Network Element Virtualization (NEV) Whitepaper

Version 1.7.1

Foreword

Xilinx, recently acquired and now a part of AMD, is the inventor of the FPGA, programmable SoCs, and now, the ACAP (Adaptive Compute Acceleration Platform). Our highly flexible, programmable silicon, enabled by a suite of advanced software and tools, drives rapid innovation across a wide span of industries and technologies - from consumer to cars to the cloud. For decades, Xilinx has been a trusted provider of programmable devices in wired and wireless networking equipment, spanning all corners of the network. Chances are that this very white paper passed through more than one Xilinx device on its way to your screen!

SD-WAN provides the policy-aware, secure routing for enhanced quality of experience. When used with Xilinx's adaptable devices, SD-WAN can deliver much faster performance and latency with innovative features. Over the next pages, Algoblu, an SD-WAN service provider, and Missing Link Electronics, a premier partner in Xilinx Alliance ecosystem, will show how they use FPGA technology to implement Network Element Virtualization (NEV) to accelerate SD-WAN adoption.

For more information on the many benefits of Xilinx products for wired and wireless communications, please visit: https://www.xilinx.com/applications/wired-wireless.html

Awanish Verma, Principal Architect and Director in the newly formed Data Center and Communications Group (DCCG) at AMD.

Problem definition

Driven by global digitalization, the daily operations of business users and the daily lives of individual consumers are increasingly dependent on the network. The personalized demand for networking services has become more and more intense. Today's carrier network is mainly designed to provide connectivity and best-effort delivery. It’s a big challenge for existing network infrastructure to support applications like online video conferencing, 4K/8K HDR streaming, streaming game, VR broadcasting and metaverse which require end to end SLA guaranteed to ensure user experiences.

The original intention of traditional Ethernet is to process non-real-time data on a shared medium that is simple to implement and low in price. It has very limited support for data services that have strict SLA requirements such as throughput, packet delay, and jitter. Although the IEEE Ethernet Working Group has since introduced a priority mechanism in the traditional Layer 2 network, added a Quality of Service (QoS) mechanism in the Layer 3 network, used VLANs and tunnels to isolate data flows, and performed port speed limits and queuing control through the token bucket algorithm to achieve control of sending and receiving, it has never been able to completely solve the problem of resource competition. This technical system essentially relies on software and hardware algorithms to achieve QoS guarantee at Layer 2 and above, which has relatively large limitations and cannot strictly guarantee SLA. The effectiveness of the algorithm, the setting of the buffer size, the packet length, the drop mechanism, the traffic model, the load and other factors will affect the final SLA delivery effect. The CIR (Committed Information Rate) promised by the operator is often hardly achieved on shared media. Delivered in accordance with CIR indicators, providing QoS guarantee based on traditional algorithms has great limitations.

In view of the fact that a pure software solution cannot completely solve the problem of resource contention under the existing technical architecture, Algoblu has taken another approach and proposed NEV, the underlying Network Element Virtualization technology based on hardware implementation.

Market Outlook

With the introduction of NEV technology, carriers can virtualize its underlying network and fine-tune the management of network resources to provide multi-tiered services with end-to-end service quality assurance to meet users' individual needs. It enables carriers to create new revenue streams by offering application and policy-oriented services rather than just sell bandwidth. In many fields such as streaming games, entertainment, streaming services, online video conference, SaaS services, IoT and Metaverse etc., users have a strong market demand for high-quality network services and are willing to pay for them. The largest potential market would be residential broadband.

The NEV solution can support multiple premium network services over users' existing broadband line, including 4K HDR or 8K UHD, streaming game, online video conference and metaverse which require end-to-end ultra-low latency, high performance, stable and reliable networking services to ensure the best user experience.

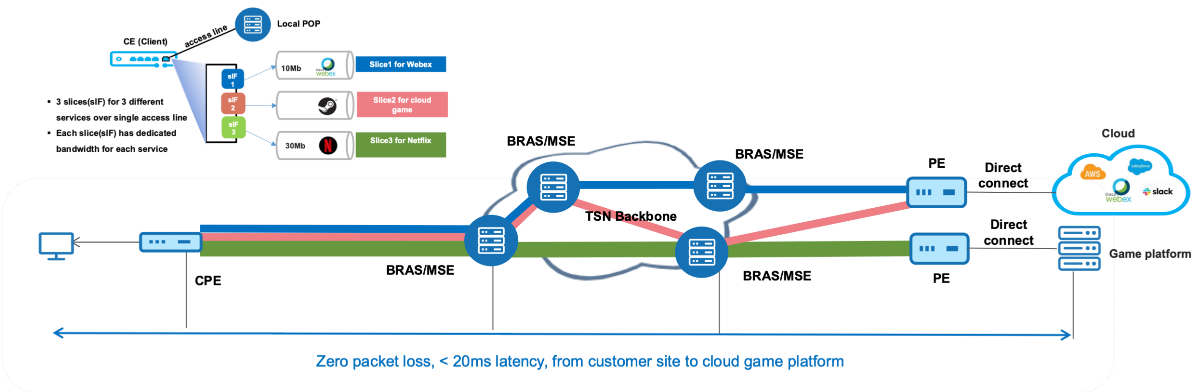

Based on OECD's statistics, the total broadband subscribers in 37 member countries are 453.9 millions by the end of 2020. Today, these broadband subscribers are charged on flat rate and there is a huge potential to charge premium by introducing application-oriented services. (Source: [OECD broadband statistics update - OECD)]

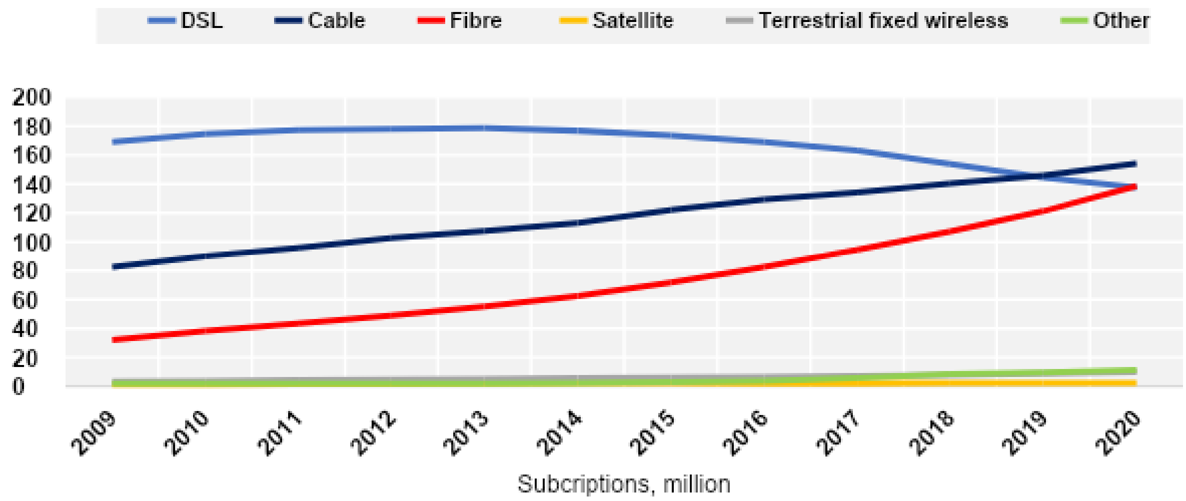

NEV (Network Element Virtualization)

NEV is Algoblu's patented technology for abstracting/virtualizing the underlying network resources, similar to VMware's concept of abstracting and virtualizing x86 server resources. NEV can virtualize the underlying network resources (e.g. fiber ports) into 100,000 individual atomic channels through self-developed FPGA-based chip. The scheduler and orchestrator allocate certain amount of atomic channels to form the sIF(Service Interface) per resource requirement. Algoblu's patented NEV network resource virtualization technology can slice a physical fiber into thousands virtual fibers, carrying specific services on each virtual fiber with strict service guarantees, and physically isolated from each other.

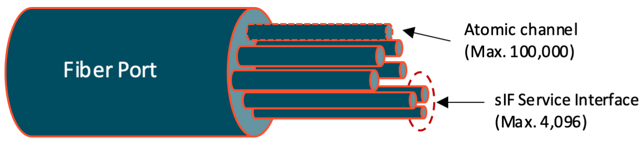

The NEV technology can enable carriers to virtualize and fine-grain management of their underlying network resources, and provide products and services with end-to-end service quality assurance for the individual needs of users. Carriers can provide multi-tiered application-oriented networking services and create new revenue streams over existing network infrastructures.

Within this document we describe NEV and present how NEV fits within the OSI network layer and how NEV's current implementation as a programmable Multi-Processor System-on-Chip (MPSoC) further increases performance, reduces costs over classical software-only SD-WAN white-box systems.

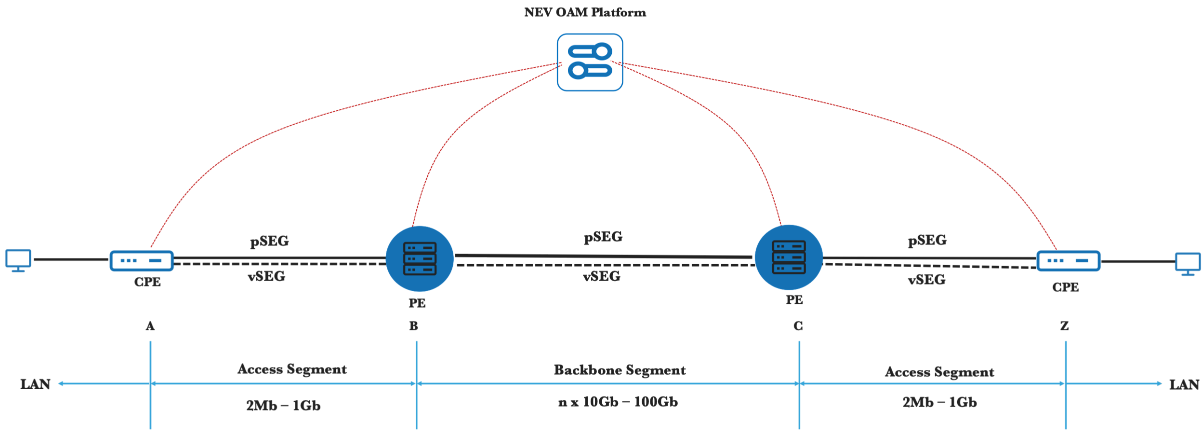

NEV fundamentally allows virtualization of networking elements and is similar, in concept, to now ubiquitous compute virtualization: Compute virtualization allows multiple so-called virtual guest computers running on one and the same physical host computer. Compute virtualization thereby provides more efficient usage of the underlying compute resources, CPU cores, main memory, mass storage, for example, while enabling dynamic resource allocation based on each virtualized guest computer's needs. NEV does the same for networking elements, cables, routers, switches, gateways, etc: Multiple virtual connections can be transported over the same physical connection while bandwidth and latency can dynamically be adjusted as needed.

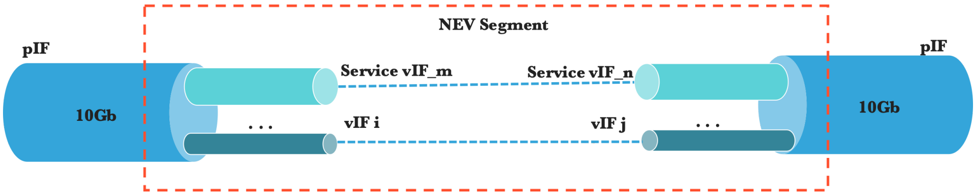

As a result, a physical network interface (pIF) can carry many so-called NEV Service Virtual Interfaces (sIF). Algoblu developed a centralized PTS(Primary Telco Services) OAM system to control and to manage this virtualization, including dynamic adjustments of bandwidth, on a network segment by network segment basis.

Built on Future-Proof Standards

NEV is compatible with existing standards such as IEEE 802.1Q and IEEE 802.3 Ethernet and is future proof by following new evolving standards such as for Time-Sensitive Networking. In the OSI Network Layer Model, NEV is operating at OSI Layer 2 and/or at OSI Layer 3. Key portions of the datapath processing are implemented in FPGA logic, such as Layer 2 and Layer 3 Encapsulation and Decapsulation or Traffic Shaping. Traffic Shaping not only reduces the burstiness of the traffic but also performs the virtualization and slicing into multiple Service Interfaces. This allows spanning an

SD-WAN by mixing OSI Layer 2 segments with OSI Layer 3 segments as this is needed when combining existing infrastructure with new 5G infrastructure.

Layer			**Protocol Data Unit (PDU)**	
Host Layers	7	Application	Data	
	6	Presentation		
	5	Session		
	4	Transport	Datagram	
Media Layers	3	Network	Packet	NEV L3
	2	Data Link	Frame	NEV L2
	1	Physical	Symbol	

Fine-Grained & Dynamic per-Application Bandwidth Allocation

For example, if a corporate SD-WAN briefly needs more (or less) bandwidth, maybe for a large virtual conference, then bandwidth can be increased (or reduced) on a per application basis, and within seconds. The telecommunication provider has the option through the NEV management software whether to adjust individually for each networking segment than spans the corporate SD-WAN or for the entire SD-WAN.

At the same time, the additional latency introduced by NEV is kept at the very minimum, and because of a digital circuit-level implementation latency is very predictable / deterministic. This is made possible by integrating into the NEV chip certain functions used in Time Sensitive Networking (TSN): Cyclic queueing (IEEE 802.1Qch), scheduled traffic (IEEE 802.1Qbv), credit-based traffic shaping (IEEE 802.1Qav). In case bandwidth increments shall be in units of 2 Mbps this translates to 10 Gbps / 2 Mbps = 5000 separate Timeslots.

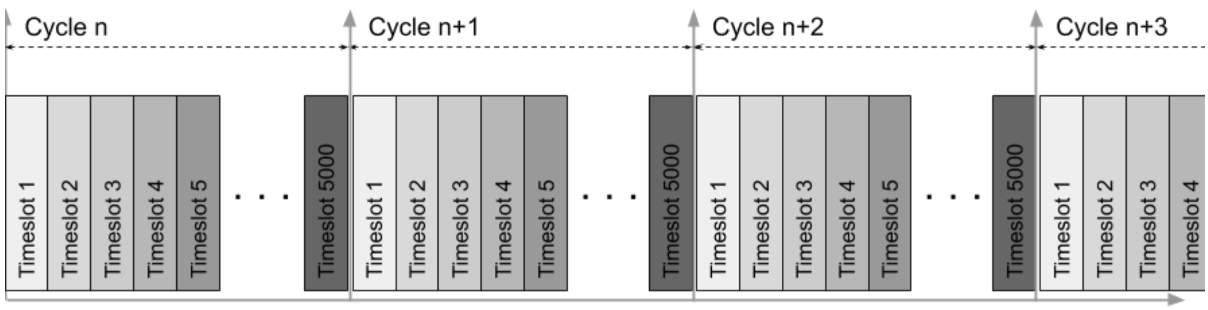

Another way to represent this can be a "wheel" with 5000 Timeslots where each sIF can occupy zero, or more, Timeslots: The following figure visualizes how different sIF can be run using NEV:

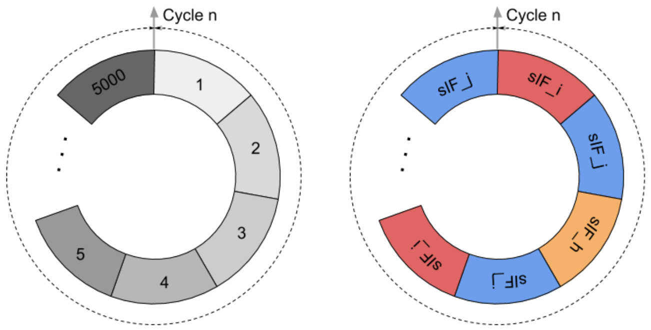

Index	sIF
0	sIF_i
1	sIF_j
2	sIF_h
3	sIF_j
4	sIF_i
...	
5000	sIF_j

Again, while this figure aims to visualize the concept for 3 sIF (here red, blue and yellow), the NEV chip can support up to 4096 sIF.

FPGA-Based Multi-Processor System-on-Chip Implementation

For cost and energy efficiency the NEV chip has been implemented in 16nm semiconductor technology on top of a modern FPGA-based Multi-Processor System-on-Chip (MPSoC). Implementation follows state-of-the-art data-in-motion processing from multiple physical 10 GigE ingress ports all the way to multiple physical 10 GigE egress ports. At electrical layer, the NEV chip deploys a mature PCS/PMA subsystem that is designed to the Ethernet requirements for 10/25 Gb/s operation specified by IEEE 802.3 Clause 49, IEEE 802.3by, and the 25G Ethernet Consortium and supports reliable copper and/or fiber connectivity for 10 GigE. Hardware / software partitioning targets best cost-to-performance ratio for NEV by mapping network processing either to programmable logic (PL) or to Linux software. Linux runs on the embedded multi-core ARM CPUs and performs network management and administration plus software APIs to Open vSwitch. Optimized hardware / software interfaces connect the FPGA-embedded AXI4 data movers with standard Open Source Linux network devices (netdev).

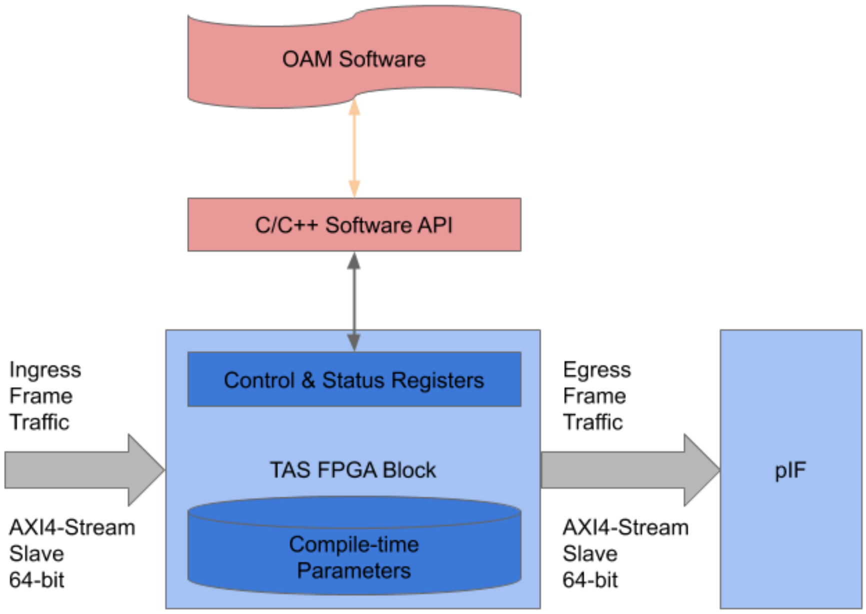

Via a software API, bandwidth, traffic priority, traffic burstiness, etc. can be controlled by assigning certain Service Interfaces to timeslots in the Traffic Shaper. Short cycle times, enabled by network protocol processing in PL, keeps processing latencies short and deterministic.

Conclusion

Algoblu's Network Element Virtualization (NEV) which has been implemented in form of the NEV chip on Xilinx Zynq UltraScale+ MPSoC Architecture offers fine-grained and dynamic provisioning of networking resources such as application bandwidth and can lead to more efficient resource utilization. Xilinx MPSoCs with multi core A53 architecture offers highly flexible implementation of hardware and software partitioning.

NEV is ideal for IoT, cloud games, VR/AR broadcasting, metaverse application which has higher requirement for low-latency technology. It is a shortcut to enable carriers to provide multi-tier services to meet users' personalized requirements with guaranteed SLA and without changing the existing network infrastructure.

Authors & Contact Information

Hua Edward Qin - Chief Product Officer – Algoblu
www.algoblu.com

Endric Schubert, PhD – Chief Technology Officer – Missing Link Electronics
www.missinglinkelectronics.com