

Red Hat and Kaloom™ scalable UPF

Surpass Terabit performance to support millions of 5G subscribers

Mobile network operators (MNOs) are keen to maximize revenue earned from their 5G assets, but have the dilemma as to where to allocate 5G capacity to support new applications that deliver high performance and ultra-low latency, and serve a high number of subscribers. Whereas new applications will monetize their investments, improved performance will drive a better user experience. Both are vitally important, and improved packet processing is key to achieving both.

One avenue for performance enhancements lies in the disaggregation of core network functions into the control plane and the user plane. This separation creates room for innovation, programmability, and optimization of each network function to realize enormous efficiencies. In 5G, the user plane function (UPF) creates the user plane for the flow of digital data traffic, while the control plane determines the ideal target for hardware acceleration and service chaining. While the mobile traffic throughput requirements per location at the MNO Edge and Core are less than one terabit per second (Tbps) today and can be served by commercial off-the-shelf servers (COTS), the 5G data explosion combined with the convergence of fixed wireless access (FWA) will create more demanding requirements per location. This could put at risk the quality of digital services and the market's overall growth for new services to be delivered from the edge. MNOs are receptive to the potential of innovative software technologies bolstered by the diversity of programmable hardware for accelerated traffic routing of a larger volume of data with incremental investments.

Achieving high-performance packet processing

5G networks may have to cope with a tsunami of packet flows choking network hardware and even virtual machines that are now migrating to cloud-native network functions (CNFs), despite the ever-growing investments in network capacity to meet bandwidth demands. The relentless growth in demand at the edge for applications and data processing capacity with growing 5G standalone (SA) networks will not be relieved by bigger pipes alone without the realignment of fragmented capacity to improve packet processing efficiency.

While the subscriber base surges to support billions of connected devices, the performance demands for the UPF may have to rise dramatically to multiple terabits per second throughput with latencies lower than four milliseconds. The hardware capacity typically far exceeds the throughput performance achieved but falls short due to the redundancies, obstacles, and speed limitations created by the complexity of the operating system (OS) and virtualization software stacks. Likewise, the performance of traditional virtualized network functions (VNF) and virtual machine (VM)-based network functions has fallen short as the overhead of their operating systems drag down throughput and latency. Moreover, bloated network stacks can impact the delivery of potential future expected data rates superior to one Tbps.

A new complementary approach to data processing hardware is needed. The Intel architecture family of processors is widely recognized as the hardware of choice for general tasks. Complementary Intel technologies allow MNOs to further scale their solutions to best fit network locations from core to edge according to their traffic requirements. Their value is further enhanced when they offload specialized tasks to more efficient, purpose-built hardware and systems for raising the overall network performance. Furthermore, the flexibility of Intel software allows functions like encryption and decryption, firewalls, packet inspection, routing, and storage networking, not only on the Intel family of processors, but can also be optimized by P4-programmable packet processors such as Intel® Tofino™ intelligent fabric processors (IFPs) combined with Intel Stratix®10 field programmable gate array (FPGAs).

Innovations for increasing the speed of packet processing

Streamlining software and protocols can achieve higher rates of packet processing throughput. When user equipment (UE) requests data, packets are processed efficiently as they cross fewer software hops before network resources process them. Additionally, parallel processing moves data more rapidly.

Hardware acceleration in the hybrid mix

VNF running on COTS equipment is the accepted means to scale. Hardware acceleration can be applied as a complementary technology to manage the exacting demands of higher throughput, lower latency, and flexibility to cope with a diversity of needs. Historically, MNOs had to choose between higher throughput and flexibility. Application-specific integrated circuits (ASICs) were faster, cheaper, and more power efficient, whereas FPGAs were more flexible, but more expensive and less power efficient. Moreover, hardware acceleration was constrained by proprietary hardware with the algorithms tied to the chip in use.

The P4-programming language makes it possible to do multiple kinds of accelerated packet processing on a single hardware platform unencumbered by proprietary implementations. P4 programmability provides limitless flexibility with intent-based networking and supports multiple use cases, ranging from the customization of network solutions for efficient scale, increased performance, and enhancement of existing datacenters. Capabilities such as real-time telemetry, security, deep insights into the networks, artificial intelligence and machine learning (AI/ML), and layer 4 load balancing are also good use cases for enterprise network solutions. MNOs also do not have to sacrifice flexibility for higher throughput when using Intel IFPs that are purpose-built to take advantage of the P4 language and are cost and power efficient.

P4 configures the network for the targeted performance

P4 creates the code that includes the protocol header definitions, the protocol parsers, the pipelines with flow tables, and the match-action semantics. It optimizes the data flow between the software and hardware elements and sets up the parameters for the flow of packets to meet the targeted latencies and service quality. A P4-programmable user plane can run a wide range of network functions, including layer-two switching, routing, load balancing, firewalls, and 5G UPF.

A UPF implemented in the P4 language provides a new way of processing data in a 5G network. It plays a key role in mapping incoming and outgoing traffic connecting distributed gNodeB (base stations) to the 5G stack with the desired targets of throughput, latency, and scalability demands of the growing data-intensive 5G applications. Since P4 code is software, it can be updated with new 3GPP releases and customized much more rapidly than hardware alternatives.

Introduction to Kaloom containerized UPF

Kaloom's Unified Edge Fabric is a fully automated networking fabric solution specifically designed for simultaneous 4G and 5G applications at the network edge, providing native support for 5G network slicing.

Kaloom provides a programmable multiple Tbps UPF that can either be integrated into its Unified Edge Fabric or deployed as a standalone solution. Designed with the most demanding 5G workloads in mind, it offers a scalable, low latency, and high-performance UPF solution. Kaloom's UPF is ideal for emerging 5G cloud edge, hybrid 4G and 5G, and 5G packet core deployments.

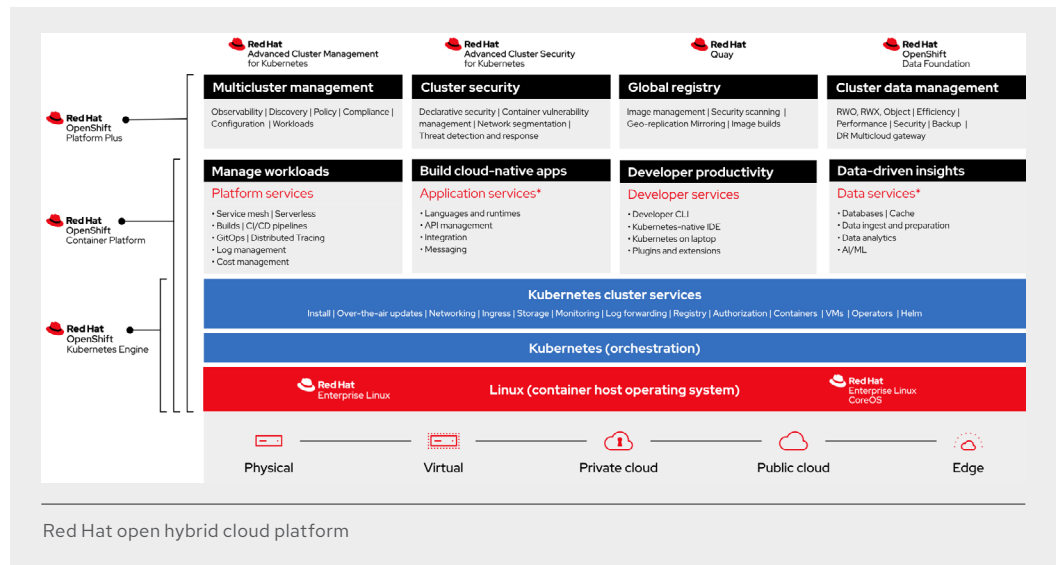
Kaloom's 3GPP Release 15-compliant UPF uses a programmable pipeline implementing packet rules (flow identification), forwarding action rules, usage reporting rules, and quality of service (QoS) enforcement rules for an entire UPF data plane. It is characterized by:

- ▶ Improved performance, while significantly lowering network latency from milliseconds to **sub-four microseconds** to meet the demands of latency-sensitive applications such as Internet of Things (IoT), augmented reality (AR), and virtual reality (VR) as well as significantly reducing the total cost of ownership (TCO).
- ▶ Scaling to **2.2 Tbps** full duplex throughput and over **four million** user sessions without affecting quality.
- ▶ The optimized cloud-native solution for container-based workloads.
- ▶ A centralized or distributed datacenter architecture.
- ▶ Supporting 5G end-to-end network slicing.
- ▶ Functioning as a fully standards-based and proven multi-vendor solution.

By being programmed in P4, the Kaloom UPF can accommodate new features and requirements of 3GPP in weeks instead of months or longer. MNOs, moreover, can innovate in less time and introduce new services to meet evolving enterprise customer requirements, allowing them to be more competitive and prepared for the future marketplace.

Red Hat OpenShift Container Platform for cloud-ready core network functions

Red Hat® OpenShift® Container Platform offers MNOs a reliable cloud-native foundation for 5G core network functions and can span the entire network from core datacenters to edge deployments. OpenShift Container Platform will help MNOs deliver new 5G services efficiently and optimize their operational model through simplified workflows.



OpenShift Container Platform has a rich set of capabilities to support the 5G core network that includes a real-time OS, hardware acceleration, and massive-scale service assurance for a growing number of ecosystem network functions. Red Hat’s partner ecosystem delivers freedom of choice for MNOs to select their preferred 5G core network functions from different vendors to fit their operational and business needs.

Kaloom’s Unified Edge Fabric and UPF is a fully containerized solution using the OpenShift Container Platform which provides a Platform-as-a-Service (PaaS) built around containers orchestrated and managed by Kubernetes on a foundation of Red Hat Enterprise Linux®.

Kaloom’s high-performance architecture

The performance tests described in this paper are based on Kaloom software running on a network switch based on the Intel Tofino programmable Ethernet switch ASIC and a server equipped with an Intel Stratix 10 FPGA card connected to one of the switch ports as a worker node. This is just one hardware configuration that can achieve high-performance results.

Kaloom’s UPF software can also run on a switch-server using the Intel Tofino series combined with an Intel Stratix 10 FPGA. This combination provides higher packet processing throughput than on just a switch or FPGA by itself. Likewise, the same software can run using an Intel Xeon® Scalable processor for an efficient and powerful control and management plane (Figure 1).

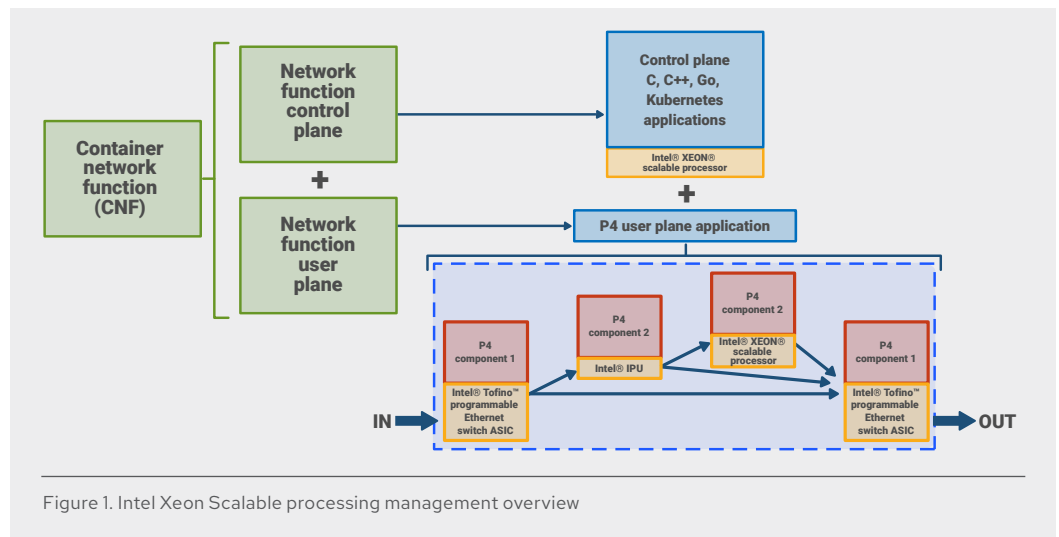


Figure 1. Intel Xeon Scalable processing management overview

Without this heterogeneous programmable hardware used for its specific strength, a similarly capable UPF would need a significantly larger number of servers and network interfaces for equivalent performance. A multiserver approach adds cost, latency, and additional power, cooling, and space requirements, which can be constrained in many edge locations.

Kaloom uses the OpenShift Container Platform to deploy the entire system using the switching, FPGA, and compute functionality built into just three switches running the OpenShift Container Platform, the Unified Edge Fabric, and the UPF.

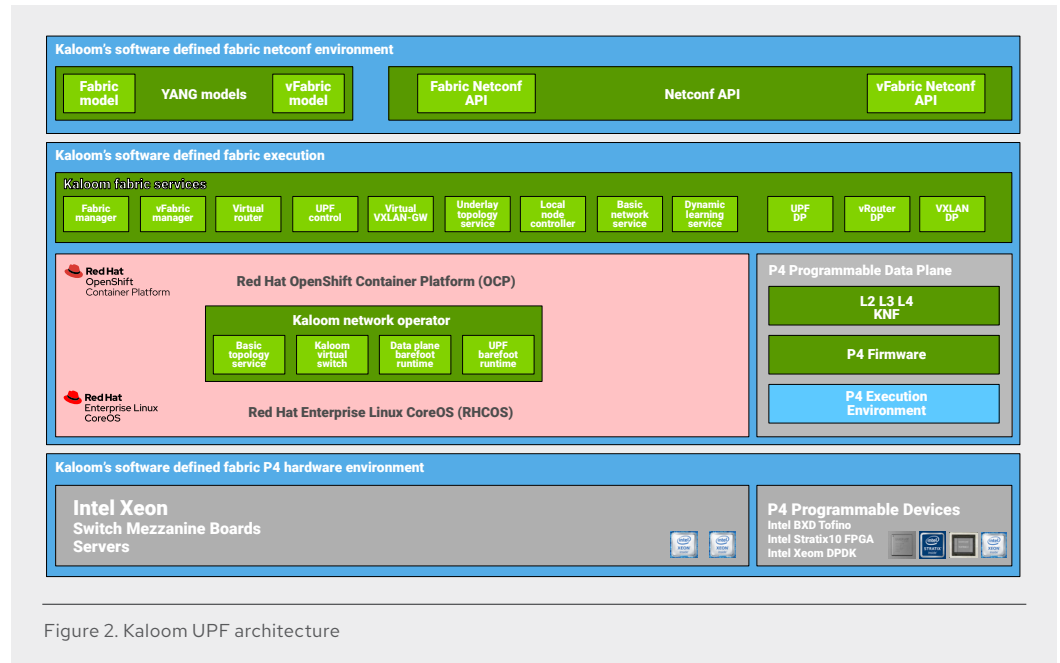


Figure 2. Kaloom UPF architecture

Kaloom’s performance levels serve the needs of one or more tenants allowing for consistent scalability. Kaloom’s Unified Edge Fabric can be subdivided into fully isolated 5G network slices to suit the needs of multiple tenants, each configured for their application and performance requirements. Every slice has a corresponding virtual datacenter partitioned from a physical edge datacenter, mirrored by a virtual fabric (vFabric) with a logical UPF slice instance (vUPF). Kaloom provides network services to the isolated vFabrics—vSwitch, vRouter, VxLAN gateway, UPF, and other 5G core network functions such as session management function (SMF) and access and mobility management function (AMF). Kaloom also works closely with ecosystem partners to integrate and service chain virtualized and containerized network functions.

System test environment

Kaloom’s solution is built on a leaf-spine switching architecture and provides further performance efficiencies by embedding and running the UPF data plane functions in the P4-programmable leaf switch. Kaloom’s 5G UPF executes user plane application on the leaf switch’s Intel Tofino Ethernet switch ASIC and Intel Stratix 10 FPGA devices. As a result, it further reduces the latencies and increases the throughput as packet routing paths are cut short.

Figure 3 provides a conceptual and high-level architectural overview illustrating Kaloom’s logical UPF implementation that consists of distributed functionality within a Unified Edge Fabric or a standalone node. Incoming packets from the radio access network (RAN) or the data network arrive on the edge switches of the fabric upon which a load balancing function distributes the packets to the various leaf switches.

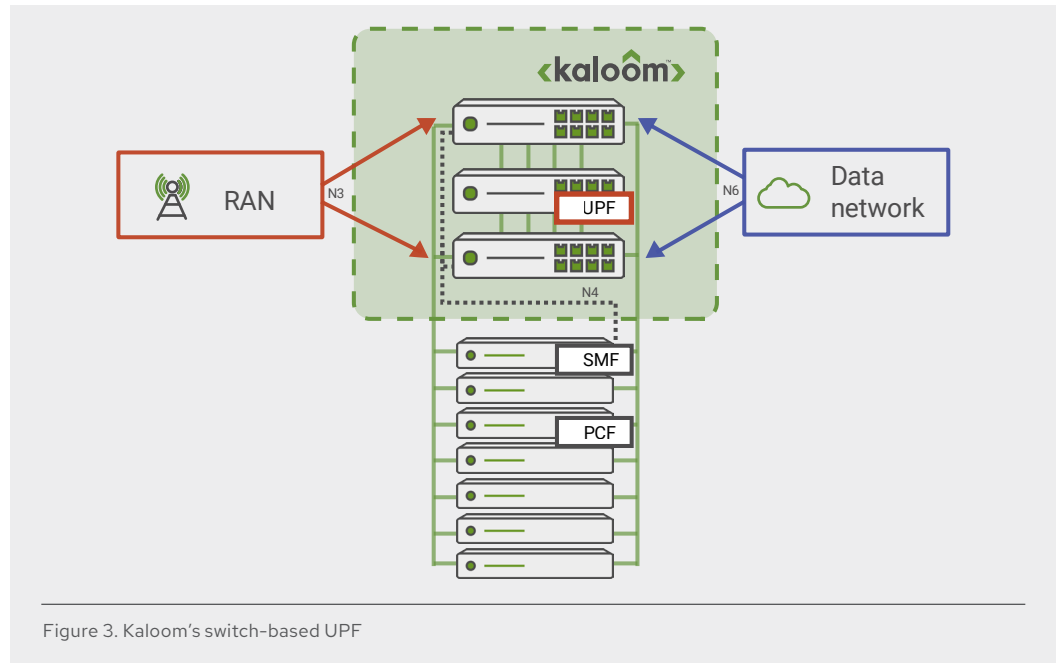


Figure 3. Kaloom's switch-based UPF

The UPF control plane scales according to the capacity and number of physical fabric nodes used as controllers. The UPF data plane is distributed over the leaf switches. The more leaf switches deployed in an edge fabric, the higher the resulting capacity of the UPF.

The system test environment that Kaloom built for this test included a mmWave air interface operating in the 20 GHz frequency band. Additionally, the edge UPF was configured to support up to 50 gNodeB base stations from which it could potentially serve up to millions of connected user equipment (UE), ranging from mobile devices to IoT sensors. The total inbound bandwidth to this switch cluster could reach 100 Tbps, with an expectation of five ms of end-to-end latency, predicated on stated numbers for the dimensions of the system.

Network topology

Figure 4 shows the data flow of the test, based on the 3GPP-standard 5G core with the UPF, located at the bottom, connected to the radio access network (RAN) via the N3 interface, to the control plane via the session management function (SMF) via the N4 interface, and to the data network via an N6 interface. The UPF processes the incoming data from the RAN by first terminating inbound traffic and encoding/decoding all the data flows. Then, the UPF communicates to the session management function to access the control plane capabilities.

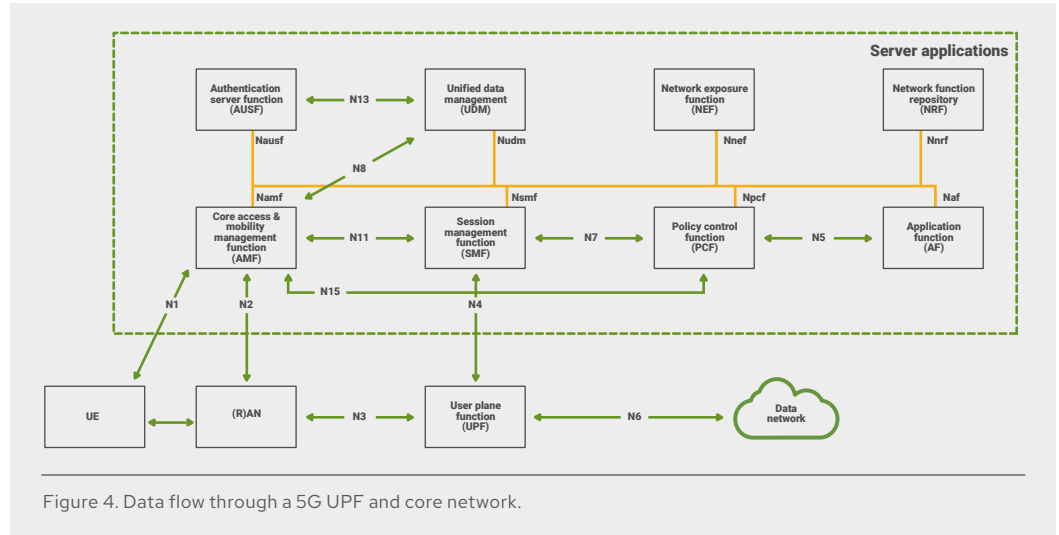


Figure 4. Data flow through a 5G UPF and core network.

Hardware configuration: Edgecore Wedge100BF-32QS

The test utilized the Edgecore Wedge100BF-32QS top-of-rack (TOR) or spine switch for high-performance datacenters. The Edgecore Wedge100BF-32QS is designed with the P4-programmable Intel Tofino Ethernet switch ASIC, part of the Intel® Intelligent Fabric Processor (Intel® IFP) family. Intel Tofino IFPs offer users customizable manageability of their packet processing, switching, and forwarding pipeline with intelligence, performance, visibility, and control optimized for datacenter fabric.

In a 1RU form factor, the Edgecore Wedge100BF-32QS provides line-rate L2 and L3 switching with a total capacity of 6.4 Tbps across 32 QSFP28 ports. When deployed as a top-of-rack (ToR) switch, it supports 25/50 GbE connections to servers with 40/50/100 GbE uplinks or as a spine switch supporting 40/50/100 GbE ToR and spine interconnects.

The switch is designed for uptime with redundant and hot-swappable power supplies, fan modules, and preloaded diagnostics. Also, open network install environment (ONIE) is preloaded for automated loading of compatible open source and commercial network operating system (NOS) offerings.

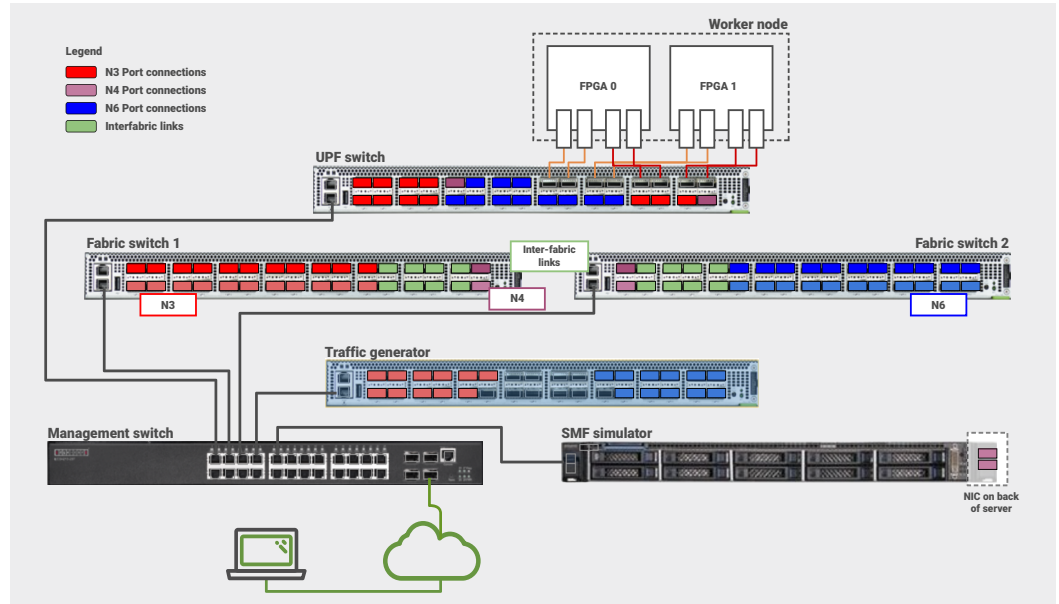


Figure 5. Test network configuration

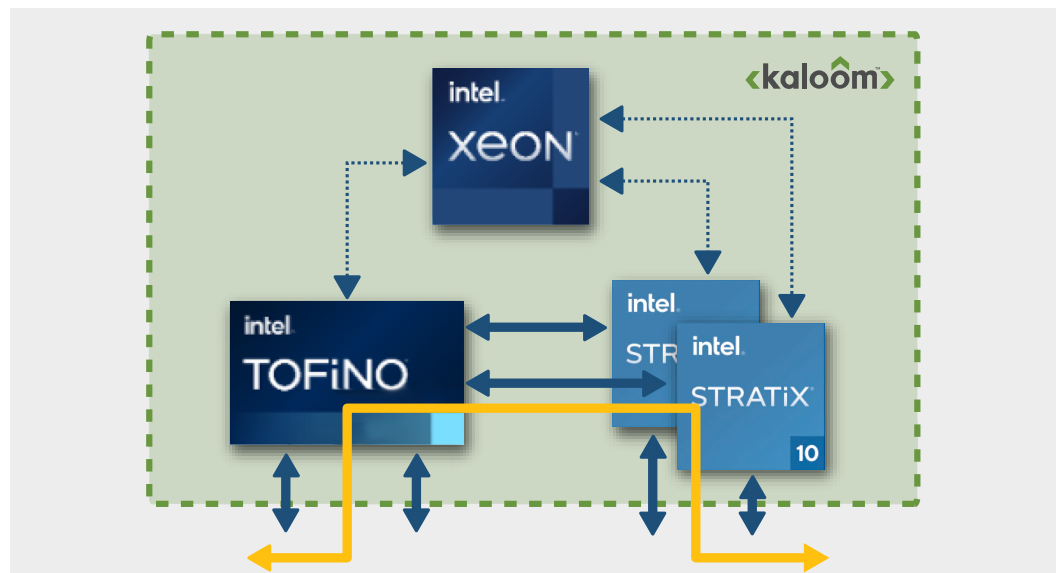


Figure 6. Test network block diagram with data flows

The performance of the Intel Tofino-based Edgecore Wedge100BF-32QS switches is further enhanced with hardware acceleration when complemented with Intel Stratix-10 FPGAs. Kaloôm's 5G UPF user plane executes on the Intel Tofino switch ASIC and FPGA devices with up to three switches. As a result, it provides a 5G UPF with much higher throughput, lower latency, and an improved cost and performance ratio than solutions based solely on traditional server-based deployments.

Testing method: The UPF system configuration consists of UPF switch and two fabric switches (Edgecore Wedge100BF-32QS) providing the L2 and L3 connectivity between the UPF and the external world. External world connectivity was emulated through an in-house traffic generator and the application server hosting the SMF simulator.

The traffic generator was also based on an Edgecore Wedge100BF-32QS switch which generated the data traffic at a terabit per second. The SMF simulator used to advance the session creations was Emblasoft's Solver hosted on a server with an Intel® Xeon® E5-2650 processor running at 2.00 GHz with 32 cores, 128 gigabytes of RAM, and 25 GbE connections provided by an Intel® Ethernet Network Adapter XXV710. The cabling between the two fabric switches and the traffic generator included 22x100G three feet fiber cables with the corresponding 100GBASE-SR4 QSFP28 transceivers on either side of the cable.

The traffic generated by the traffic generator is both N3 traffic (UE simulated traffic) which is encapsulated with GPRS Tunneling Protocol (GTP) and N6 traffic for the downlink traffic destined to the UE. For the traffic to reach the UPF a bundle of 22x100G fibers was used to handle both the N3 and N6 traffic. In addition, 2x100G fiber cables were used for N4 connectivity to reach the UPF. The SMF simulator connectivity from the fabric switches used two 4x25G breakout cables with QSFP28 transceivers on the fabric switches and SFP28 transceivers on the server side.

The session scaling tests involved Silicom FPGA SmartNIC N5010 series with Intel Stratix 10 FPGAs. The FPGAs were hosted in a server with an Intel Xeon E5-2650 processor running at 2.00 GHz with 32 cores and 256 Gigabytes of RAM. The FPGA card was provisioned with two million entries for uplink and two million entries for downlink traffic for a total of two million user sessions per FPGA. The connectivity from the UPF switch to the FPGAs used 4x100G fiber cables with transceivers previously mentioned for each FPGA. The provisioning was done by simulating the UPF host gRPC signaling to the FPGA component responsible for populating the entries into the FPGA.

Performance result

By using Kaloom Unified Edge Fabric, OpenShift Container Platform, and P4-programmable switches from Edgecore based on the Intel Tofino programmable Ethernet switch ASIC and Intel Stratix-10 FPGA, the test demonstrated an edge UPF that can handle over four million simultaneous sessions while maintaining full duplex throughput values of 2.2 Tbps.

Conclusion

Data flows across the UPF have scaled with the proliferation of software-defined network functions, virtualized machines, bare metal, and containerized infrastructure. Additionally, the software overlay on hardware frees MNOs to route traffic to avoid congestion and penalties on the quality of service. On the other hand, the software sprawl has increased cost and complexity as overheads accumulate with the increasing need for software clusters to interwork.

Bloated software stacks cost MNOs who realize lower returns from their investments. UPF software taking advantage of heterogeneous P4-programmable hardware for packet processing achieves the desired performance at significantly lower costs. P4 programming rationalizes the forwarding of packets, and hardware acceleration attains the desired terabits rates of packet processing.

The test shows that Kaloom's programmable UPF has the performance and user session capacity to meet the above-mentioned industry scalability requirements.

Learn more

[Programmable networking fabrics for core and edge data sheet](#)

[Kaloom and Red Hat Solution Brief "A unified solution for the distributed edge"](#)

[Red Hat OpenShift](#)

[Edgecore Wedge100BF-32QS](#)

[Intel Network Builders](#)

[Intel Tofino Intelligent Fabric Processors](#)

[P4 programmable language](#)

[Intel Stratix 10 FPGAs](#)

[Intel Xeon Processors](#)

[Press release: Kaloom delivers UPF functionality to Telenor's multi-vendor standalone 5G core solution](#)

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.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

Copyright 2022 Kaloom, Inc. The information contained herein is subject to change without notice and is correct to the best of Kaloom's knowledge at the time of publication. Kaloom shall not be liable for technical or editorial errors or omissions contained herein. Kaloom, the Kaloom logo, Software Defined Fabric and Cloud Edge Fabric are trademarks of Kaloom Inc. Other product or service names may be trademarks or service marks of others.

522/TM/HO9/PDF

351143-001US



About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

[f facebook.com/redhatinc](https://facebook.com/redhatinc)
[@RedHat](https://twitter.com/RedHat)
[in linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

North America
1888 REDHAT1

**Europe, Middle East,
and Africa**
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com

redhat.com
#F32063_1022

Copyright © 2022 Red Hat, Inc. Red Hat, the Red Hat logo, and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. All other trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.