using the industry's best processor, platform and software technologies
  - Deep knowledge of programmable networking at every layer of the solution stack
  - Fully integrated platform to deliver networking, compute and storage that is carrier-grade, easy to deploy and economical to operate

NoviFlow Lanner

The Future of Networks

With P4 and network intelligence customers can customize data flows so they can rapidly innovate, adapt, and differentiate next-generation workloads.

Learn more >

References

1. www.intel.com/IFP

2. Product brief: https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-3-brief.html

3. Intel P4 Studio: https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/p4-suite/p4-studio.html

4. Intel P4 Insights: https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/p4-suite/p4-insight.html

Notices and Disclaimers

- Intel technologies may require enabled hardware, software or service activation.
- No product or component can be absolutely secure.
- Your costs and results may vary.
- For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.
- Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.
- All product plans and roadmaps are subject to change without notice.
- Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.
- Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and SEC filings at www.intc.com.
- No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document, with the sole exception that code included in this document is licensed subject to the Zero-Clause BSD open source license (0BSD), https://opensource.org/licenses/0BSD
-

Thank You!

NoviFlow Lanner