

# Lighting Up the 5G Core with a High-Speed User Plane on Intel® Architecture

**The Metaswitch Fusion Core 5G user plane function (UPF) hits a packet throughput of 500 Gbps when running on dual-socket 2nd Generation Intel® Xeon® Scalable 18-core processors.<sup>1</sup>**



metaswitch

With the promise of high-bandwidth, low-latency connectivity for both fixed and mobile endpoints, ubiquitous 5G coverage will fundamentally change the way people and devices communicate. While there has been much focus on new radio (NR) network technologies, the benefits of 5G can only be fully realized with commensurate innovations in the mobile core infrastructure. To meet the diverse demands of end users, massive sensor grids, and communications service providers, the 5G Core (5GC) completely separates control and user plane functions, moving significant data plane components into distributed multi-access edge computing (MEC) facilities while supporting a centralized service based architecture (SBA) for delivering intelligent network operations.

Demanding the granular scalability of end-to-end network slices, the promise of 5G can only be realized with a core built entirely on cloud native constructs with virtualized network functions (VNFs). These VNFs must employ highly distributed software architectures, be implemented on dynamic compute resource pools, and feature massive deployment and configuration automation. While SBA elements must be built from scratch using microservices methodologies to support stateless instantiation in container-based environments, cost-effectively implementing powerful user plane packet processors on general purpose compute platforms requires significant innovations in data plane acceleration techniques.

## The Role of a UPF

The user plane function (UPF) is a fundamental component of a 3GPP 5G Core infrastructure system architecture and represents the data plane evolution of a Control and User Plane Separation (CUPS) strategy. The 3GPP's technical specifications define the UPF as both the packet data unit (PDU) session anchor (PSA) point for providing mobility within and between radio access technologies (RATs) and the interconnect point between the mobile infrastructure and the data network (DN). Where there is more than one UPF-PSA employed for multi-homing, an Intermediate UPF (I-UPF) is deployed. To perform this role, it is tasked with providing encapsulation and decapsulation of the transport protocols employed in the data plane, such as the GPRS Tunneling Protocol (GTP). Prior to forwarding packets, the UPF implements independent classifiers on both the uplink and downlink to direct flows to specific DNs or PSAs based on traffic matching filters.

With services ranging from enhanced mobile broadband (eMBB), massive machine type communications (mMTC), and ultra-reliable low latency communication (uRLLC), 5G must support the requirements of a broad range of new mobile applications plus fixed wireless access (FWA) and fixed mobile convergence (FMC) endpoints. The 5G UPF supporting these disparate connectivity demands must be deployable in a myriad of distinct container-centric or serverless private, public, or hybrid edge/core cloud infrastructures, requiring complete ubiquity and total portability.

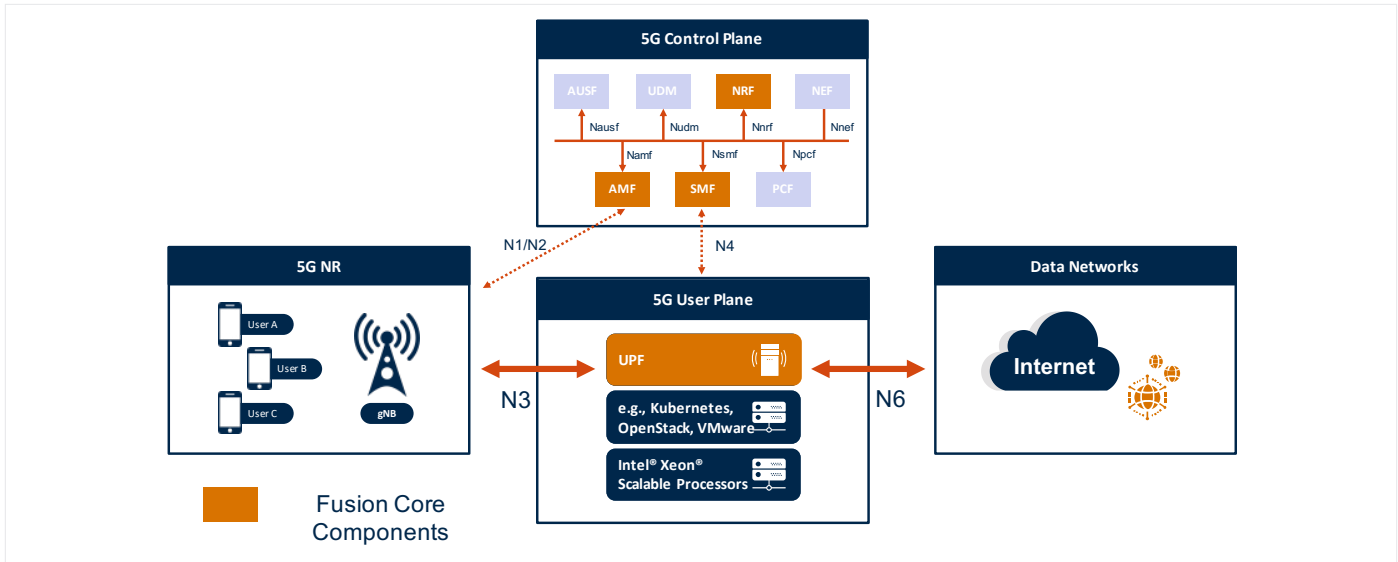


Figure 1. Metaswitch Fusion Core block diagram.<sup>2</sup>

### UPF Testing Profile and Results

Metaswitch has demonstrated how software-based 5G UPF can deliver a superior 5G user plane deployed within containers on standard server platforms leveraging Intel® Xeon® Scalable processors. Employing Internet MIX (IMIX) traffic profiles and using a total of 34 Intel Xeon Gold 6254 processor cores, a single-server UPF achieved a packet throughput of over 500 Gbps on a simplified pipeline.<sup>1</sup>

The test comprised 100,000 sessions from 160 simulated gNBs with a total of 10,000 Access Control Lists (ACLs) provisioned. A complete UPF packet processing pipeline was utilized with full GPRS Tunneling Protocol (GTP) encapsulation/decapsulation plus session-based policing and metering. Two asymmetric traffic profiles were used, with differing ratios of uplink to downlink traffic (see Figure 3).

Although quite exceptional, this level of performance is particularly critical in constrained multi-access edge computing environments where the majority of UPF instances would be deployed typically one step further

out than previous generation EPC gateway functions. With these traffic characteristics, the gating factor for achieving a performance much beyond 500 Gbps was the number of available PCIe lanes, indicating a pathway to Tbps UPF throughput on a single server with PCIe Gen4 upgrades and the next generation of processors.

#### A UPF must

- have superior performance on the lowest possible footprint
- be highly orchestrated and automated for ease of deployment
- granularly scale in network slices to meet changing traffic patterns
- feature a highly programmable packet pipeline for new services

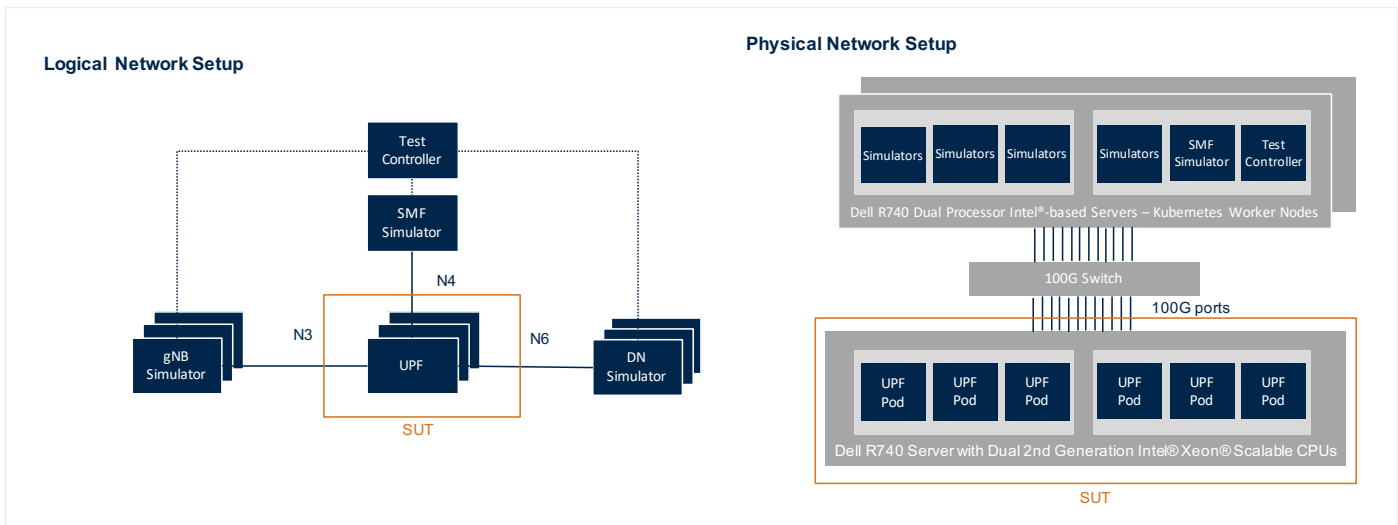


Figure 2. User Plane Testing Profile: Logical and Physical Network Setups.<sup>2</sup>

The complete UPF consists of poll-mode packet processing containers handling the user plane load, along with management and PFCP server processing components running on additional cores. User plane traffic is load balanced across cores using hash-based distribution on the NIC.

	CONFIGURATION #1	CONFIGURATION #2	CONFIGURATION #3
TOTAL UPF CORES	6	34	34
PACKET PROCESSING CORES	4	25	25
AGGREGATE BW (N3 INTERFACE)	95 Gbps	502 Gbps	503 Gbps
AGGREGATE BW (N6 INTERFACE)	88 Gbps	467 Gbps	478 Gbps
PACKET RATE	18.5 Mpps	98 Mpps	71 Mpps
AVERAGE PACKET SIZE (UL & DL)	596 bytes	596 bytes	839 bytes
PACKET RATE DL : UL RATIO	2.4 : 1	2.4 : 1	2.1 : 1
BANDWIDTH DL : UL RATIO	3 : 1	3 : 1	9 : 1
PACKET PROCESSING CORE UTILIZATION (AVERAGE)	61%	67%	48 %

Figure 3. User Plane Testing Results.<sup>1</sup>

### Test Scenario

- A Dell R740 server with two Intel® Xeon® Gold 6254 processors (microcode 0x5000021), 384 GB RAM (24 x 16 GB DIMMs) and a 1.6 TB SSD Self-Encrypting SATA Drive.
- The server featured Dual-Port 100G Mellanox ConnectX-5 NICs and a Quad-Port Intel® 10G X550T embedded NIC.
- Metaswitch demonstrated 500 Gbps throughput using 34 processor cores.<sup>1</sup>
- Software included Docker 18.06, Kubernetes 1.14, and DPDK 18.02 (application).
- The testing environment was set up with 100,000 emulated subscribers (100,000 PDU Sessions), 160 emulated gNodeBs and with an attach/detach rate of 1,000 subscribers per second. IPsec, buffering and L7 DPI were not enabled.
- The user plane packet processing included full UPF function [session matching on IP/F-TEID, GTU-U encapsulation and decapsulation, usage reporting (metering)] with policing (MBR QoS enforcement), forwarding actions and a representative set of 10,000 interface ACLs.
- Two sample IMIX packet distributions with a mix between 64 and 1,500 byte packets.

### Solution Overview: Metaswitch Fusion Core 5G UPF

Typically the domain of customized application-specific integrated circuits (ASICs) and network processor (NPU) hardware in dedicated appliances, the 5G UPF is a pure software-based packet-forwarding platform. It must implement rapid encoding and decoding of tunneled IP payloads, applying complex pipelines and enforcing stringent Quality of Service (QoS) metrics while consuming as few compute resources as possible.

Intel® Network Builders partner Metaswitch has developed a complete cloud-native 5G core solution featuring a pure software UPF that meets the performance and resource efficiency needs of high-throughput 5G networks. To fulfill the UPF functionality, the Metaswitch solution is built to run on commercial off-the-shelf (COTS) servers powered by Intel Xeon processors. It features an industry-leading, high-performance, forwarding engine and a unique distributed state implementation for resiliency and granular scalability cloud infrastructures. This is achieved on an extraordinarily small processing footprint at a very low price point, which is increasingly important for platforms deployed within MEC environments.

#### Metaswitch Composable Network Application Processor

While the open-source Data Plane Development Kit (DPDK) has helped accelerate packet processing within COTS servers, such polled zero-copy direct memory access techniques cannot, on their own, provide the levels of price/performance required of core network routing functions with numerous forwarding rules and multi-layer encapsulation.

The Composable Network Application Processor (CNAP) is a flexible and powerful software packet processing engine that Metaswitch developed to provide the foundation for virtualized network functions that handle data plane traffic. With a unique combination of programmability and performance, CNAP implements a packet-processing pipeline with match-action classifiers that allow lookups on multiple elements of the packet header in a single operation. With careful design of the packet pipeline, this allows multiple logical operations to be collapsed into a single match-action step.

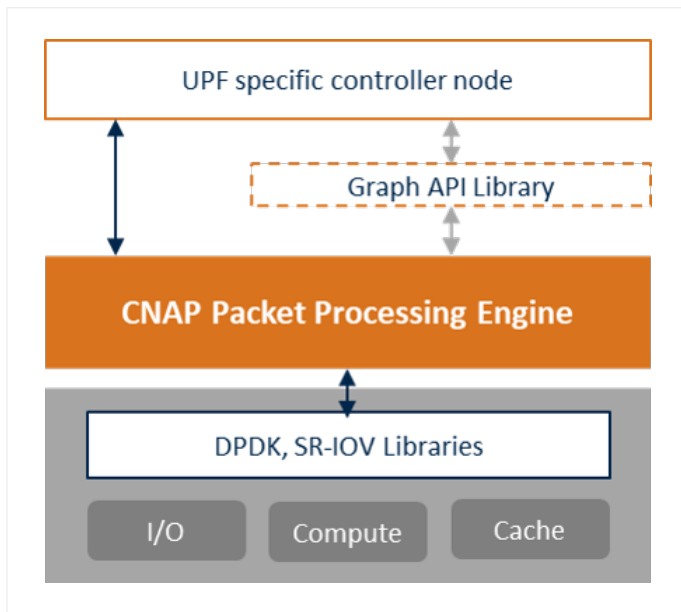


Figure 4. Metaswitch UPF block diagram.<sup>2</sup>

CNAP makes the best possible use of instruction and data caches in Intel® x86 architecture CPUs to maximize efficiency. In addition to processing packets in batches, CNAP leverages explicit pre-fetch instructions and process interleaving to avoid cache stalls while ensuring that all match-action instructions, packet header data, and lookup table data are in-cache for the fastest possible execution. The pipeline is described by a gRPC API with P4-like semantics and applied using a series of API calls to the CNAP engine. This allows complete flexibility in configuring the UPF to implement new data plane protocols or tunneling techniques without the typically long development cycle. A fully containerized Metaswitch UPF can be deployed on a minimum of three processor cores, allocated respectively to management and control plus an uplink and downlink packet processing engine.

#### Powered by Intel® Xeon® Scalable Processors

For the servers that run the Metaswitch 5G Core, Metaswitch uses Intel Xeon Scalable processors, which are Intel's second generation platform for cloud-optimized, 5G-ready networks. The processors have an open architecture and a platform that provides a future-ready foundation for agile networks that can operate with cloud economics, be highly automated and responsive, and support rapid and more secure delivery of new and enhanced services enabled by 5G.

Intel Xeon Scalable processors provide capabilities that the Metaswitch 5G Core uses to improve throughput. For example:

- Vector instructions speed up various packet manipulations and data structure look-up functions by making use of Intel® Advanced Vector Extensions (Intel® AVX) CPU instruction set. Intel® Advanced Vector Extensions 2 (Intel® AVX2) instructions, which operate on vectors up to 256 bits in size, are widely available in servers. Support for the latest Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instructions will allow the software to operate on larger vectors with single instructions and have some other performance benefits, such as faster access to unaligned data.
- The expanded per-core L2 cache of 1 MB in the Intel Xeon Scalable processors allows for improved prefetching algorithms to take advantage of the larger cache size.

#### Microservices-Based User Plane for Cloud-Native Deployment

As a true cloud-native application, the software components that make up the Metaswitch 5G Core solution—including the user plane packet processing engine—are delivered as container images and designed for lifecycle automation via Kubernetes. Affording maximum deployment flexibility, these containers may be deployed inside virtual machines (VMs) or on bare metal in any private, public, or hybrid cloud.

Metaswitch has engineered its 5G core user plane using interchangeable and disposable containers to provide outstanding performance while reducing the software maintenance required in production environments. A process that is not functioning properly can be easily killed and replaced, without any impact on the overall system. The key benefits of UPF containerization include:

- The system is very robust and fault tolerant, providing active-active N+k redundancy.
- The system can scale down to 4 cores of a single CPU for far edge use case applications.
- Lifecycle management is completely automated, healing software malfunctions and simplifying software upgrades.
- The software can scale elastically through the addition of new containers.
- Containers can be remotely initiated or modified. This approach is inherently cloud friendly and high performance.

From the MEC to the core compute cloud infrastructure, intermediate and PDU session anchor user plane functions must be dynamically deployable and granularly scalable in autonomous network slices. This can only be achieved through the innovative application of cloud-native design methodologies and automation techniques employed by the Metaswitch Fusion Core 5G UPF.

#### Distributed State for Cloud-Native Environments

Metaswitch developed a unique distributed state functionality for the 5G core user plane, allowing the software to run on multiple servers in a cloud environment. The information required to process packets through a 5G core includes tunnel endpoint state, subscriber policy state, and routing table state. This data is accessed rapidly to process each packet very quickly and must be highly available for session preservation in the event of failures or other lifecycle events.

Metaswitch's application of cloud native principles to stateful real-time packet processing includes integrated and automated "sharding"—or partitioning—of state information to all UPF instances involved in implementing a forwarding pipeline. This allows state processing to be distributed and redistributed as needed between UPF instantiations while facilitating backups and lifecycle management activities.

Automated distribution of shards between instances, combined with multiple wellness checks, ensures user sessions are preserved across instance failures without impacting user experience.

## Conclusion: Accelerating the Path to 5G Services

With the Metaswitch 5G Core solution and its high-performance user plane, running on 2nd generation Intel Xeon Scalable processor-based servers, Metaswitch has addressed communications service providers' (CommSPs) cloud deployment, service innovation, and network evolution pain points by employing revolutionary cloud native technologies. This functionality and performance will help CommSPs to accelerate the rollout of 5G packet core technology and new revenue-generating services with unmatched economics and deployment flexibility.

## About Metaswitch

Metaswitch is a leading cloud-native communications software company. The company develops commercial and open-source software solutions that are constructively disrupting the way that service providers build, scale, innovate, and account for communication services. By working with Metaswitch, visionary service providers are realizing the full economic, operational, and technology benefits of becoming cloud-based and software-centric. Metaswitch's award-winning solutions are powering more than 1,000 service providers in today's global, ultra-competitive, and rapidly changing communications marketplace. For more information, please visit [www.metaswitch.com](http://www.metaswitch.com).

## About Intel® Network Builders

Intel Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at <http://networkbuilders.intel.com>.



<sup>1</sup> Testing conducted by Metaswitch as of July 31, 2019. See Test Scenario for configurations.

<sup>2</sup> Figure courtesy of Metaswitch.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of July 31, 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be totally secure.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](http://intel.com).

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice Revision #20110804

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

Other names and brands may be claimed as the property of others.