White Paper

Wireline Networks
Telco Edge Cloud

Transforming Fixed Access Using WT-474 Traffic Steering on a Cloud-Native Architecture

Authors

Jonathan Newton,
Vodafone Distinguished Technical Expert, Vodafone

Paul Mannion,
Wireline and Converged Access Business Director, Intel

Table of Contents

Subscriber traffic steering at the telco edge cloud

Edge computing is not new, but it has taken on new importance with the digitization of operator networks. The associated opportunity depends on the efficient steering of enterprise or subscriber traffic to the correct edge application or service. Operators and new independent software entrants are exploring innovative use cases and services. The operator community is developing new architectural approaches through standard bodies such as the Broadband Forum (BBF) to enhance their current fixed-access network architectures and take advantage of these rapidly emerging opportunities.

The industry is converging on defining the "edge" as locations with a maximum round trip time (RTT) to the end user of 20 milliseconds (ms).[1] This degree of access latency accommodates major edge use cases such as augmented/virtual reality (AR/VR), edge video analytics and vehicle-to-any (V2X) communications. Edge service locations will include the following:

- **Communication service provider (CoSP)-operated sites** including central offices and regional data centers (DCs) or leased space in colocation or neutral host-provided DCs.
- **Cloud service provider (CSP)-operated sites**, including leased space in colocation provider DCs.
- **The enterprise edge**, including branch offices, industrial sites, regional DCs and leased space in colocation provider DCs.

As the industry begins to explore different use cases, it becomes increasingly clear that the latency and bandwidth advantages bestowed by the fixed-access fiber edge will be critical to the rollout of high-value edge services such as edge internet of things (IoT), AR/VR and video analytics. These attributes make the CoSP-operated edge critical to the edge service-delivery value chain.

The pandemic has made robust home broadband more important than ever, with remote work and school from home putting new demands on fixed-access networks. In 2020 to 2021, fixed traffic grew at an unprecedented 40% compound annual growth rate.[2] During the same period, there were also dramatic changes in the symmetry of the traffic, with online collaboration driving dramatic growth in uplink traffic. Prior to the pandemic, the typical uplink-to-downlink ratio was 1:8; during the pandemic, it reached as low as 1:5,[3] placing new demands on the fixed-access network architecture.

The larger CSPs are also on the march. Amazon launched Outpost, Microsoft launched Azure Stack for network and Google has launched Google Anthos for the edge. These approaches are based on the premise of deploying edge compute platforms which are connected to and compatible with the edge cloud strategies offered by the CoSPs. New application and service outcomes coupled with high quality of experience will play increasingly important roles compared to traditional bandwidth speed tests when it comes to selecting a broadband supplier.

CSPs have distinct advantages from economies of scale for service creation, but they lack insight about traditional fixed-access consumer behaviors and their needs. That disconnect makes it difficult or impossible for CSPs to adapt or readjust access technologies to offer more bandwidth, lower latencies or different mobile and fixed-access solutions depending on the applications or services offered. CoSP expertise provides an advantage. The landscape is similar for enterprise use cases. Edge IoT will demand more flexible, application-specific provisioning, which can be enabled through new fixed-access technology and a complimentary evolution in broadband standards.

The way forward is to merge technical and operational approaches to offer both cloud applications and network services on the same platform, in a variety of edge locations. This goal is possible using the distributed cloud central office (CO) technologies from the broadband forum coupled with new application-specific steering capabilities that will allow the operator to automatically provision application-specific characteristics through the fixed-access network.

With modern application development cycles, the new digital fixed-access network must be able to rapidly adjust to new application demands in real time. The implication is that the classic timeframes to develop new network-based services are no longer viable. New edge traffic-steering and deployment methods must be adopted in network design and deployment for next-generation edge service enablement.

Figure 1 illustrates how operators are now looking at deploying access user plane technologies in different parts of their edge networks depending on the services offered. This is enabled by the adoption of Control and User Plane Separation (CUPS) as well as new wireline and wireless convergence technologies such as the Access Gateway Function (AGF), which enables core services through a fixed-access network.

Upgrading a traditional fixed access network can be a time consuming and difficult process. These networks have often been deployed with statically configured connectivity between the customer and the broadband network gateway (BNG) that connects the customer to the operator's network. This static provisioning means that maintenance activities such as network upgrades and adding new capacity are often executed late at night and with long planning and deployment cycles that often involve "man in the van" services or forklift upgrades. Consequently, networks are often overprovisioned to minimize maintenance requirements, resulting in power, capacity and space inefficiencies.

The BBF Working Text 474 (WT-474) offers an evolutionary path for fixed-access architecture that enables continuous integration and continuous deployment (CI/CD) development cycles and zero-touch upgrades. Live subscriber sessions are dynamically moved without impacting service, and the resources consumed by the system can scale up and down to meet current traffic demands and power usage constraints.

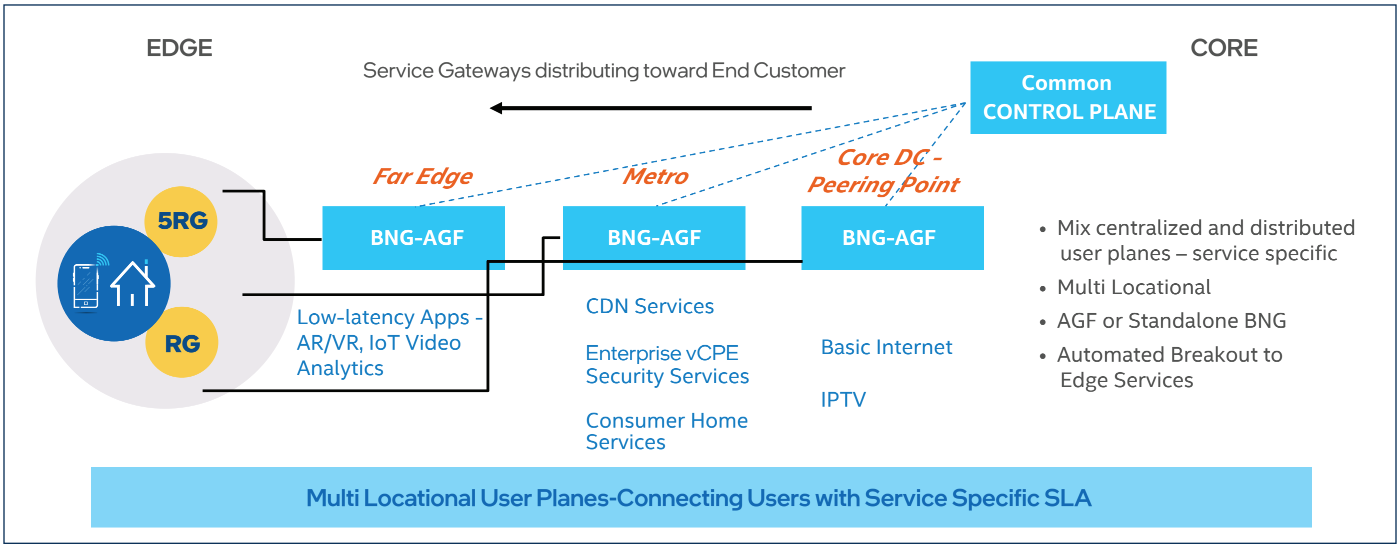

Figure 1. Multi-location user planes and services.

Network convergence is a significant consideration, as represented in Figure 2. Since 2017, the BBF and 3GPP have been collaborating on a series of standards that will enable fixed and mobile convergence in the coming years. In effect, the fixed-access and mobile-access networks will converge. The evolution of the fixed-access network to a more mobile-centric control platform means that fixed-network operators must embrace the same cloud-native approaches that are being deployed for 5G today.

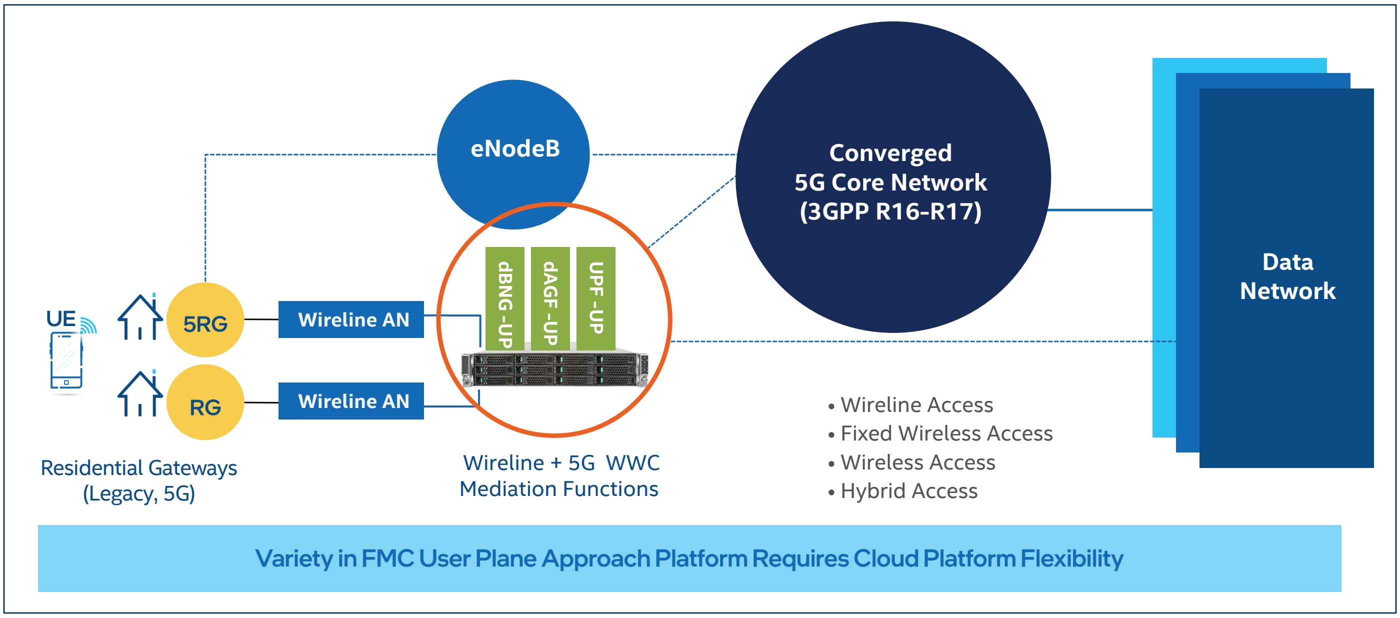

Figure 2. Fixed and mobile convergence architecture.

This will result in a platform that allows them to deploy the correct gateway technology (UPF, AGF, UPF-AGF, BNG-FMIF) depending on the connectivity environment or country in which they are operating.

Architecture fundamentals of fixed-access edge traffic steering

The architecture shown in Figure 3 is a simplified view that integrates the key work being done by the BBF in its Cloud Central Office, Disaggregated BNG and Subscriber Session Steering projects. The architecture seamlessly enables the evolution to greater fixed/5G network convergence as the service gateway can also be an AGF, as defined in the BBF Wireless Wireline Convergence work.

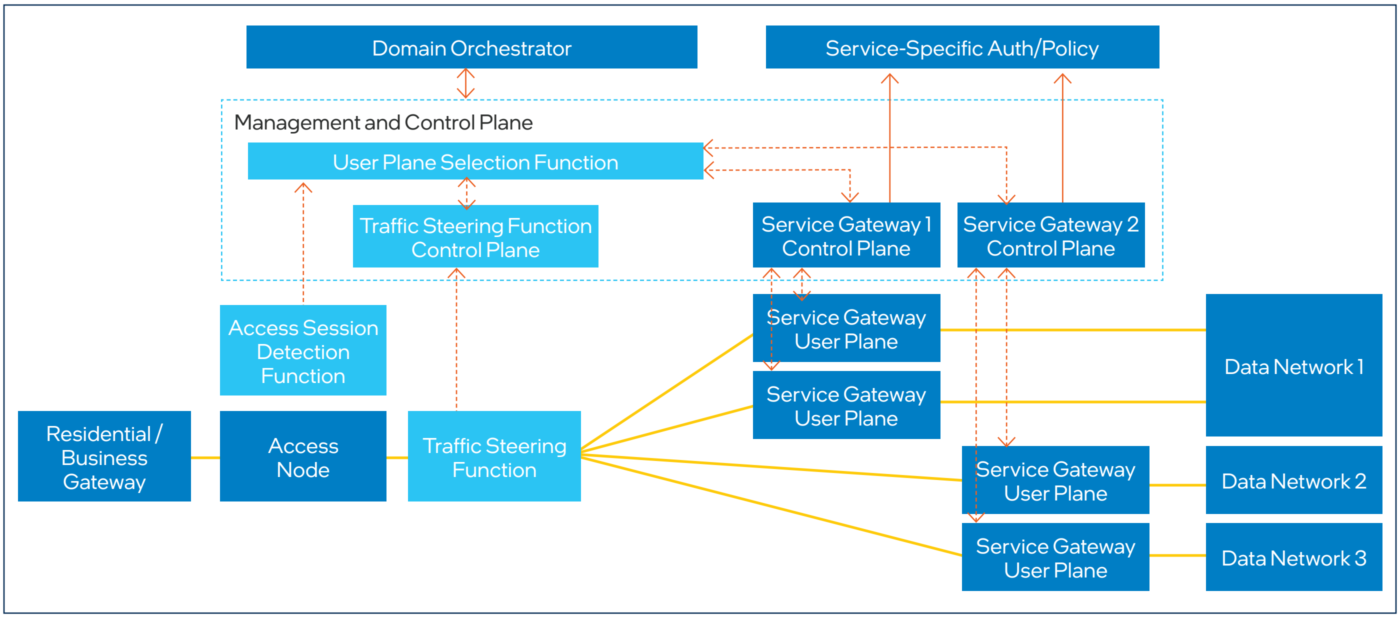

Figure 3. Fixed-access telco edge traffic steering architecture.

A key advantage associated with this change in architecture is that it moves away from the traditional statically provisioned connectivity between a subscriber and services. Instead, the network provides a dynamic ingress user plane, service selection and session load balancing that simplify network operations, improve resilience and ensure that subscriber sessions are connected to service gateways that can meet customer service requirements. In particular, this architecture provides low-latency edge services and optimized user plane resources to match dynamic traffic requirements.

The architecture also enables a cloud-native and disaggregated implementation, where the functions are deployed in software on cloud infrastructure (in geographically separate locations), enabling improved resiliency and more rapid development and deployment of new capabilities and connected services. The major components in this new architectural approach include the following:

- **Service gateway (SG)** is a generic name for the function responsible for providing the required network to the subscriber and providing access to application services. Examples of a service gateway include BNG, AGF and Provider Edge Router (PE). The service gateway is decomposed into control plane and user plane functions based upon Control User Plane Separation (CUPS) protocols.
- **Service gateway control plane (SG-CP)** is the control component of the service gateway, with each control plane capable of controlling many service gateway user planes. It is responsible for functions such as authenticating subscribers, allocating IP addresses and taking user planes in and out of service.
- **Service gateway user plane (SG-UP)** is the user plane component of the SG and is responsible for forwarding traffic and providing access for the user to the required network and application services. User planes will be deployed in both edge and core locations, allowing services to be provided with appropriate latencies to deliver the required end-user application experience. Furthermore, user planes may be implemented as hardware or software functions as required to meet the flexibility and traffic forwarding requirements of the operator.
- **Traffic steering function control plane (TSF-CP)** is responsible for identifying the user plane to which any one subscriber should be connected. Key to this is the User Plane Selection Function (UPSF) which is queried whenever a new subscriber session is brought up, identifying the service gateway and user plane that can meet the subscriber's service and latency requirements while maintaining balanced load across the domain. The UPSF is also responsible for proactively and reactively identifying any required change in service gateway and user plane mapping in response to network changes and maintenance activities.
- **Traffic steering function (TSF)** is responsible for directing the packets for a particular session to and from the correct user plane. This is a relatively simple cross-connect function that can be built into the physical access node, or an aggregation switch or router.
- **Access Session Detection Function (ADSF)** is the function that recognizes that a new session is active and needs to be connected to the correct service gateway and user plane. The first sign of life for a new session will depend upon the type of subscriber session but will be activity such as a port coming up on the physical access node or session request packets received from a home or business site. As with the traffic steering function, ASDF may be provided by the physical access node, or by any other element that can recognize a new session first sign of life.
- **Access Node (AN)** is responsible for terminating the fiber or copper access connections from homes and businesses.

What has changed? Supporting demonstration and major use cases

To support an architecture proof of concept, Intel, Vodafone and BISDN built the functional lab prototype illustrated in Figure 4. Intel and Vodafone worked with BISDN to provide the BNG Session CP (SG-CP) extensions to enable interactions with the user plane selection function (UPSF), which was designed and built by Vodafone.

Intel also provided the traffic steering function (TSF) implemented on an Intel® Tofino™ 64 x 100G ports P4 programmable switch using the P4 programming language, which enabled a programmatic and flexible approach to access traffic classification. The Intel Tofino switch directs the traffic related to a specific subscriber context from the access network to the specific BNG user plane instance identified through the steering process and in the reverse direction from the BNG user plane to the access network.

The demo was executed in Intel's lab and recorded for the BBF website in 2021, then was extended for the key uses cases described below in 2022. The demo showed that some of the issues identified in the introduction can be addressed and solved using this new flexible approach. This approach enables interesting new user cases that give operators greater flexibility in deploying services and provide significant operational benefits.

The first demo use case addressed shows how the full system, including all the key control plane and user plane components shown in Figure 4, can be instantiated on a Kubernetes edge cloud. Subsequently, new subscribers are added to the system. The major difference here is the involvement of the UPSF, which the SG-CP queries to understand current SG-UP loading and uses that information to decide which SG-UP a new subscriber should be connected to. This approach gives the operator dynamic control over resource usage with respect to quality of service and minimizes the impacts of an outage, where an overloaded SG-UP goes out of service and disconnects active subscribers. This new provisioning approach is more dynamic and cloud-like, using the same approaches used today in many 5G core deployments.

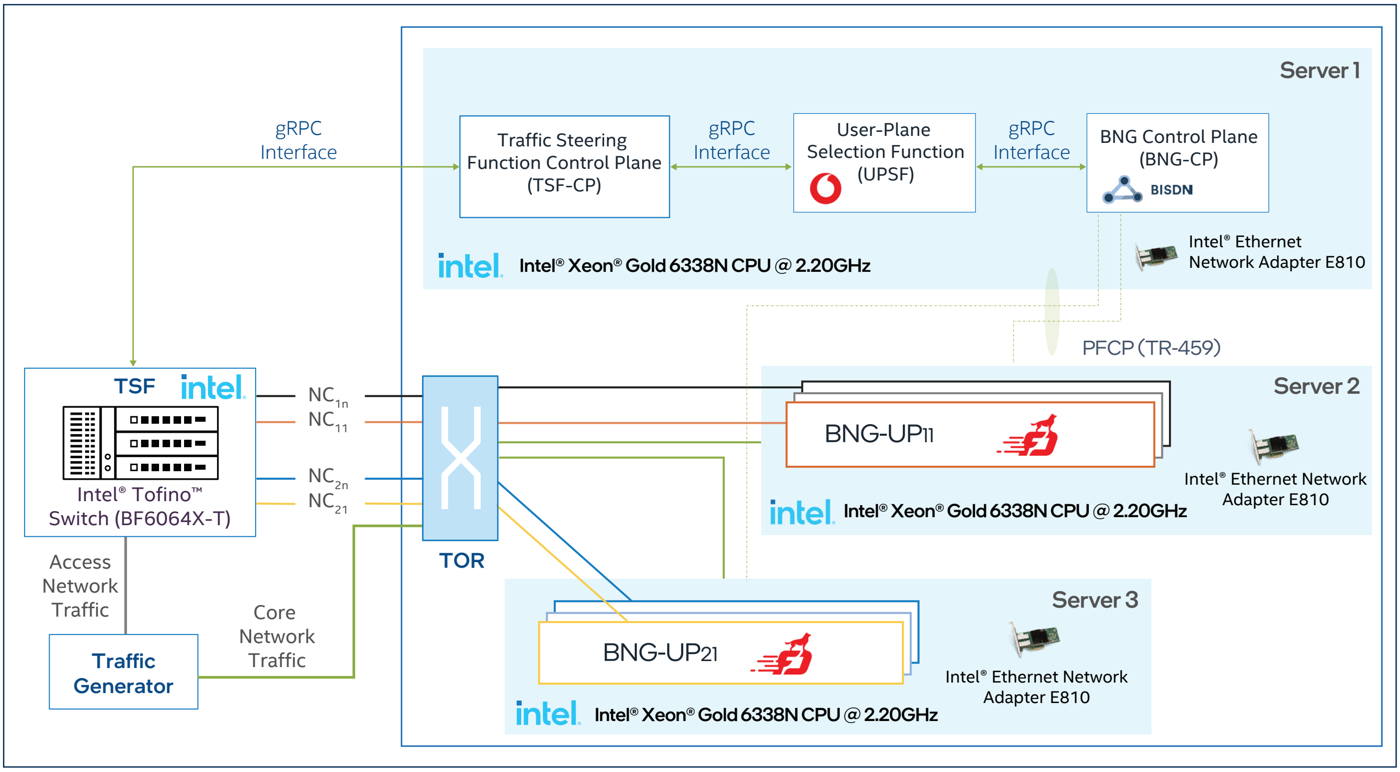

Figure 4. Cloud-native edge traffic steering reference architecture.

The use case was about service-based selection and demonstrates how the system can dynamically connect subscribers to new edge services. In this scenario, the SG-CP creates a linkage in the UPSF that contains a list of service group IDs and the SG-UPs that can provide access to these services. The new-session setup is similar to the sequence described above, but in this case, the operational control (e.g., Radius server) will include additional information about the services (service group IDs) required by the new subscriber.

Again, the UPSF decides, based on current loading, which SG-UP to connect the subscriber to and enables access to the requested edge services. This allows the operators to dynamically create and assign new SG-UPs based on the characteristics of the new services being offered. The access network becomes more flexible and service-aware, allowing the operators to match their network provisioning and spend with revenue-generating services.

The second use case addresses field maintenance: the removal of an in-service SG-UP or the upgrading of a new SG-UP to a later software version containing new features or enhancements. This frequently occurs when operators need to roll out new features (e.g., IPv6 services) or bug fixes that require a truck roll or significant downtime late at night. A group of subscribers treated in a similar manner is termed a “subscriber group.” The operator will initiate a SG-UP deletion request and the UPSF then identifies another in-service SG-UP that can support the currently active subscribers. Then the operator, via the correct SG-CP, will install the existing subscriber states onto the newly selected SG-UP. The UPSF then notifies the TSF to redirect the affected groups to the newly selected SG-UPs and reconfigures the downstream traffic toward the end user.

The final use case in the demo addresses green strategy and power optimization. The Kubernetes Horizontal Pod Autoscaler is configured to periodically check the load on each of the in-service SG-UPs and takes time-of-day traffic heuristics into account. At busy hours, the HPA will scale out SG-UP instances to accommodate peak-traffic demands. Conversely, as homes “switch off,” the HPA can re-balance the subscriber shards and scale in the SG-UP nodes, turning off the underlying compute resources to reduce power consumption.

To enable this new flexibility and scale, the SG-UP nodes are implemented in a cloud-native microservice fashion, as shown in Figure 5. This architecture allows the user planes to be deployed onto multi-locational Kubernetes clouds, sized and scaled appropriately for the throughput and latency needs of the services hosted at these locations.

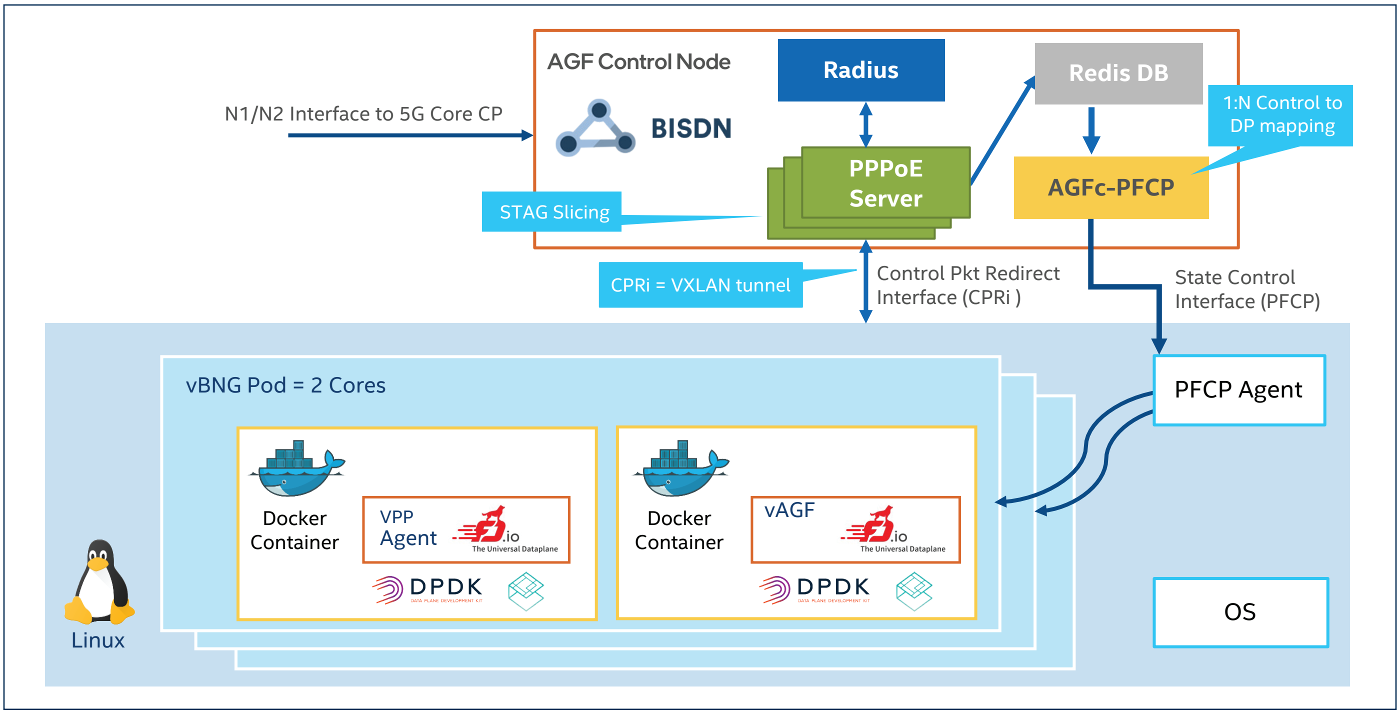

Figure 5. Cloud Native BNG

This cloud-native user plane architecture is similar to that in 5G implementations; it implements each BNG microservice in software using the vector processing (VPP) technologies available in the FD.io project. Each BNG-UP is implemented as a two-instance docker container Pod, each of which consumes two cores on the Intel Xeon® Gold 6338N processors.

For the BNG application, the telecommunication (Comms) Dynamic Device Personalization (DDP) package is used with the Intel Ethernet Network Adapter E810. This package enables the Ethernet controller to steer traffic based on PPPoE header fields so the NIC can route packets to specific virtual functions/queues based on the unique PPPoE header fields, namely the protocol ID.

The cloud BNG-CP instances used in the demonstration were developed by the Berlin institute of Software Defined Networking (BISDN). The cloud BNG-UP instances were based on a reference architecture developed by Intel. Figure 6 shows the throughput of an Intel Xeon processor-based server based on two Intel Xeon Gold 6338N processors running vBNG container instances.

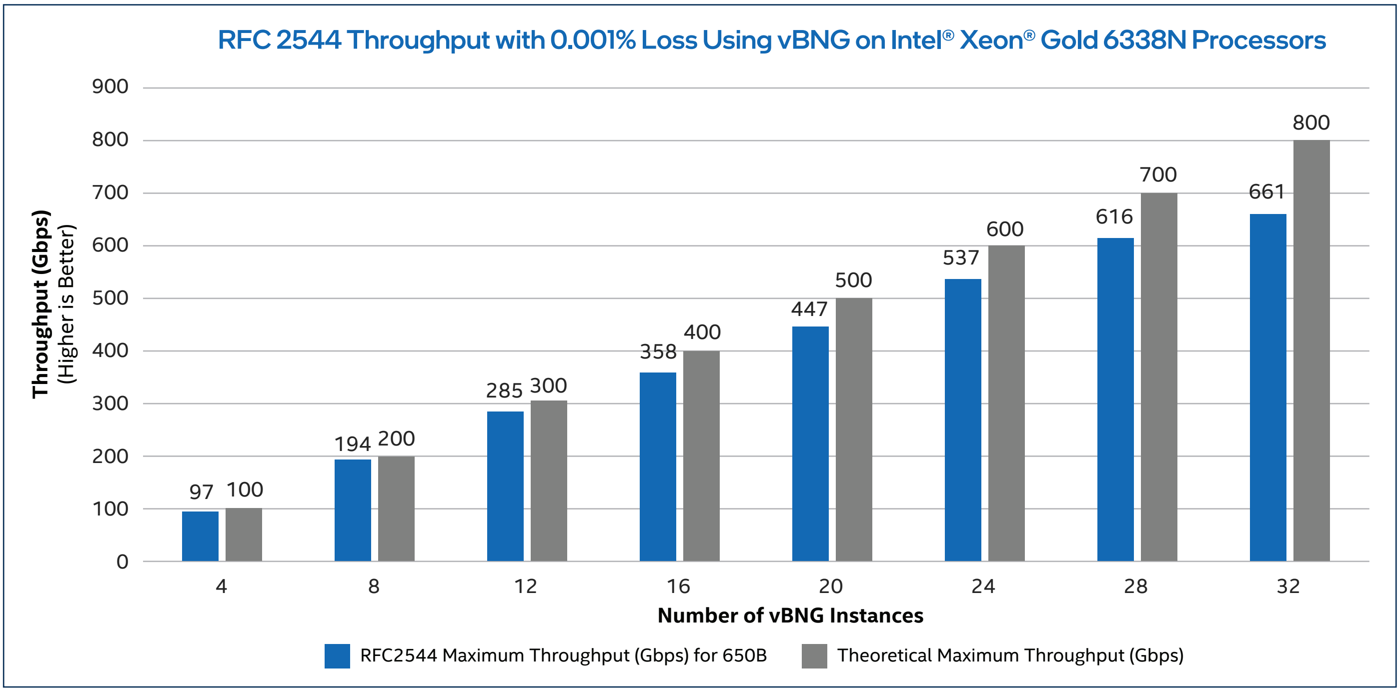

Figure 6. Cloud-native BNG-UP scaling.

The throughput scales linearly as the deployment increases from four to 32 vBNG instances in increments of four instances. With the full 32 instances deployed, the throughput is 661 Gbps using RFC2544 test methodology with 0.001% packet loss. This result is achieved using 96 data processing cores (1.5 cores per instance for 32 instances). All resources used by the BNG application are local to the socket. It is found to be I/O bound but not CPU bound. The flexibility of the approach works well for edge traffic steering on a multi-location architecture. Further details about the key use cases in the demonstration videos are described in the following sections.

Seamless subscriber group steering

The first use case advancement demonstrated is user plane maintenance, including BNG upgrades and security/bug fixes. Correcting these issues can be resource-intensive and costly, often requiring truck rolls or off-hours work. Traffic-steering enhancements and cloud-native infrastructure make this process much easier, even enabling real-time, in-service upgrades. The process follows the make-before-break sequence in Figure 7.

In the BBF demonstration, there are eight active cloud-native user planes. User plane 5 is marked for maintenance by the UPSF. The active subscribers on this user plane are moved to one of the other seven active user planes. The TSF-CP reconfigures the upstream path to the newly selected user plane, and the BNG-CP reconfigures the downstream direction. All active sessions are maintained. User plane 5 is then removed from the Kubernetes Pod for maintenance and is re-introduced when subsequently upgraded.

Resource scaling and power management

In this use case, the cloud BNG user plane uses the Kubernetes Horizontal Pod Autoscaler (HPA) to modulate the user plane resources required based on time-of-day traffic, as illustrated in Figure 8.

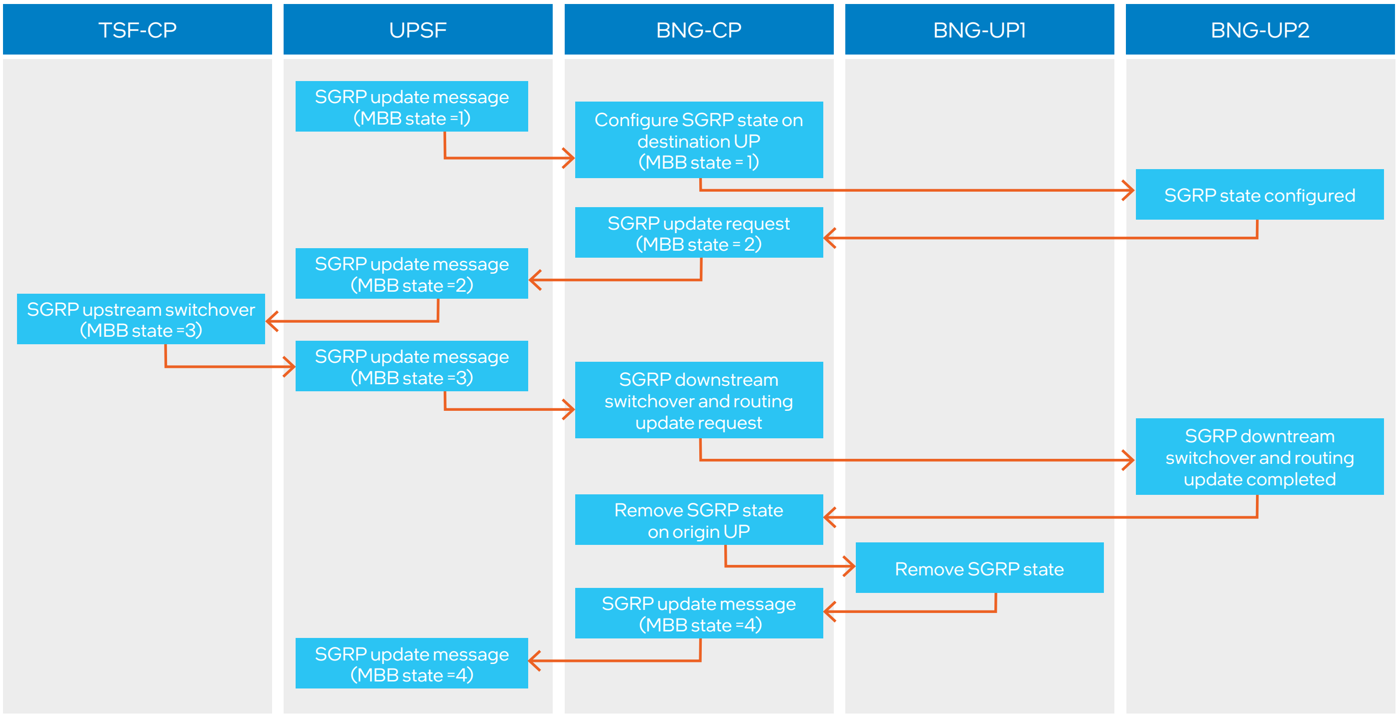

Figure 7. Make-before-break subscriber session management.

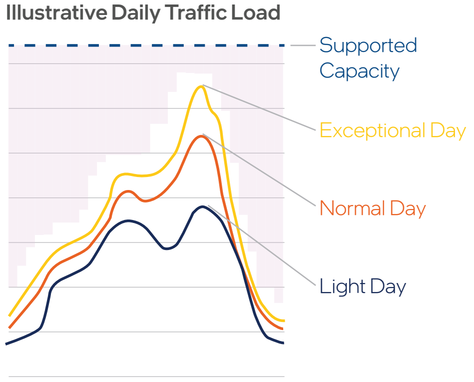

- Broadband service providers need to build capacity above the exceptional peaks in demand.
- Daily average usage (especially on light days) may approach 50% of an exceptional peak.
- Trend toward mass viewing events such as sports finals on OTT streaming video services can increase the normal-to-exceptional ratio.
- Session steering and Control-User Plane Separation together enable active sessions to be moved between user plane instances.
- This unlocks the capability to reduce power consumption by progressively disabling and powering down unused servers, CPUs and cores.
- Machine learning may be applied to predict demand and proactively disable or enable capacity as required.

Figure 8. Auto-scaling the BNG user plane to optimize power consumption.

Broadband providers today must build their networks to take peak usage into account, handling mass-viewing events such as the World Cup although average usage may be just 50% of this peak. Using appliance- and ASIC-based user plane technologies, the network is deployed statically for peak capacity, and the power and operational expenses required remain constantly at these exceptional peak levels.

Dramatic fluctuations in energy costs since the Summer of 2021 make this peak-sizing approach problematic because of its unpredictability. For example, the average wholesale electricity price in Germany was €175.45 per megawatt-hour in November 2022, roughly the same as a year prior but down from a record of over €469 in August 2022.[4]

In contrast, when using WT-474 traffic steering and a Kubernetes cloud-native infrastructure approach, the HPA is configured for average traffic thresholds. The BBF demo meets this average user plane need using three Kubernetes BNG user planes, each capable of 25 Gbps. When this average threshold is exceeded, the Kubernetes HPA scales out another BNG user plane, and the UPSF rebalances subscriber groups from the currently overloaded user planes to this newly created BNG Pod.

In the demo, the scale-up parameter is set to 8, meaning the network is capable of scaling to 8x25 = 200 Gbps of capacity and the scale-down parameter is set to 1 = 25 Gbps. The opposite is also true: when traffic begins to fall below the configured threshold, the HPA works in tandem with the UPSF to reduce the number of required user planes, rebalance subscribers and scale in the system to support the network traffic. Scaling down BNG user planes to the minimum appropriate level can provide significant cost savings and environmental benefits by reducing power consumption.

Rolling in-service upgrades

The final use case demonstrated is when an operator, due to bug fixes, new features etc., wishes to update their estate of in-service BNG user planes. The use case demonstrates the real power of cloud services and CI/CD. The Kubernetes deployment controller performs automated software rolling upgrades of BNG Pods without any manual intervention, as illustrated in Figure 9.

During the update process, a new software manifest file is made available, and old BNG Pods are replaced by new Pods adopting this new software, all while ensuring service is maintained. Two parameters control the rolling update: the maxSurge parameter is configured to 1, meaning we have one new user plane at any time and can upgrade one in-service BNG user plane at a time. Following the make-before-break sequence, the older user plane is sent a “drain and delete” message, and the subscribers are moved from that user plane to the newly configured one via the USFP and TSF. The process continues in this manner until all old versions of the BNG user plane have been removed and the newly introduced ones all run the latest BNG software version.

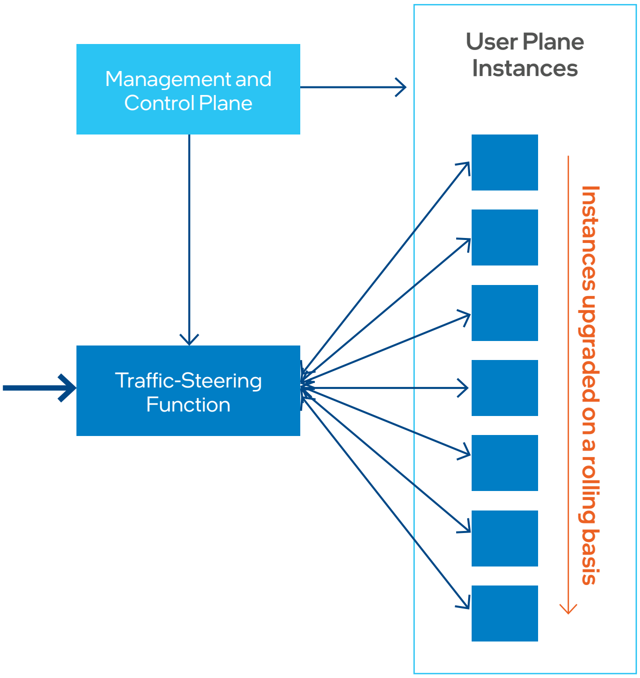

- The stateful nature of broadband subscriber sessions currently makes it challenging to upgrade traditional broadband gateways.
- Steering enables the user plane estate to be upgraded on a rolling basis without any impact to customers.
- Service providers can now resolve issues and add new software features at any time, without negatively impacting customer services.
- Enables a continuous deployment approach where automated deployment of new user plane software can be achieved as an integrated part of a rapid development cycle.

Figure 9. Cloud service-based hitless rolling upgrades.

Conclusion

Fixed-access service providers have a huge opportunity to differentiate and service-enrich their networks by moving away from the traditional static mapping of subscriber to service. Doing so will allow individual subscriber sessions to be steered to the right location to support their service needs, including appropriate latency for edge applications.

At the same time, a cloud-native approach enables the platform to adopt new CI/CD-like software capabilities into the network, removing the need for long planning upgrade cycles and minimizing network outages. Cloud-native architecture also allows the sizing of network resources for traffic levels, helping to reduce operating costs and support corporate climate initiatives.

The BBF WT-474 defines a dynamic fixed-access architecture that supports cloud-native principles to enable the network to scale up or down in rapid response to customer traffic load or service needs. It also enables new features to be rolled out in a CI/CD approach without costly and time-consuming outages for maintenance.

From an individual subscriber perspective, the network can now dynamically connect their sessions to a service gateway at the right location to meet their application requirements, including emerging applications that can benefit from deployment at low-latency edge-compute locations.

Appendix

Table 1. vBNG server.

Platform	Intel® Server System M50CYP Family
CPU	2x Intel® Xeon® Gold 6338N Processor, 2.2 GHz, 32 Cores
BIOS, Microcode	SE5C6200.86B.0020.P24.2104020811, 04/02/2021 , 0xd0002c1
Memory	16x 32GB DDR4
Hard Drive	Intel® SSD DC S4600 Series SSDSC2KG96 (960 GB)
Network interface Card	4x Intel® Ethernet Network Adapter E810 -2CQDA2

Table 2. Software.

Host OS	Red Hat Enterprise Linux 8.2 (Ootpa)
vBNG	vBNG 20.11
Linux Container	Docker version 20.10.5, build 55c4c88
DPDK	DPDK-v20.11
BIOS Settings	P-state disabled, Intel® Hyper-Threading Technology enabled, Enhanced Intel SpeedStep® Technology disabled, Intel Turbo Boost Technology disabled, C-States disabled, SR-IOV and VT-d enabled

Table 3. Application configuration per instance.

Uplink	**Frame Size**	650 bytes*
	Subscribers per Instance	4000
	vCPUs per Instance	1
	ACL	Blacklist with 150 rules
	Flow Classification	Flows classified on vlan tag pair
	Policer/Metering	Two-rate three-color marker
	Routing	Single forwarding rule
Downlink	**Frame Size**	504 bytes*
	Subscribers per Instance	4000
	vCPUs per Instance	2
	ACL	Reverse path forwarding — one rule per subscriber (4k)
	HQoS	4 Level HQOS — port, pipe, traffic class and queue
	Routing	One route per subscriber (4K)

* Frame size quoted is max size of frame at any point in processing (e.g., uplink 128 bytes =120 bytes + {2x4-byte access vlan tags}).

Table 4. Test environment configuration information and relevant variables.

Traffic Generator	Ixia Novus 100GE8Q28
Connection Details	Ixia ports and DUT ports connected back-to-back (eight connections)

[1] GSM Association, April 26, 2019. “Cloud AR/VR Whitepaper.” https://www.gsma.com/futurenetworks/wiki/cloud-ar-vr-whitepaper/.

[2] Mike Robuck, Fierce Telecom, February 10, 2021. “OpenVault: Covid-19 pandemic drives 51% spike in broadband traffic in 2020.” https://www.fiercetelecom.com/telecom/openvault-covid-19-pandemic-drives-51-spike-broadband-traffic-2020.

[3] Rupert Wood and Jakub Konieczny, Analysis Mason, August 6, 2021. “Fixed network data traffic: worldwide trends and forecasts 2020–2026.” https://www.analysysmason.com/research/content/regional-forecasts-/fixed-network-data-rdfi0-rdmb0/.

[4] Bruna Alves, Statista, December 21, 2022. “Average monthly electricity wholesale price in Germany from January 2019 to November 2022.” https://www.statista.com/statistics/1267541/germany-monthly-wholesale-electricity-price/.

0523/BB/MESH/PDF 348767-002US