Broadband Evolution: Wireline and Wireless Convergence on an Intel Cloud Native Platform

Authors

Paul Mannion
Intel Business Director for
Broadband and Converged Access

Andrew Duignan
Platform Solutions Architect for
Broadband and Converged Access

Table of Contents

Introduction

In 2017, the Broadband Forum began looking at how communications service providers might converge their fixed and mobile infrastructure to deliver multi-access connectivity (fixed and mobile) to better utilize their fixed and mobile assets.

This concept of wireless and wireline convergence was born out of the desire to utilize fixed and mobile assets more efficiently and to effectively provide multi-access connectivity for customers, as well as to reduce operational complexities across multiple network systems.

Another goal of this proposed transition is to provide better service experience and network optimization by embracing recent technology advances such as NFV, SDN and cloud-native network infrastructures, to enable a programmable, converged core allowing providers to reduce costs by simplifying the unifying architecture design of traditional network infrastructure and coupling services implemented on the same cloud native architecture.

This involved reviewing the current 5G system architecture and identifying the changes necessary to enable wireline access to a unified 5G core and then working with the 3GPP standards body to ensure that the necessary changes are standardized and released.

This approach specified several access architectural deployment options and the underlying infrastructure sharing aspects and provided new specifications for 5G Customer Premises Equipment (such as 5G-Residential Gateway), designed to integrate to the 5G core and provide true multi-access capabilities.

WWC Platform Architecture

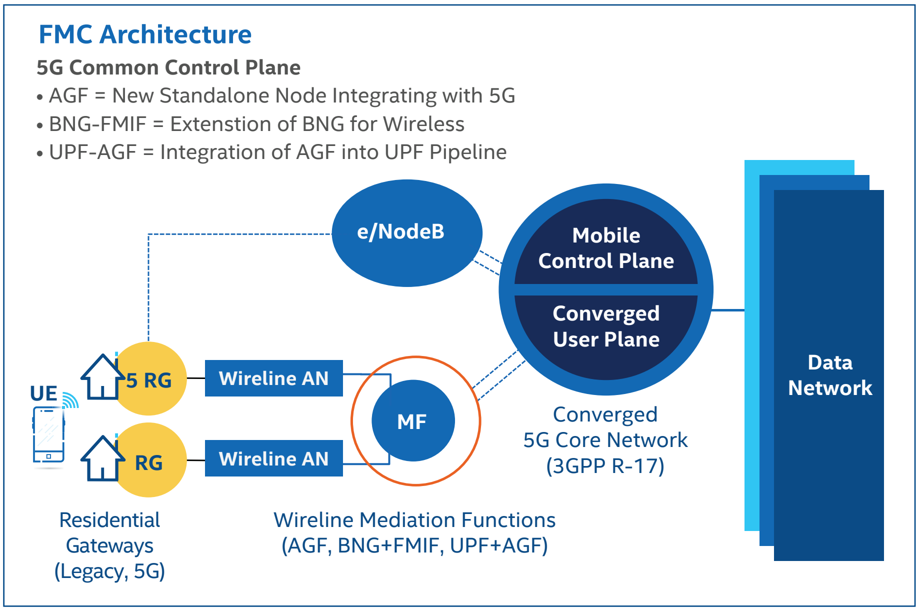

Figure 1. WWC Architecture.

The new WWC standards allow access connectivity via a legacy residential gateway (RG) or a new 5G-enabled RG. These gateways are deployed in the consumer premises or at the enterprise and enable wireline access over which 5G services can also be offered.

The mediation functions (MFs) shown in Figure 1 are used to enable the various WWC interworking options. The broadband network gateway (BNG) fixed mobile interworking function (BNG-FMIF) involves adding the requisite mobile user plane extensions into the legacy BNG. Operators who primarily own fixed assets and deploy or partner for a 5G control plane service will adopt this approach. This deployment typically happens in the central office or edge/metro location of the operator's edge estate.

It is also possible to deploy a standalone access gateway function (AGF) which mediates between the wireline access nodes and the mobile core user plane function (UPF). This makes sense for operators who have deployed a 5G core but wish to also offer new wireline services through a wireline access network. This deployment can be enabled at the central office, edge/metro or in the mobile location where the 5G UPF and core resides.

The final deployment option is the combined UPF/AGF. In this scenario, the AGF functionality is consumed or integrated into the 5G UPF. This type of deployment typically takes place in the 5G mobile core location.

In early 2021, Intel, in co-operation with Heavy Reading, carried out a survey among 80+ operators on their preference in approach for offering WWC services. The results vary quite dramatically, as shown in Figure 2. The survey also covered multiple systems operators (MSOs, cable operators) who primarily offer DOCSIS services. In their case, the WWC standardization work is taking place in CableLabs.

The variety in deployment options means that underlying platform flexibility is one of the key attributes required in WWC rollout.

As shown in Figure 3, the operator survey also revealed the operational (OpEx) and capital expenditure (CapEx) savings that the operators are hoping to realize from this convergence.

The key ambition from the operators' perspective is to realize commonality in the service assurance and security operational model (OpEx) and the underlying software stack for WWC (CapEx) implementations and also from the underlying hardware implementation platform (CapEx).

The remainder of this white paper will describe how Intel, through the adoption of the key tenets of cloud native and a common server-based hardware, is making these transformational gains possible.

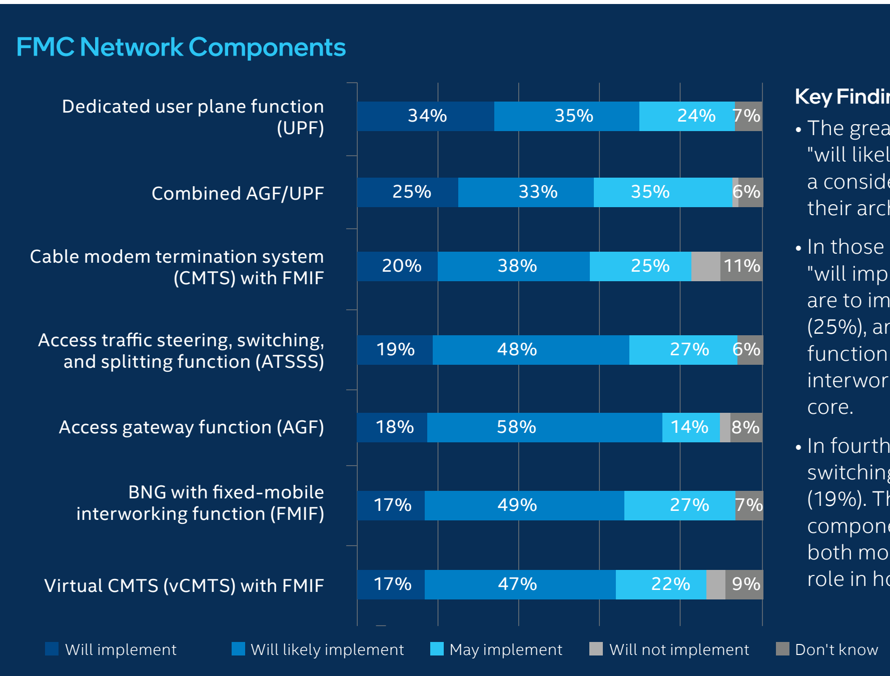

Key Findings

- The greatest range of inputs (33%-58%) fall into the "will likely implement" category, which reflects that a considerable number of SPs are still mulling over their architecture requirements.
- In those cases where SPs have opinions based on "will implement" inputs, the top three preferences are to implement a UPF (34%), a combined AGF/UPF (25%), and a CMTS with fixed-mobile interworking function (FMIF; 20%). The FMIF is required to interwork fixed line subscriber traffic with the 5G core.
- In fourth place was the 5G core access traffic steering switching splitting (ATSSS) component function (19%). The ATSSS is a strategically important component since it enables traffice steering over both mobile and fixed networks and therefore plays a role in how edge compute resources are utilized.

Question: How likely is your organization to implement the following FMC network components? [n=84-85]
Source: Heavy Reading

Figure 2. WWC Deployment Options. Results from a 2021 survey in cooperation with Heavy Reading.

FMC Savings

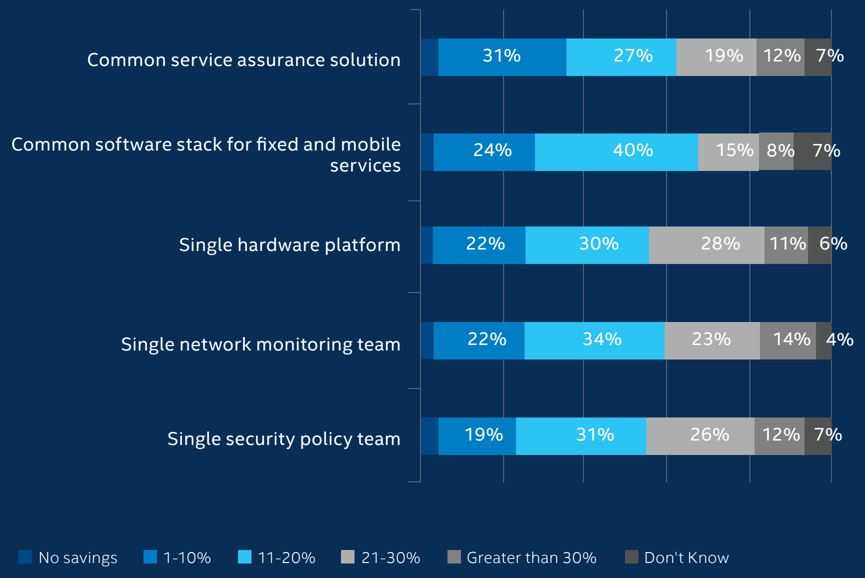

Key Findings

- SPs are forecasting considerable savings. On the upper end, the top three choices in the 21-30% saving range are a single hardware platform (28%), a single security policy team (26%), and a single network monitoring team (23%).
- These three also lead in the greater than 30% savings range, with common service assurance solutions scoring highly too (12%).
- This input reinforces that the ability to run a common network on common hardware with a single security and monitoring team is a game-changer from a savings perspective.

Question: What level of savings do you expect FMC to deliver within two years of commercial deployment for each of the following (n=84-86)
Source: Heavy Reading

Figure 3. Wireline and Wireless Convergence Benefits. Results from a 2021 survey in cooperation with Heavy Reading.

WWC and Intel's Platform Flexibility

By way of example, the WWC-defined Access Gateway Function User Plane (AGF-UP) routes the traffic from subscribers connected to the wireline access network to the 5G core network of communication service providers. It then implements functions such as policy rule enforcement and per subscriber service level agreements (SLAs), packet classification, header processing, transport encapsulation/decapsulation, hierarchical quality of service traffic management and more. The pipeline complexity depends upon the service provider requirements.

Intel has taken what it learned through its experience with cloud native deployments and adopted the Kubernetes container as a service (CaaS) platform to implement a number of the WWC gateways described above.

The intel AGF-UP reference implementation uses separate uplink and downlink packet processing pipelines. Each pipeline applies a set of functions to each packet that enters the pipeline. As shown in Figure 4, the uplink packet processing pipeline handles the packets flowing from the subscriber access network to the ISP 5G core network, while the downlink pipeline deals with the packets running from the core network to the access network. The average packet size of upstream traffic is typically smaller than for downstream, and the amount of upstream traffic is normally five to eight times less than downstream traffic. In recent years, the traffic gap between upstream and downstream has reduced significantly due to increased use of applications such as Instagram, Snapchat, TikTok, etc.

The Intel reference AGF-UP is developed using the performance-optimized Vector Packet Processing (VPP) framework, which leverages Data Plane Development Kit (DPDK) drivers to take advantage of high-speed I/O. The VPP is designed around an extensible and modular packet-processing graph architecture in which each independent graph node does limited packet processing on a vector of packets. The VPP framework allows application developers to plug in new graph nodes without changing core or kernel code to build customized packet processing solutions. The AGF uplink and downlink pipelines are implemented as separate packet processing graphs around a number of nodes where each graph node implements a specific packet processing function.

Each instance of AGF-UP is run as a Kubernetes cloud microservice. The instance runs inside a Kubernetes pod, within which there is a container running the VPP pipeline for both the uplink and downlink traffic (separate pipelines running on separate CPU Hyperthreads). Kubernetes is used to orchestrate the deployment of these AGF instances across the user plane server nodes.

The reference AGF-UP application has a compile-time switch so that it can be compiled as an AGF-UP, or alternatively, as a legacy BNG-UP. The same code base and software package is used for both and demonstrates the flexibility of the approach while using the same user plane software stack and underlying hardware platform.

Uplink Packet Processing Pipeline

The reference implementation of the AGF uplink packet processing pipeline consists of the following functions constructed as a directed graph of VPP nodes, as shown in Figure 4:

- **Packet Rx (Receive):** Packets from the access network are receieved from the NIC ports using DPDX poll mode drivers (PMD) and sent to the next stage to begin packet processing.
- **Access Control List (ACL):** This stage employs an access control list (ACL) table to implement firewall policies, such as block rules, on the incoming traffic. A table lookup operation is performed on each received packet, and in case of rule match, the packet is dropped.
- **Flow Classification:** Exact-match classification is performed on the 5-tuple header fields of the input packets (source and destination IP addresses, source and destination UDP/TCP ports and transport layer protocol ID) to identify the session. The session information is stored as packet metadata to be used later in the pipeline, and access network encapsulations (QinQ + PPPoE header) are stripped off the packets.
- **Metering and Policing:** This function meters the subscriber traffic flows to determine compliance with a service contract and applies traffic policing to enforce the contract. As a result, packets that conform to a specified rate are sent to the next stage of the pipeline while packets that violate the rate are dropped.
- **GTP-U Encapsulation and Routing:** Packets are encapsulated with a GPRS Tunnelling Protocol User Plane (GTP-U) header at this stage and routed to the 5G core network through the correct network interface port.
- **Packet Tx (Transmit):** With the help of DPDK poll mode drivers, packets are transmitted out of the system through the NIC ports connected to the 5G core network.

Downlink Packet Processing Pipeline

The reference implementation of the AGF downlink packet processing pipeline consists of the following functions constructed as a directed graph of VPP nodes, as shown in Figure 4:

- **Packet Rx (Receive):** Packets from the 5G core network are received from the NIC ports using DPDK PMD drivers and sent to the next stage to begin packet processing.
- **Access Control List (ACL)**: This stage employs an ACL table to implement firewall policies, such as allow rules, on the incoming traffic. A table lookup operation is performed on each received packet, and in case of rule match, the packet proceeds to the next stage.
- **Flow Classification:** The 5G core network encapsulations (IP + GTPU Header) are stripped off the packets and exact-match classification is performed on the 5-tuple header fields (source and destination IP addresses, source and destination UDP/TCP ports and transport layer protocol ID) to identify the session. The session information is stored as packet metadata to be used later in the pipeline.
- **QinQ Encapsulation and Routing:** Packets are encapsulated with a QinQ (IEEE 802.1ad) header and routed to the access network through the correct network interface port.
- **Hierarchical QoS Traffic Management:** Each packet runs through a hierarchical QoS (HQoS) scheduler to ensure that thousands of subscribers can get the desired broadband capacity as per the service contract. It supports scalable five-level hierarchical construction of traffic shapers and schedulers to allow finer-grain traffic control than a traditional single-level QoS scheduler.
- **Packet Tx (Transmit):** With the help of DPDK poll mode drivers, packets are transmitted out of the system through the NIC ports connected to the access network.

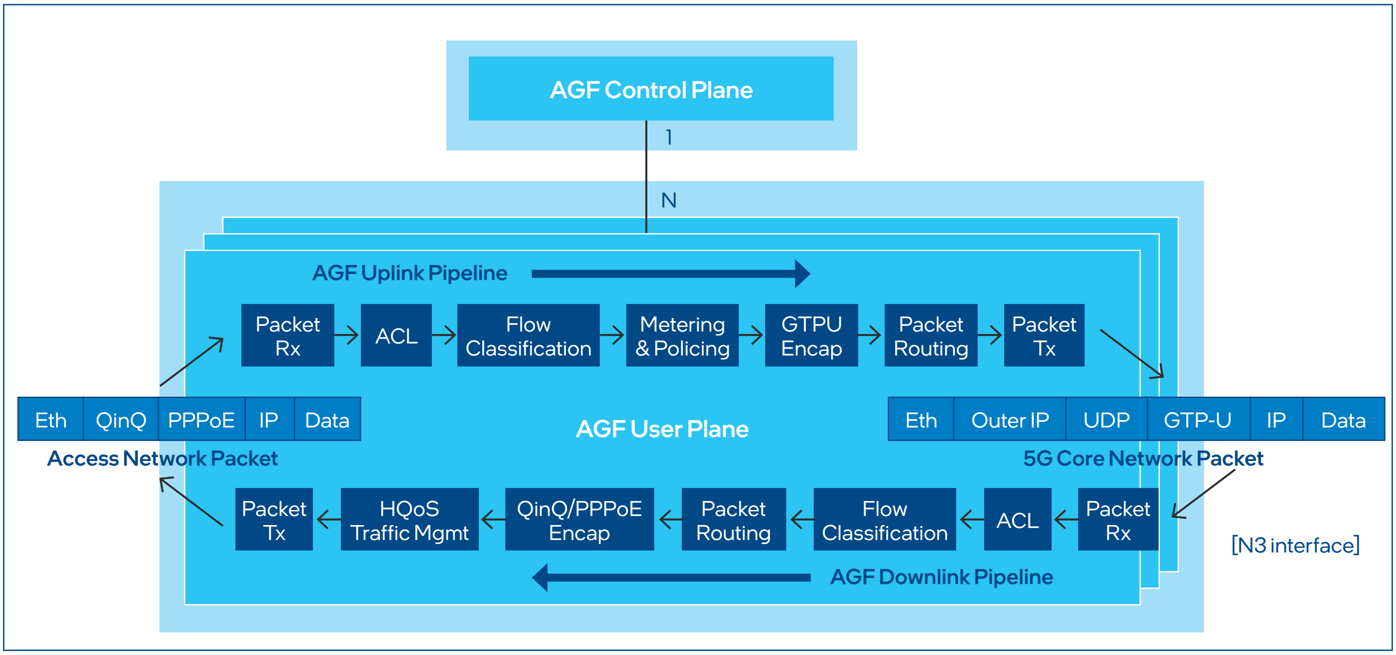

Figure 4. AGF Uplink and Downlink Pipeline

Cloud Native Deployment Approach

The AGF deployment strategy fully embraces the architectural and operational principles of cloud-native networking, where a singular data plane networking entity in the stack is called a Cloud-Native Network Function (CNF). A deployed CNF generally should follow these conventions:

- **Highly Performant:** The CNF must take advantage of Enhanced Platform Awareness (EPA) features to ensure low latency and high throughput.
- **Agile Placement:** The CNF must allow for flexible placement to deploy on any EPA-feature-ready platform and must be generic to the underlying EPA infrastructure.
- **In-Service Lifecycle Management:** Using automatic telemetry-aware controllers, the CNF must ensure that it can scale resources under increasing workloads and retract resources under decreasing workloads.
- **Highly Available:** The CNF must be highly available and fault-tolerant to meet the service level agreement of near-zero downtime. The CNF must also utilize the HA schema to maintain interfaces for quick and simple service upgrades without affecting the service it provides.
- **Observability:** The CNF must ensure that all network and performance metrics of workloads are exposed through an easy-to-consume platform, allowing for rapid network debugging and modification.

As is common with many cloud native deployments, the AGF uses the Kubernetes container-orchestration engine to ensure that all the CNF conventions above are met.

To ensure the methodology of the Control and User Plane Separation (CUPS) architecture and simplify the management of each infrastructure-agnostic microservice, the AGF stack makes a macro-level segmentation between the CNF infrastructure and the AGF application. Within the AGF application segment, the deployment is again segregated into functional sections based on their role in the CUPS architecture:

- Control plane
- User plane management
- User plane

From the perspective of a network operations engineer, the AGF cluster is deployed and maintained from a secure admin portal. This remote configuration is enabled through the use of Helm Charts. For Day-0 app deployment, Helm Charts are rolled out on top of the requested edge cloud infrastructure.

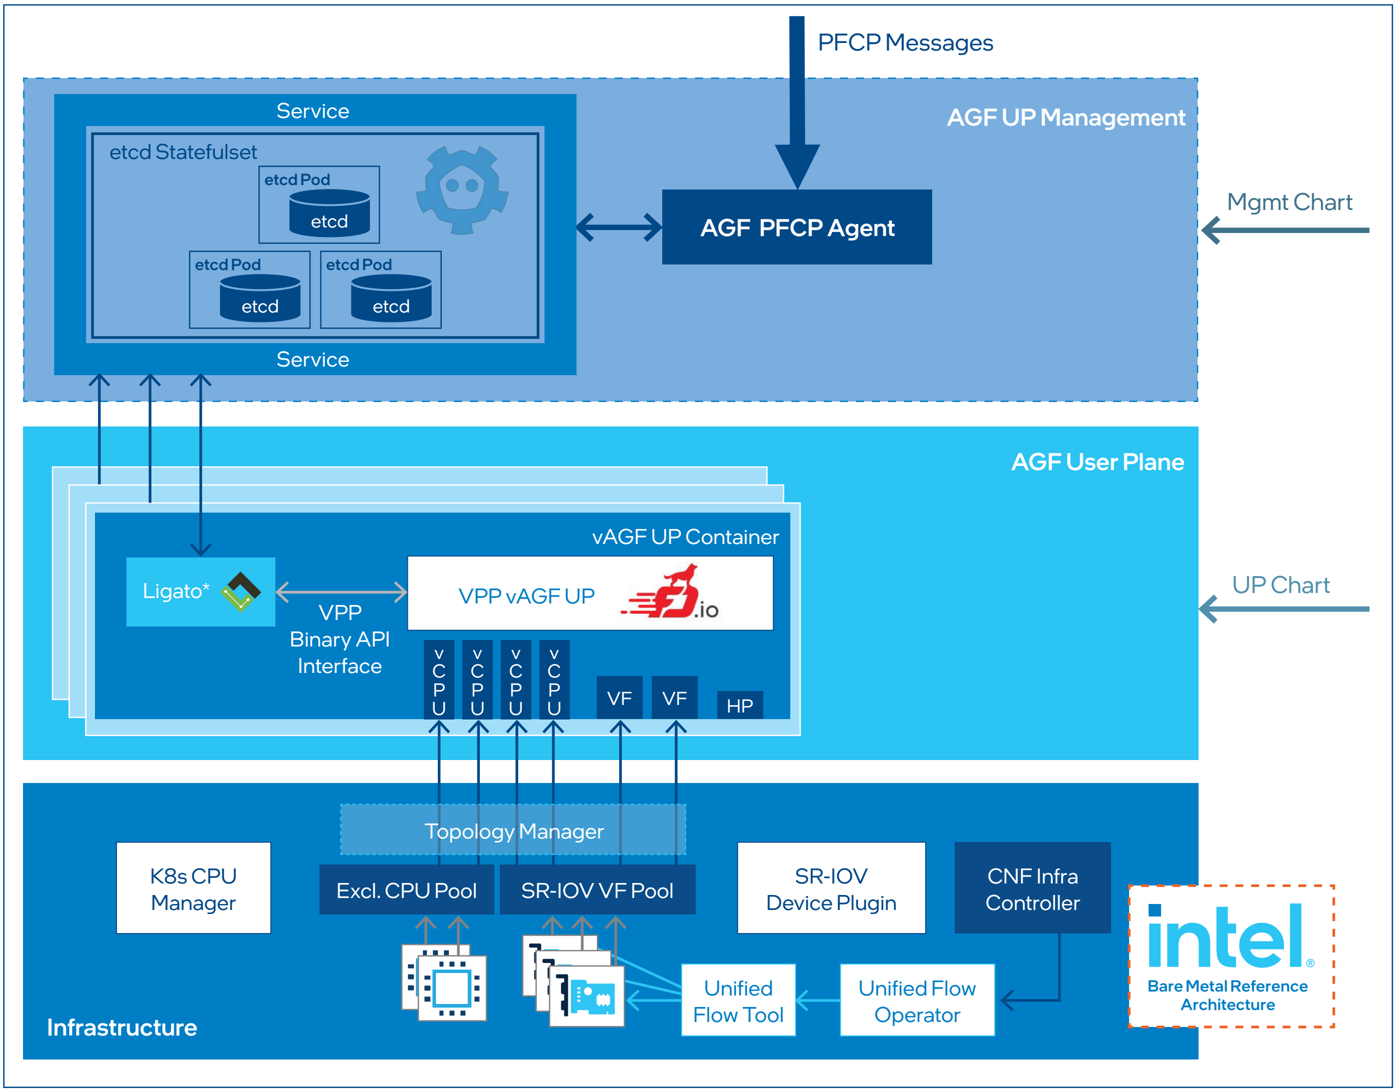

Figure 5. Detailed view of AGF UP deployment

The network operations engineer is also able to monitor and change lifecycle elements of the AGF cluster using Helm Charts and through user-friendly interfaces.

The steps of deployment for the AGF cluster are as follows:

1. The network operations engineer requests the cloud platform infrastructure that will host the AGF cluster using an edge cloud of their choice. In the AGF reference application, the default infrastructure deployment is handled by the Container Bare Metal Reference Architecture (BMRA). Note that by design, the AGF application is agnostic to the underlying infrastructure and thus does not require a specific infrastructure deployment platform such as the BMRA.
2. Once the infrastructure is in a ready state, the network operations engineer will deploy each of the application-plane charts.
3. Once all the application elements are running, the network operations engineer will configure the transport network on the 5G core and access side to route subscriber traffic to their assigned AGF instance while monitoring application and infrastructure metrics to ensure that peak performance is in line with the service level agreement (SLA). The end state of this deployment can be seen in Figure 5.

How Intel NiCs enable WWC

Receive Side Scaling (RSS) is a network driver technology that enables traffic streams to be distributed among hardware queues on a network card. This enables multiple CPU cores to poll the hardware queues separately, which is more efficient than requiring a single CPU core to poll all the NIC traffic and distribute it among the remaining CPU cores.

Dynamic Device Personalization (DDP) enables the Intel® Ethernet 800 series network card to parse headers on a number of communications protocols. It is possible to steer traffic to specific NIC hardware queues based on this. Under the default configuration (default DDP package), the Intel Ethernet Controller E810 can support traffic steering on packet headers that are common to most networking domains. To support custom packet types, a Dynamic Device Personalization (DDP) package can be loaded for immediate use without reloading the Ethernet controller non-volatile memory (NVM) image, as shown in Figure 6.

For the cloud native AGF function, it is necessary to process GTP-U encapsulated packets on the core network side (as the packets are being sent to/from the 5G core UPF). On the access side, the AGF must process PPPoE packet headers, which have QinQ format VLAN tags (C-Tag and S-Tag). This necessitates that the NIC can steer the incoming packets

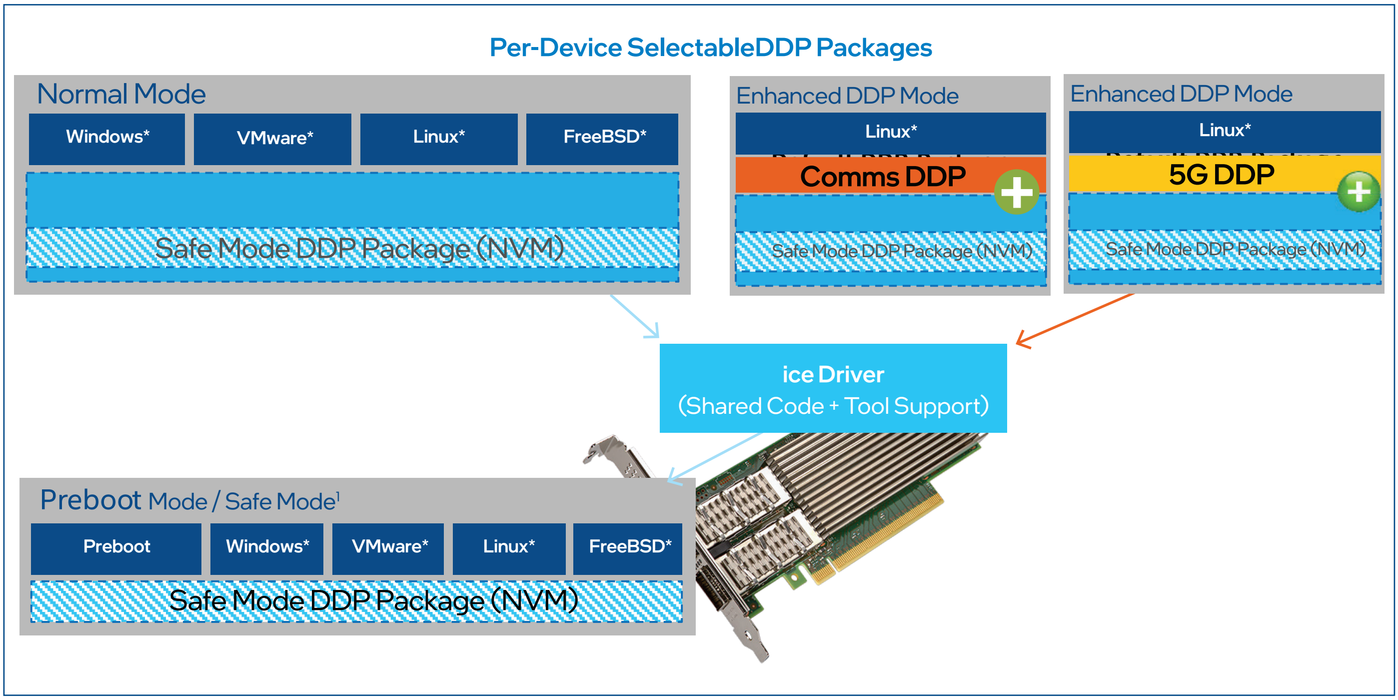

Figure 6. E810 Network card DDP packages

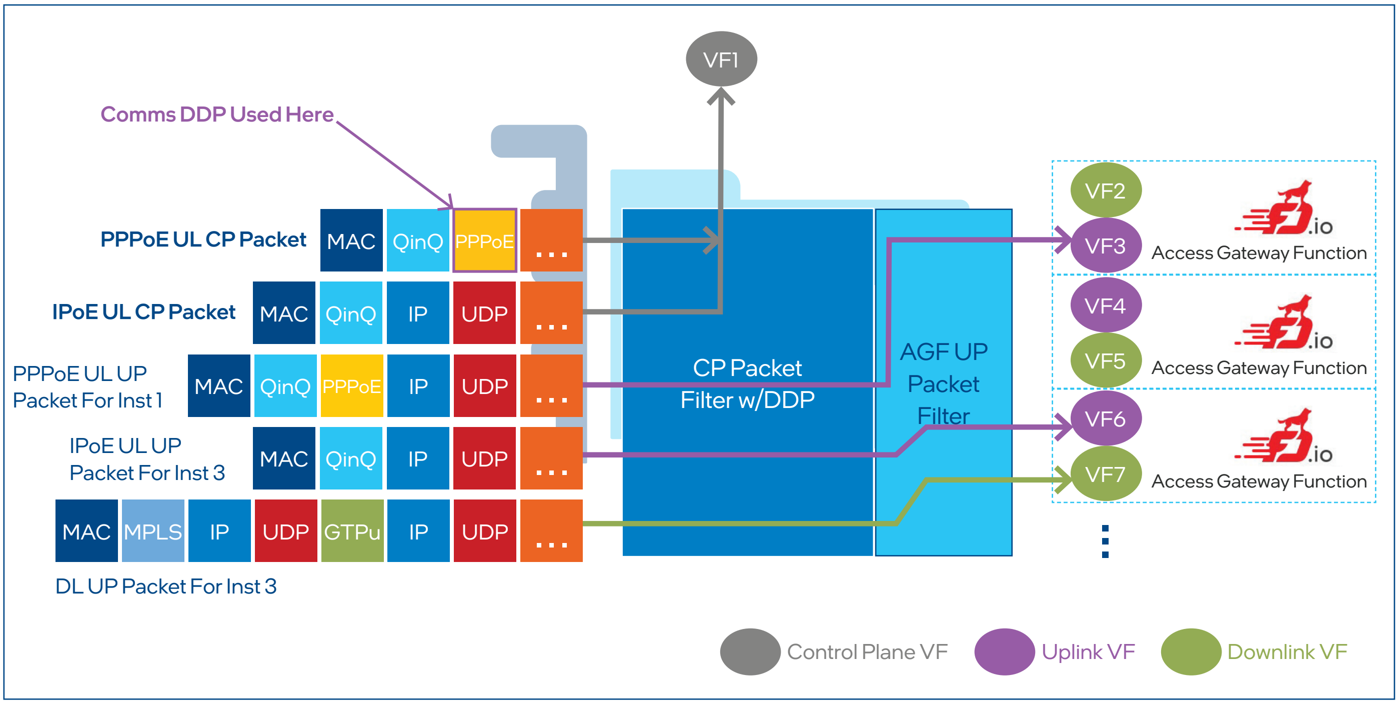

Figure 7. Example of packets being routed to appropriate VF by DDP packet filter

to the appropriate NIC virtual function (VF). Each AGF instance is running as a container within a Kubernetes pod. The container is assigned a VF for each pipeline (uplink and downlink). Each of these containers within the Kubernetes pod is pinned to specific CPU cores. The uplink pipeline is pinned to one hyperthread, and the downlink pipeline is pinned to two hyperthreads. These three hyperthreads are running on two physical CPU cores. Each AGF instance has its own dedicated CPU core resources. In this cloud native environment, each AGF instance operates without resource contention to another instance.

The packet steering in the NIC optimizes packet throughput (rather than using a CPU core to distribute traffic to the appropriate AGF container instance).

In the context of the AGF deployment, the Telecommunication (Comms) Dynamic Device Personalization Package is used. Once added, this package allows the Ethernet controller to steer traffic based on PPPoE header fields, thus supporting control plane offloading described in the Device Config Function. Figure 7 provides an example of what the AGF can achieve with NIC traffic steering and DDP.

Device Config Function

The AGF uses the SR-IOV and DDP functionality on the Intel Ethernet Controller E810 to complement the Control and User Plane Separation (CUPS) architecture. The AGF uses traffic steering on the NIC to perform two macro level functions:

- Route control plane discovery and session packets directly to the control plane (VF1) without the cost of going through the data plane (CPRI, known as Control Plane Redirect Interface).
- Route user plane packets to the correct vAGF instance running on an SR-IOV VF on that Physical Function (PF) (physical port).

To program these rules, the AGF uses the Device Config Function (DCF), available on the Intel Ethernet Controller E810. With DCF, DDP rules can be sent to the NIC controller which ultimately programs different hardware offload elements for WWC wireline and wireless protocols into the Intel Ethernet Controller E810.

Intel Cloud Native Reference Architecture Enables All WWC Mediation Function Models

The Container Bare Metal Reference Architecture (BMRA) is a CaaS-based experience kit which Intel maintains and releases to partners. It ensures quick and scalable MF (AGF, BNG, BNG-FMIF) deployments and helps in applying Intel's best-known configurations to harness the power of servers with Intel® Xeon® processors and Intel® Ethernet Controllers for optimized telco data plane performance in line with industry performance and flexibility expectations for WWF MFs.

The BMRA is the default platform infrastructure component for the AGF as it contains an Intel architecture-optimized reference profile dedicated to the AGF/BNG deployments. The BMRA represents a baseline configuration of components that are optimized to achieve maximum system performance mainly for cloud native and CNF use cases when running on Intel® Xeon® Scalable Processors.

For more information on BMRA, visit: https://networkbuilders.intel.com/intel-technologies/container-experience-kits

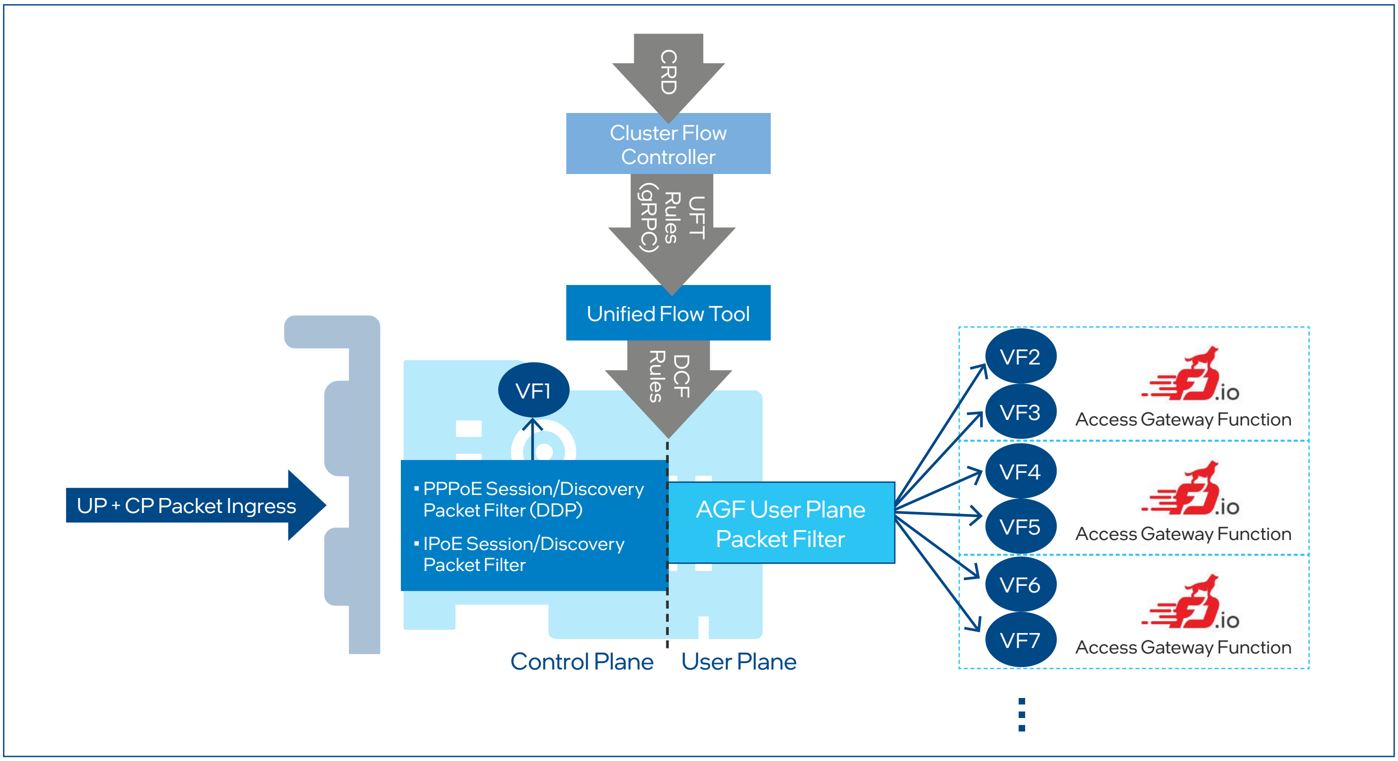

Figure 8. Kubernetes Ethernet Operator

Using this cloud native approach, the achievable performance is quite impressive, as shown in Figure 9. On the Intel Xeon Scalable Processor 6338N, the performance scales with 2 core instance. Running 30 instances (60 cores) across two CPUs gives an aggregate performance of about 500 Gbps.

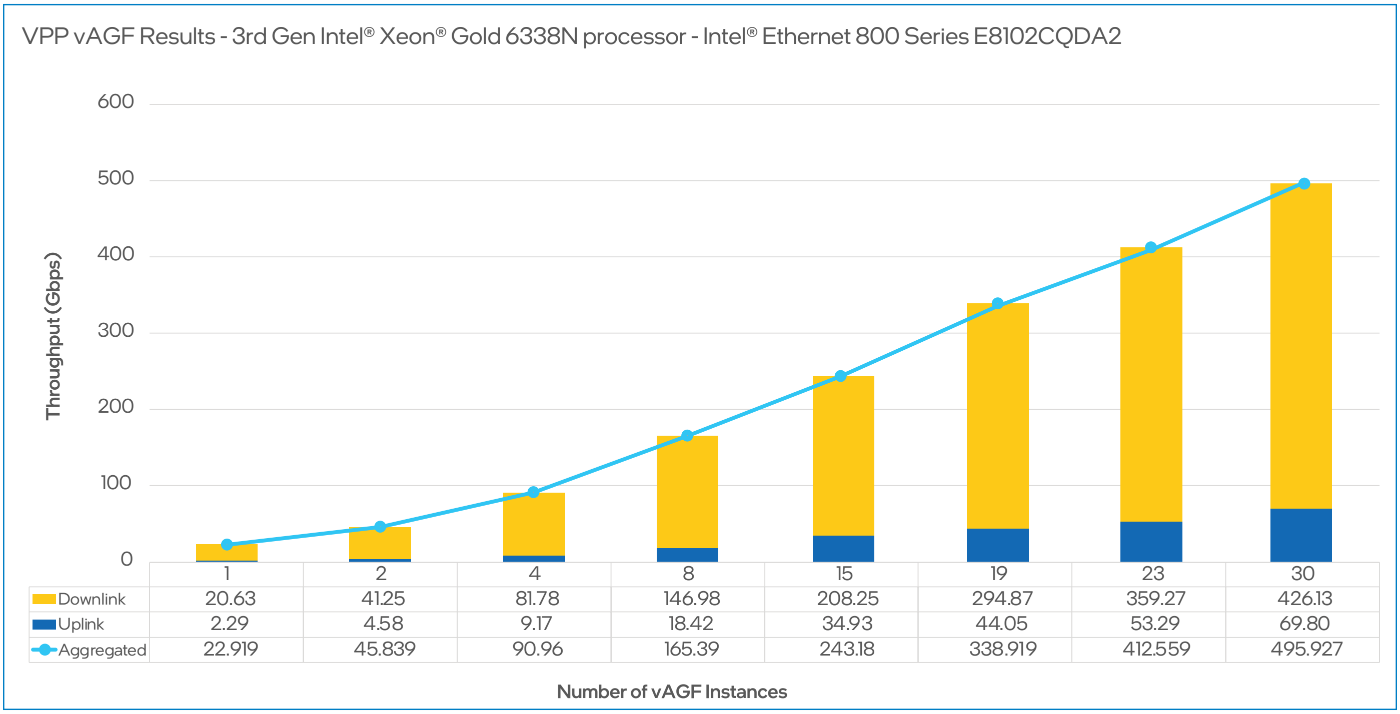

	1	2	4	8	15	19	23	30
Downlink	20.63	41.25	81.78	146.98	208.25	294.87	359.27	426.13
Uplink	2.29	4.58	9.17	18.42	34.93	44.05	53.29	69.80
Aggregated	22.919	45.839	90.96	165.39	243.18	338.919	412.559	495.927

Figure 9. Cloud native AGF performance

Conclusion

The drive to a wireline and wireless converged access architecture is well underway. The standard bodies of BBWF and 3GPP have collaborated together closely and the necessary standards were in place from 2022. What's now required is for the ecosystem to deliver the products and platforms required to release the operational and capital expenditure savings expected by operator adoption of Wireline and Wireless Convergence.

Intel has been working for two years with the ecosystem to deliver these new capabilities on a cloud native platform. The breath of the WWC solution choices will drive a new need for a flexible approach akin to those realized in the recent 5G core platforms through the NFV adoption period. The additional complexity in WWC is the added performance requirements, driven by the traditional wireline access speeds and associated QoS services which must now be addressed by the WWC platforms.

Intel has enabled all of the necessary building blocks from the latest generation Xeon SP processors, which unleash these high-performance capabilities, to its investment in public CaaS reference architectures which expose and implement the hardware features necessary to enable the flexibility and convergence benefits required in unleashing the expected operational and capital savings. By adopting a Kubernetes CaaS approach, the platform unleashes the commonality in the service assurance, deployment, and security models and also enables the software commonality through the compile time implementations of the WWC user plane gateways.

0922/BB/MESH/PDF 352868-001US