

The 100-Terabit Broadband Network Gateway (vBNG)



Table of Contents

Introduction	1
Cloud-Native BNG	2
Hardware-Based Versus Software-Based BNGs	3
Deployment Architecture and Benefits	4
Hardware and Software Disaggregation	4
Control and User Plane Separation (CUPS)	5
Cloud-Native Software-Based Solution	5
Discrete Network Design	6
Switch-based vs x86-based BNG ..	7
Performance Testing and Results Using an Intel-based Server	7
A BNG Delivering Over 100 Tbps of Throughput	9
Summary and Conclusion	10

Introduction

Service provider network architects and engineers are under intense pressure to transform their networks to meet the challenges presented by an ever-changing telecommunications business environment.

The growth in demand for bandwidth and broadband services by wireline and wireless subscribers continues unabated and this will accelerate as the network evolves and 5G network deployment expands and edge use cases and services become more prevalent.

The pandemic has exacerbated this traffic growth effect, and in 2021, fixed traffic growth outpaced mobile traffic growth for the first time.

Today's subscribers have a choice of services that can be delivered by over-the-top (OTT) providers and other new cloud entrants from outside the traditional telecom industry. This has rendered obsolete the historical "bundled services" model service providers previously relied on to increase average revenue per user (ARPU) and reduce churn. Further pressure is being applied by the pandemic-driven shift to working from home and home schooling where enterprise-grade solutions for security, reliability, and connectivity means residential subscribers require the same level of security and performance as would be expected from an enterprise service.

Many operators believe the topic of wireless-wireline convergence (WWC) is upon us and all architectural decisions for wireline network evolution must take this into account and not preclude architecture convergence between fixed and mobile access technologies.

How do network providers meet the evolving needs of customers, the increased demand for specialized services, and do so with a technical and economic model that allows them to future-proof their network architecture and grow their business while creating new shareholder value?

The answer is to invest now in a flexible network model that enables them to succeed today and evolve in the future. That model demands transformation at the service edge of the network. Historically, service providers have depended on specialized hardware to manage subscriber sessions and to deliver specific services. This was expensive from a capital and operational expense standpoint. Yet it has been the best solution for decades, until recently.

While the early days of virtualized solutions saw fast adoption in some areas, like data center computing, it was slow to be adopted in networking. It first became economical in mobile networks around 2015, and today most 4G and 5G core networks are virtualized – either with virtual machines or cloud-native implementations in the 5G timeframe. In contrast, the high bandwidth demands of fixed networks have constrained virtual solutions from being practical. Tremendous strides have been made with virtualized systems and in the ecosystems which support them such that those systems now have the performance required for edge usage to not only provide a more cost-effective solution, but also a myriad of other benefits.

Broadband Platform Performance Per Socket

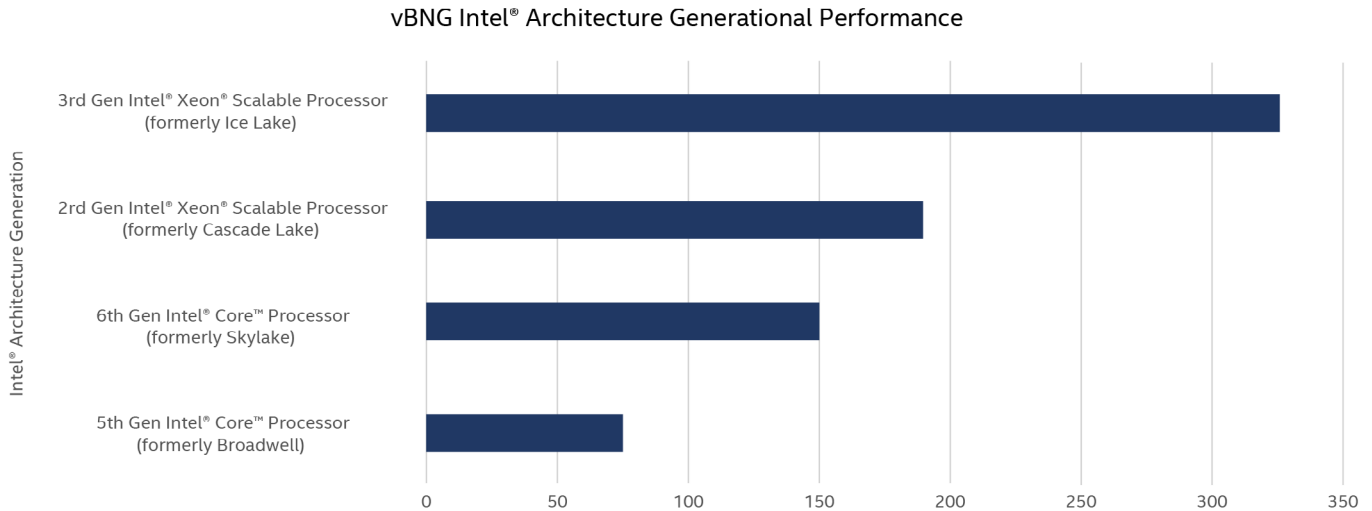


Figure 1. By investing in its CPU, Network Interface and easy to use open-source developer toolkits (DPDK) Intel has made huge strides enabling generational improvements in virtual BNG user plane performance. The latest 3rd Gen Intel® Xeon® Scalable processors are enabling up to 10 Gbps of BNG user plane performance per core. That is up to 320 Gbps on a single socket and twice that in a dual socket system.

The fixed broadband network has many critical components, but the linchpin of it all is the broadband network gateway (BNG). Many service providers are now deciding which path to take for their next-generation BNG. Should the industry move away from a proprietary hardware-based approach to a new open, cloud-native software-based approach using commercial off-the-shelf (COTS) hardware? The evidence points to yes. With a cloud-native BNG architecture, carriers can start to achieve improved economics and deliver new services at the speed enjoyed by the hyperscalers.

“Carriers are jointly cooperating on open architectures because they recognize the need for the industry to come together to advance flexible, agile, cloud-native BNGs so that they can grow the size of the telecom pie, rather than let other players capture all the value.”

- Michael McFarland, VP Product & Marketing, Benu Networks

Cloud-Native BNG

Transformation at the Edge

The Broadband Network Gateway (BNG) establishes and manages subscriber sessions by acting as the authentication point through which subscribers can connect to a carrier’s broadband network. The BNG aggregates subscriber traffic from the access network and routes it to the carrier’s edge service, core network and/or on to the Internet. All traffic

in-bound from the Internet passes through the BNG on its way back to the home subscriber.

Historically, a BNG performed primarily subscriber identification and authorization with the main service provided being basic Internet access. But the world has changed, and broadband-offered services have evolved to include streaming video, Voice over IP (VoIP), gaming, home office application access, and a multitude of other applications. This has not only increased the volume of transmitted information, but the traffic model has become very mixed and ever changing with varying protocols and performance requirements. This has driven new functionality requirements for Quality of Service (QoS), security and other critical functions. As 5G networks and services evolve, this trend will continue. Traffic patterns have also changed dramatically with streaming video often consuming 80% of the bandwidth. Hence, moving content delivery networks (CDNs) closer to the subscriber can dramatically reduce the bandwidth required in the aggregation and core network. Having compute closer to the customer also establishes a critical footprint for multi-access edge computing (MEC) to support low-latency applications for gaming, augmented reality (AR), virtual reality (VR), and numerous business applications like remote robotic surgery and other latency-sensitive applications. Consequently, a BNG platform must be able to scale up and scale down, support a variety of centralized and edge distributed deployment models (potentially concurrently), and even have different user planes tailored in their configurations to optimize for different types of services – basic Internet vs streaming video vs low-latency edge applications. This allows the operators to adjust their edge service architecture to align with new and better APRU opportunities. In general, a BNG today needs to deploy quickly, scale up/down rapidly, handle many different types of services, enable faster feature velocity, provide in-depth security, and deliver a great user experience.

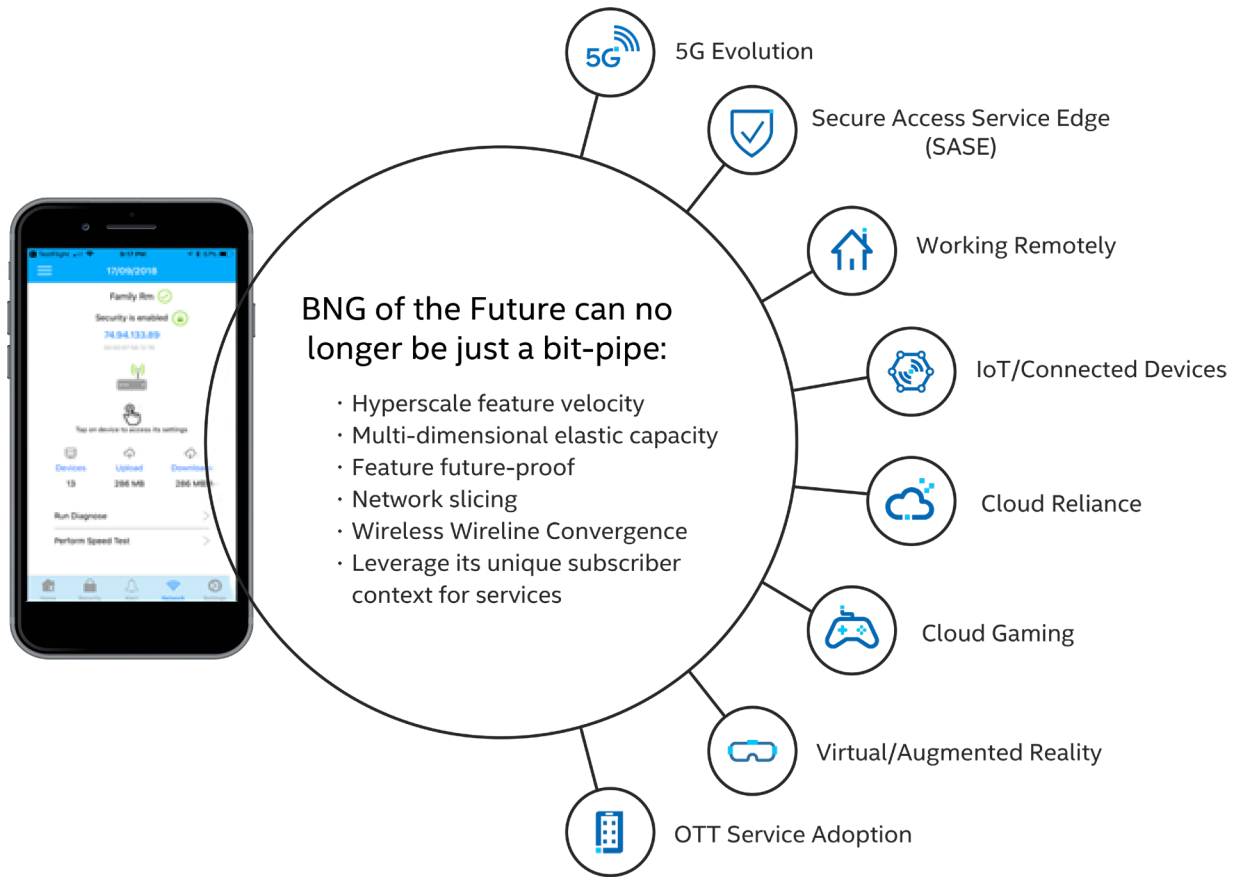


Figure 2. Key Market Trends Driving BNG Requirements. BNGs of the future can no longer just be a bit-pipe. The BNG has unique subscriber context and traffic visibility, enabling it to become a platform for innovation.

Hardware-Based Versus Software-Based BNGs

BNGs can be hardware-based or software-based. Hardware-based BNGs can be either built on proprietary ASIC-based vendor silicon from a telecommunications equipment manufacturer (TEM) or make use of merchant silicon from a third-party chip provider (e.g. Broadcom, Intel). In either case, the hardware-based switching platforms typically use Application Specific Integrated Circuits (ASICs) for the user plane, and these require ASIC-specific software development.

Porting software from one ASIC to another is poorly supported because different ASICs have different architectures, interfaces, and programming models. So, while merchant silicon is a step in the direction of hardware/software “disaggregation,” the user plane software is still very much tied to not only the chip vendor, but to a particular family of ASICs and programming (e.g. P4) toolkits.

Hardware-based BNGs are limited to what the hardware supports and many of the monolithic platforms make it very difficult to introduce innovation into the network from other vendors. This carries into many areas including the lack of an open hardware model to run other types of software, lack of an open software architecture for feature flexibility and extensibility and this slows the time to market for new services. Hardware-based deployment models come in a limited number of hardware configurations and capacities which means the decision to centralize or distribute the equipment becomes a long-term choice which is not easily

changed – the opposite of flexibility, and the ability to adapt the user plane characteristics to new services becomes more difficult or expensive.

Lastly, hardware-based BNGs do not support horizontal “scale-out” of additional compute within a single system, and “scale-up” requires either forklift upgrades or adding new IP nodes which impacts overall network configurations. And control and user plane scaling cannot occur independently (unless the control plane is pulled out of the user plane platform). Bottom line, scaling is a more complicated matter than it should be. As MEC use cases become more and more predominant, the ability to scale down and then scale out as the services take hold, becomes a more critical parameter in the decision choice for a BNG platform.

In contrast, a software-based BNG runs in a virtualized or cloud-native environment on COTS hardware using x86 general purpose processors (GPP). Because these chips are not “application-specific,” they provide complete flexibility as a true software solution. Hardware and software disaggregation is possible, with the ability to port software from one GPP to another. With the rise of new use cases like augmented reality, virtual reality, cloud-gaming, IoT, and others, there is a strong need to future-proof the BNG to support new service capabilities and this can be achieved by simply updating the software, without the constraints of any particular ASIC.

Attribute	Hardware-Based BNG	Open/Virtualized BNG
True Hw/Sw disaggregation	Not really - custom ASIC software	Yes - using general purpose processors (GPP)
Service Innovation	Difficult - closed architecture	Fast - architected for innovation
Feature Extensibility	Hindered - lack of open software model	Easy - open software model & APIs
Multi-Vendor Integration	Expensive and time consuming	Less complex - built on open standards
Scaling	Scaling can be complex (forklift upgrade or additional IP nodes)	Seamless scale-out by adding CPU & memory resources without impacting network. Control and user plane can be scaled independently
Deployment Options	Constrained based on limited box sizes and cost points	Unconstrained - any size system
Familiarity	High	Low - fixed network teams are new to virtualization (mobile teams more familiar)

Table 1. Attribute Comparison of Hardware vs Software-Based BNG [Note: Hardware-based BNGs are designed to move volumes of bits at a low cost. An open or virtualized BNG optimizes a more highly-functioning model of cost per subscriber. The primary drivers being the number of subscriber queues, schedulers, traffic shaping requirements, meters, counters, and logical interfaces.]

In addition, the separation of control and user plane becomes straightforward, enabling independent scaling of each with the ability to distribute the user plane without increasing network complexity. Finally, and recently, we have seen that the cost per subscriber is now actually lower than hardware-based BNGs.

Deployment Architecture and Benefits

Many large global service providers are supporters of the open BNG architectures defined by the Broadband Forum and Telecom Infrastructure Projection (TIP).

This architecture provides the option to separate the control plane from the user plane and run each in a virtualized environment using COTS components and open software. A disaggregated and open BNG allows operators a choice of different hardware platforms, a choice of different control plane applications and the ability to choose the type of Network Operating System (NOS) they want to use. Using a cloud-native architecture, carriers have a platform for innovation which can enable feature velocity comparable to what cloud hyperscalers have achieved. This ability to develop new capabilities quickly gives them a long-term competitive advantage, particularly as the competitive landscape broadens to fast-moving OTT players.

Furthermore, the cloud-native BNG architecture results in a low total cost of ownership which leads to a lower cost per subscriber and reduces dependencies on individual large suppliers. Disaggregation also provides service providers with more flexibility in service offerings allowing organizations to easily adapt to dynamic and changing business requirements.

These benefits apply to all communication service providers and are primarily derived from four cornerstones of the architecture:

- 1) Hardware and Software Disaggregation
- 2) Control Plane and User Plane Separation (CUPS)
- 3) Cloud-Native Software-Based Solution
- 4) Discrete Network Design

Each is examined below.

Hardware and Software Disaggregation

The disaggregation or separation of hardware and software delivers many benefits.

- It frees service providers from being locked-in to single-vendor hardware.

"We see the key benefits delivered by a cloud-native virtual BNG as being transformational to a service provider's business. We are working with service providers who are now transitioning their broadband access technology to the cloud BNG model."

-Paul Mannion, Wireline Access Segment Manager, Intel

- It runs on COTS server hardware which provides better flexibility and the ability to match system sizing with the bandwidth and session requirements driven by subscriber service needs.
- The separation of software drives innovation by allowing third party applications to be introduced without the limitations of the ASIC or hardware, thereby speeding innovation.

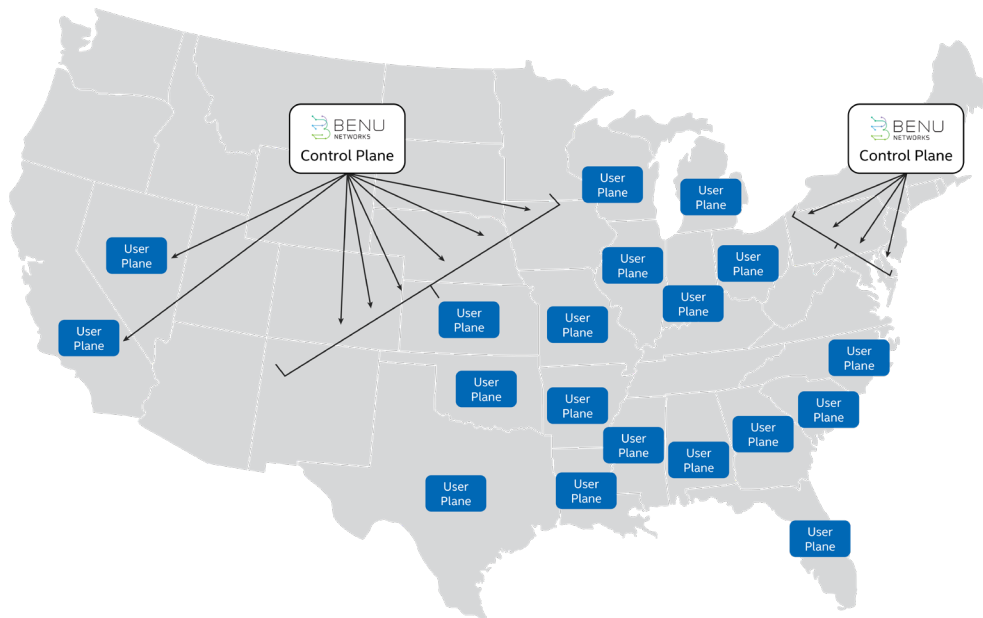


Figure 3: Control Plane and User Plane Separation: Hundreds of user planes can be managed by a pair of geo-redundant control planes, essentially creating a “virtual chassis” that can scale to millions of users and has better resiliency than the traditional BNG chassis. Furthermore, it is much easier to deploy, configure, update, and monitor than hundreds of individual BNGs.

- Provides diversity in the hardware supply chain, more competitive pricing, and better risk management in the event of a supply chain disruption.

Better economics, freedom from hardware vendor lock-in, innovation, and flexibility are some hallmarks brought by the disaggregation of hardware and software.

Control and User Plane Separation (CUPS)

Separation of the control plane from the user plane delivers adaptability, easy scale-out for growth, and true in-service upgrades.

Key benefits include:

- Allows the user plane (or multiple user planes) to be distributed geographically and deployed where they can be most effective. Cost-effective low-scale BNGs can be deployed at the MEC edge, making it easier to reduce network traffic in the aggregation and core, deliver low latency, provide user-centric services, and reduce overall cost. Virtual Customer Premise Equipment/Virtual Residential Gateway solutions also become possible.
- A centralized control plane simplifies subscriber management, management of IP address pools, and the implementation of northbound integrations to Operations Support System/Business Support System systems.
- Ability to have user plane “slices” for specific use cases like WWC and Access Traffic, Subscriber, Switching, Steering (ATSSS) support for low-latency MEC services, IoT, cloud-gaming, or streaming video.
- Fifty percent (50%+) of broadband traffic in many carrier networks is from streaming video and with distributed user planes, CDN services can be moved out closer to the subscriber to significantly reduce overall network traffic and cost.

- The CUPS architecture allows operators to dynamically apply CPU resources where they are needed resulting in better efficiency and cost-effectiveness.
- The separation means that the control plane and user plane can be scaled independently.
- Independent control and user planes provide better resiliency such that failovers go unnoticed by subscribers.

In summary, a BNG with CUPS will be able to provide a better user experience, easier management, and lower operational costs (see Figure 3).

Cloud-Native Software-Based Solution

While Benu Networks supports bare metal deployments as a BNG network appliance, many carriers are taking advantage of virtual machines and cloud-native software-based solutions. With cloud-native, there are numerous benefits:

- Cloud-native BNGs offer lower operational expenditures, ease of management and seamless scaling, giving service providers the flexibility and agility, they need. Enhanced resiliency and in-service upgrades make use of micro-services and multiple containers for a service. The cloud native deployment engines (e.g. Kubernetes) are now well understood and adopted, removing one of the big initial barriers to NFV adoption by providing management and orchestration.
- A common operational environment in addition to cloud-native 5G core and cloud-native edge services.
- Capacity can be scaled up and down as needed to allow a “pay as you grow” model for improved economics.
- With a cloud-native solution, service providers can share compute with other applications, establishing a footprint for MEC.

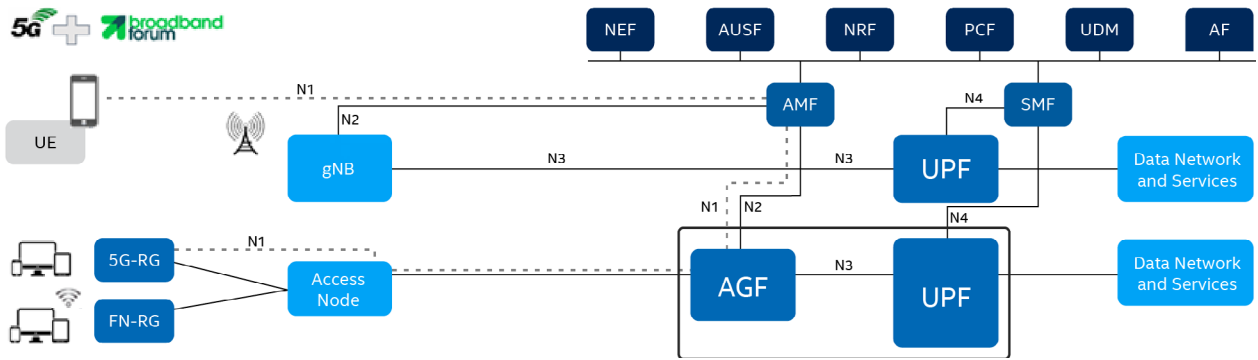


Figure 4: 3GPP and the Broadband Forum (BBF) have collaborated to define the evolution of the BNG for Wireless-Wireline Convergence (WWC).

- It provides a clear future-proofing benefit and can be easily evolved to provide FMC through a 5G Access Gateway Function (AGF) and User Plane Function (UPF). It also provides the same access network functions, hierarchical QoS, and provider edge core routing functions, but with 5G control plane interfaces (N1 and N2) to the 5G Access and Mobility Function (AMF) and the (via N4) Session Management Function (SMF). See figure 4.

In summary, a software-based BNG is economical today and can support any user service in the future, in contrast to hardware-based BNGs that are limited by the hardware ASICs and hard-wired resources inside them. Services can include enterprise-level security for remote corporate workers and other types of IP services and offer a clear architectural upgrade path to WWC implementations.

Since the services can be run in the access edge or at an aggregation point, operators can deploy a service-delivery architecture optimized for low latency, high bandwidth, and cost efficiency to provide a consistent and better user experience across 5G, fixed line, and Wi-Fi access. New services can be added quickly, and device-level policies can be set, helping to rapidly activate new services and reduce churn.

Discrete Network Design

Discrete network design has been done for years by hyperscalers. They separate the switching and connectivity function from the “brains” of the data center. We take this same approach, separating the basic port aggregation switching from the BNG, which is essentially the “brains” of the broadband network with unique subscriber context and visibility (see Figure 5). As such, we gain the following benefits:

- Significant capex and opex savings:
 - Ability to use low-cost, commoditized switching platforms for high port density.
 - Flexibility to interchange vendors in the aggregation layer where the number of nodes (due to number of ports can be high).
- Independent scaling of aggregation vs BNG platforms. This is valuable given the uncertain mix of fixed and mobile traffic over time, particularly with the rise of 5G and base station connectivity.
- Independent deployment and lifecycle of aggregation switching and BNG to better meet changing market requirements.
- Dramatically increased feature velocity on the BNG, where we have unique subscriber context and traffic visibility.

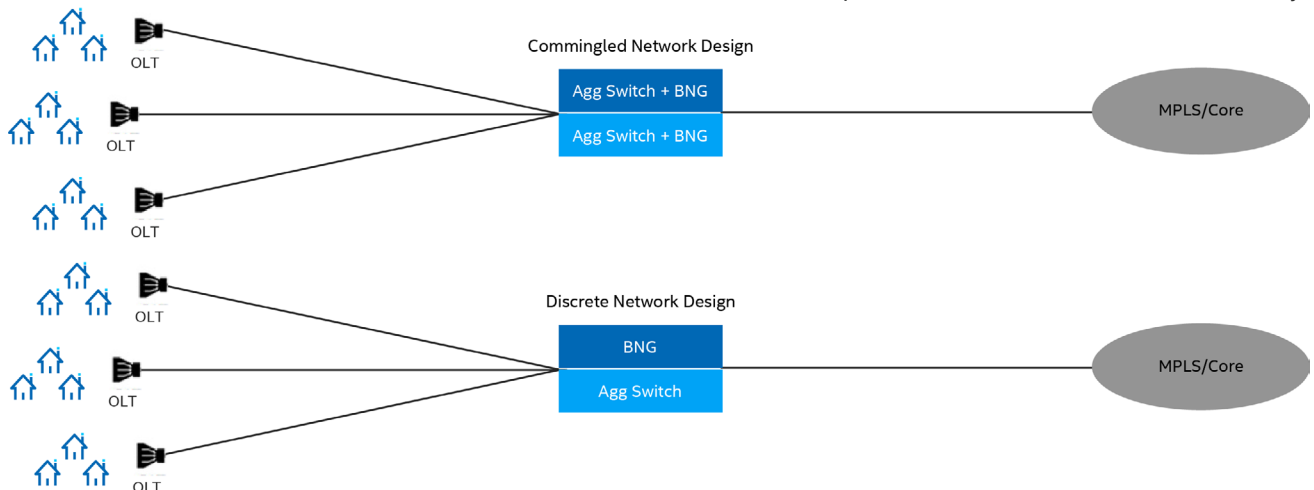


Figure 5: Cloud-Native BNG provides “discrete network design,” which keeps different functional elements separate (in this case, Aggregation Switches and BNG), so they can be scaled independently, upgraded separately with different lifecycles, provide for better software/hardware disaggregation, allow more choice in vendors, and are cost-optimized.

- No ASIC or hardware limitations on the BNG, and an architecture pre-built for MEC.
- A standard network design across any access network - DSL, PON, DOCSIS, and mobile.

In summary, when aggregation and BNG functions are combined, carriers ultimately get pinned into a single vendor or single chipset. Moving BNG complexity to an x86-based COTS platform allows deployment of simplified access and aggregation switches which can be used for any access scenario thus providing flexibility for the BNG to evolve at a different pace.

While networks will ultimately leverage cloud technologies, Benu understands the need for carriers to have transition technology. Benu supports a hardware abstraction layer that places the BNG user plane onto aggregation switches, as well as OLTs that include a switch. This provides customers with a familiar switch-based BNG and can save on space and power for very small edge deployments with only a few thousand subscribers. As edge locations scale, the benefits of discrete network design can be realized to drive better economics and flexibility (see Figure 5). Benu can deliver either solution to carriers and enable them to evolve.

Switch-based vs x86-based BNG

Benu Networks SD-Edge platform software supports a hardware abstraction layer so that it can support merchant silicon or x86-based platforms. Choosing the right platform requires examining the network requirements. Switch-based BNGs rely on ASICs which are optimized for a high rate of packet processing and forwarding. Therefore, on the coarse measurement of "cost per bit," they win because they are optimized for this function. However, there is far more functionality needed to deliver high-value subscriber services on a BNG and when considering the more important metric of "cost per subscriber," the software-based BNG running on General Purpose Processors (Intel) can be lower cost.

BNGs require per subscriber queues, schedulers, shapers, meters, counters and tunnel interfaces. As a result, even the most advanced switches cannot support large numbers of

BNG subscribers which results in "racking and stacking" a lot of switches to simply handle the subscriber scale due to the diverse resources required within the switch silicon to support a BNG subscriber. This drives higher costs and increases the number of physical interfaces, and the total power and cooling consumption. Thus, for high volume subscriber BNGs, an x86-based solution often has a lower TCO (see Table 2).

On the other hand, in situations where there is limited space and subscriber service is simpler, the switch-based BNG can be more cost-effective and provide lower TCO.

Performance Testing and Results Using an Intel-based Server

Even with the many value points of deploying a cloud-native BNG architecture, service providers have been reluctant to adopt the architecture, often questioning its ability to perform at the level required when deploying high-bandwidth services to thousands of fixed-line subscribers.

Benu Networks tackled this issue head-on by working with Intel to define and execute performance testing on an Intel-based cloud-native Benu Networks' BNG platform. The remainder of this white paper defines the performance testing environment (see Figure 7), results, and conclusion.

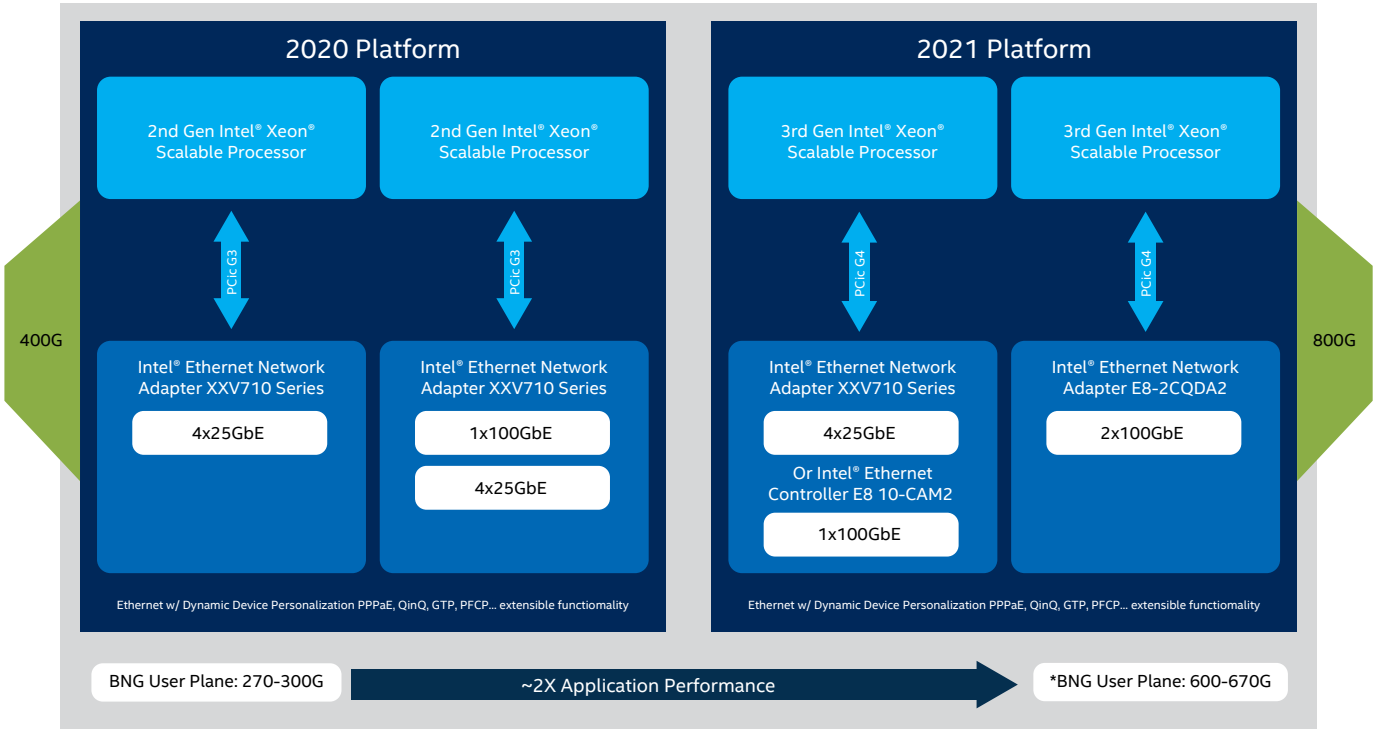
The latest generation of Intel platform hardware is the 3rd Gen Intel® Xeon® Scalable processor. The CPU moved from PCI gen 3 to a PCI gen 4 architecture (see Figure 6). This means that each CPU in a dual socket system can now attach 400G of I/O giving a theoretical server performance of 800G. Intel also introduced the E810 network interface card (NIC) which provides 200G of NIC throughput. By attaching these NICs to each CPU, the platform target throughput of 800G is realized.

On the previous 2nd Gen Intel® Xeon® Scalable processor, Intel and partners achieved 270-300G of BNG performance on a server. On this latest generation CPU we expect BNG performance to reach between 600-670G assuming a 30-32 core CPU server complex.

Solution	Cost Drivers
x86-based BNG	Bandwidth
Switch-based BNG	Per Subscriber Elements: - Queues - Schedulers - Shapers - Meters - Counters - Interfaces

Table 2: The cost of BNG and Provider Edge Routing is driven by bandwidth if using an x86-based BNG. Alternatively, the cost is driven by the number of subscribers if using a switch-based BNG. Some use cases will be better-served by an x86 BNG, and others better-served by a switch-based BNG. Benu Networks SD-Edge platform can support either type of hardware.

FMC HW Platform Evolution (BNG and AGF)



*Estimated BNG traffic improvements - projected not measured
 Note: roadmap subject to change without notice

Figure 6: 3rd Gen Intel® Xeon® Scalable processor platform for BNG

“We have focused heavily in the last 18 months to increase the packet-per-second forwarding throughput in our 3rd Gen Intel® Xeon® Scalable processor and DPDK solution.”

-Paul Mannion, Wireline Access Segment Manager, Intel

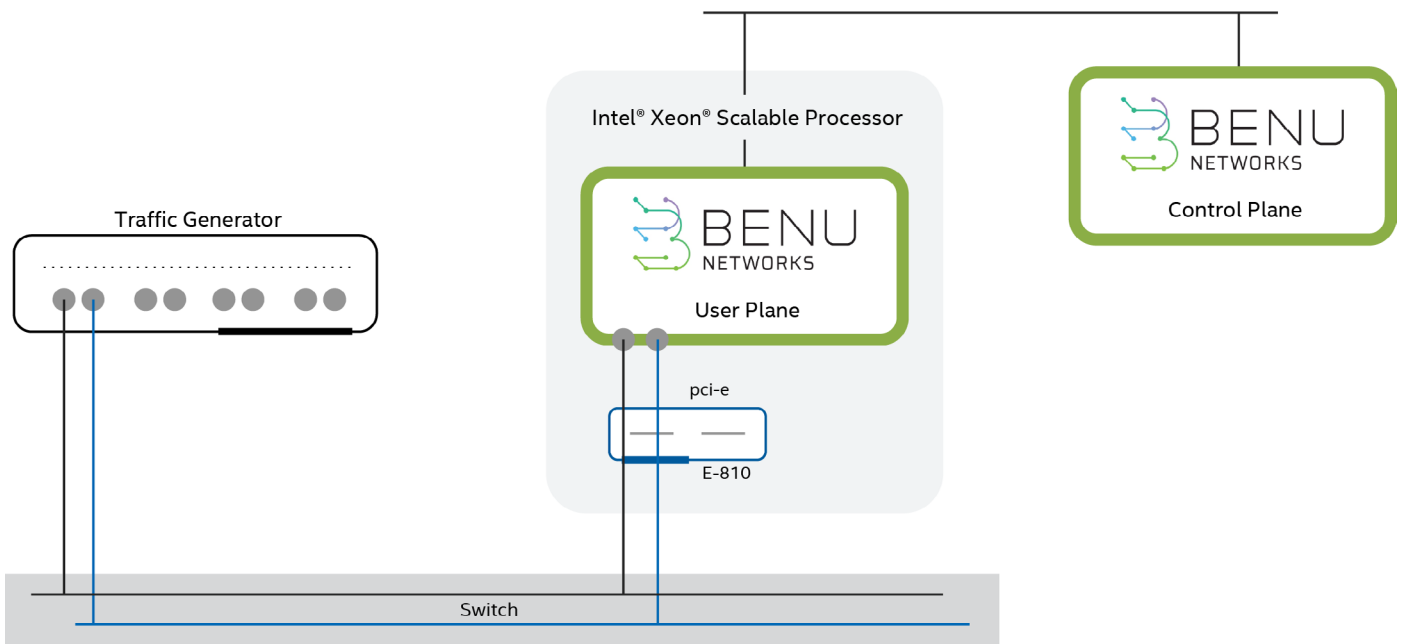


Figure 7: Performance Testing Configuration

CPU'S	Memory	Storage	Network Interface	BIOS	Other
2x 6338N CPUs	16x 16GBx3200 DIMMs for total of 256GB	1x SSD Drive 960GB	4x E810-2CQDA2	BIOS build date 02/08/21 BIOS SE5C6200.86B.0020.P09	Controller: BMC 2p74 CPLD v3p1 8x QSFP28 Active Optic cables

Table 3: Specifications of BNG User Plane system tested.

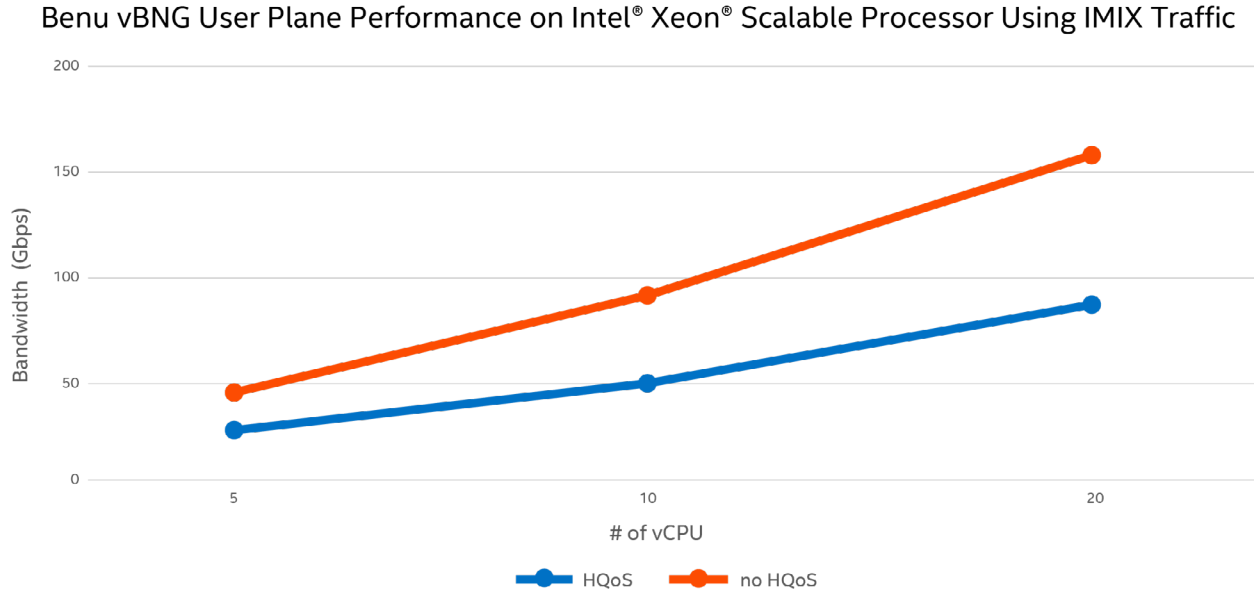


Figure 8: Cloud-native BNG performance is roughly 5 Gbps per vCPU when performing per-subscriber HQoS, and about 10 Gbps per vCPU when not using HQoS. Thus, a large 2-socket server with 50 cores could support over 500 Gbps of throughput.

The server used in the testing is shown in Table 3 above.

Benu's BNG system testing shows this 10G/Core performance is achievable in this system. With servers supporting 50 cores or more, this indicates that a user plane could achieve 500G+ in performance (see Figure 8).

A BNG Delivering Over 100 Tbps of Throughput

Traditional BNGs have the construct of a physical chassis with slots for line cards and control/route processor cards (control plane hardware responsible for control plane operations, chassis management, and routing). Sometimes vendors provide a fixed configuration (not a chassis) that integrates the control/route processor module and line card in a single 2U system.

These traditional BNGs have several disadvantages:

- Scaling is limited to the number of slots in the chassis. Once used up, a second entire BNG system must be added, and two systems must be added, if providing redundancy.
- Scaling is monolithic, requiring the control plane and user planes to be scaled at the same incremental large cost step upgrades when adding more BNG systems.

- Redundancy is expensive with typically 1:1 redundancy, requiring a costly 100% duplication of hardware.
- System failovers impact both the control plane and the user plane, slowing recovery time.
- Each BNG system is managed separately, creating enormous management overhead.
- Traditional BNGs often rely on ASICs that may not support future requirements, such as combined support for BNG, 5G AGF, 5G UPF, and SSE / SASE security services all at once.

With a cloud-native CUPS-based vBNG, these limitations are eliminated. The analogy in the CUPS architecture is that the user plane is the line card, and the control plane is the route processor, but this architecture is not constrained by the physical limitations of a chassis with a fixed number of slots. There is no physical chassis. The CUPS vBNG creates a virtual chassis that is unconstrained in geographical location and scalability with superior resiliency and greatly simplified management. As shown in Figure 3, hundreds of user planes can be associated with a pair of geo-redundant control planes. Each control plane can support up to 256 distinct user planes that have M:N resiliency.

White Paper: The 100-Terabit Broadband Network Gateway (vBNG)

This is like having a chassis with 256 line cards that are distributed geographically. As noted in Figure 8, a large 2-socket server with 50 cores can support over 500 Gbps of throughput. The performance testing shown here is a single user plane. But when 256 user planes are part of a vBNG the entire vBNG system supports over 100 Terabits of throughput.

The 100-Terabit vBNG has these additional advantages:

- Scales to millions of users, unrestricted by physical location
- Scaling can be done incrementally, not in big chassis-based steps
- Scale control plane and user plane independently at different rates and with automation
- M:N user plane resiliency is very cost effective
- User planes can failover to a user plane that is in a completely different POP or data center location
- Failover of a control plane does not impact the user plane forwarding, or vice versa
- Simplify management by reducing a network to only a small number of very large vBNG systems that have superior resiliency

The flexibility, scalability and agility of this architecture, coupled with the proven performance of the Intel Xeon processor family are strong arguments to deploy a disaggregated, cloud-native, CUPS-based BNG architecture.

Summary and Conclusion

With the performance delivered by the latest Intel Xeon Scalable processor and Benu Networks' virtualized BNG, service providers can be confident when deploying a cloud-native, disaggregated virtual BNG. The following benefits can be realized:

- By separating the control plane from the user plane (CUPS) and running on commercial hardware a service provider can improve service availability and simplify network management.
- Network costs decrease.
- Flexibility increases. User planes can be deployed where they can be most geographically and functionally efficient.

- No longer constrained to vendor-specific hardware the carrier can scale up or scale down as needed.
- Innovate and add new services quickly.
- Employ an architecture built for wireless-wireline convergence and create a footprint for MEC.

To give confidence to service providers that the latest generation of commercial servers (x86) have the power and capability to support the speed, service mix, and traffic volume of the ever-growing needs of broadband subscribers, Benu Networks and Intel have jointly partnered on performance testing and tuning of the architecture running on Intel servers. Benu found the Intel platform had more than enough power and capability to handle the load and service mix.

There is a transformation happening at the edge of the network, and the time to invest and win the future is now. Cloud-native infrastructure will address many of the challenges service providers face in today's hyper-competitive environment.



Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. 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 Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.