

Cloudifying the Broadband Network Gateway

From dedicated appliances to an agile,
horizontally scalable Telco Cloud

Technical architecture proposal supported by
production-relevant vBNG validation

White Paper



intel

Radisys

DELL



Cloud Networking for AI

www.6wind.com



Table of Content

Executive Summary 3

1. Broadband Networks Are Changing 5

2. The Limits of a Vertically Integrated BNG Model 7

3. Cloudification Changes the Infrastructure Model 8

4. Four Sources of Value from BNG Cloudification 10

5. Efficiency Is a System Property 12

6. Making the Architecture Real: COTS + 6WIND vBNG 14

7. Can the Service Layer Be Preserved?. 16

8. Functional Validation of the Broadband Service Layer 17

9. Validation Under Production-Relevant Conditions 19

10. Evidence: Preserving the Tested Service Objectives 21

11. What the Evidence Establishes 23

12. What Cloudification Enables at Infrastructure Scale 25

13. Decoupling Service, Software and Hardware Lifecycles 26

14. From Infrastructure Efficiency to TCO 28

15. Progressive Transformation and Asset Protection. 29

Conclusion: Cloudify the Function, Preserve the Service 31

References and Further Reading 32

Executive Summary

Broadband networks are evolving toward fiber-based access while telecommunications infrastructure is simultaneously moving toward increasingly software-defined and cloud-oriented architectures. The Broadband Network Gateway (BNG) is a critical step in this transformation because it combines high-throughput forwarding with subscriber management, authentication, policy enforcement and Quality of Service.

The central question

“Can service providers cloudify the BNG and gain the agility and infrastructure efficiency of a software-based architecture without compromising the broadband service characteristics expected from the service edge?”

A software-based BNG deployed on standard COTS infrastructure changes the infrastructure model. Rather than coupling service functionality, forwarding capacity and hardware lifecycle inside a dedicated appliance, the BNG becomes a software workload consuming resource from a horizontal compute and networking layer. This also changes the programmability model. Packet-processing behavior can evolve in software rather than being constrained by the fixed or partially programmable processing stages exposed by a particular forwarding ASIC. This provides greater flexibility to introduce, modify and combine broadband functions as service requirements evolve.

Cloudification value	What changes	Potential operator value
Service continuity	BNG moves to software while preserving tested service objectives	Production viability
Lifecycle independence	Service software becomes less dependent on proprietary HW capabilities	Faster service evolution and shorter time to service
Programmability	Packet-processing behavior evolves in software rather than being constrained by fixed or partially programmable ASIC processing stages	Greater flexibility to introduce and adapt broadband functions
Elastic scalability	Capacity can scale horizontally rather than inside one fixed appliance	Capacity follows subscriber and traffic demand
Infrastructure mutualization	COTS resources can support multiple instances and network functions	Higher infrastructure utilization and service composition

A laboratory validation was used as technical evidence for this architecture, not as the objective of the paper. It tested whether a vBNG on x86 COTS infrastructure could preserve production-relevant throughput, subscriber scale, QoS and latency. The validation reached approximately the 100 GbE target connectivity envelope, scaled to 60,000 PPPoE sessions, and exercised two- and three-queue QoS. Average latency remained below the defined 80µs objective across all reported scenarios, while rare higher-latency excursions were observed and are addressed in the detailed results.

Importantly, reaching the 100 GbE test envelope required only a limited subset of the compute resources available on the server platform. The server provided 96 physical CPU cores in total, while the 60,000-subscriber configuration used 10 physical cores for the vBNG fast path. The comparison is not intended to imply linear scaling from 10 to 96 cores, but it demonstrates that the 100G result was achieved without exhausting the server's compute capacity. The approximately 100G result should therefore be understood as the connectivity envelope exercised by this validation, not as an identified processing limit of the vBNG or the underlying COTS platform.

This available compute headroom is particularly relevant to the cloud model. Additional resources can potentially be assigned to increase vBNG processing capacity, while capacity can also be scaled horizontally through additional vBNG instances and COTS nodes. Subject to production dimensioning, NIC capacity, workload isolation and service requirements, the same infrastructure model can also enable co-location of other network functions, allowing otherwise available compute capacity to be used for additional services rather than reserving an entire server for the BNG alone.

The broader efficiency opportunity is therefore a system property. It combines optimized and programmable packet-processing software, power-aware COTS hardware, efficient use of compute resources, horizontal scalability and infrastructure mutualization across multiple virtual network functions. These architectural benefits are not all quantified by the present benchmark; the validation establishes the technical feasibility and resource-efficiency evidence on which broader production dimensioning and TCO assessment can be built.

1. Broadband Networks Are Changing

FTTH is progressively becoming the foundation of fixed broadband. The transition occurs over time: fiber networks coexist with xDSL and cable infrastructure while subscribers migrate progressively. The ecosystem is also becoming more heterogeneous through alternative fiber operators, wholesale networks, Open Access and Layer 2 Bitstream Access models.

Subscriber distribution, traffic concentration and capacity requirements can therefore evolve by geography and over time. The service edge must be able to follow this evolution without forcing capacity to be committed too early to a fixed physical platform.

Broadband in a gigabit society

As access capacity increases, the technical quality of residential broadband increasingly converges around three fundamentals: available bandwidth, low and predictable latency, and reliability. Commercial differentiation increasingly occurs above the connectivity layer through content, security, cloud services, OTT offers and emerging digital services.

Service-layer implication

The connectivity layer must become more agile and economically efficient while continuing to behave as dependable infrastructure. Cloudification only creates value if those service characteristics are preserved.

Planning for an uncertain bandwidth future

Gigabit-class access already provides significant headroom for many current residential applications. The timing and magnitude of the next major traffic growth cycle are less predictable. AI and other emerging applications may create new traffic patterns, but operators cannot efficiently dimension every service-edge platform around an uncertain long-term peak.

Why the BNG infrastructure model must evolve

Network change creates a need for service agility without sacrificing service quality.




Figure 1. Drivers for a more adaptable broadband service edge.

Broaden the ecosystem and improving choice for broadband service platforms

The velocity in service features require functional and performance options go hand in hand. Service diversity on top of quick updates on AI, security, etc. is hampered by integrating anything on a single platform where a lot of these innovations are driven by non-broadband ecosystem players more commonly working with IT or data center technologies rather than telecommunication primitives.

The architectural question becomes

how can an operator add or reallocate capacity when requirements change without prematurely replacing or over-dimensioning an entire service-edge platform?

2. The Limits of a Vertically Integrated BNG Model

Traditional BNG architectures closely associate the broadband network function with a dedicated hardware platform. This vertically integrated model has delivered predictable performance and operational maturity, but it also creates a tight relationship between service capability, platform capacity and hardware lifecycle.




Figure 2: Tight coupling between service capability, platform capacity and hardware lifecycle in a traditional BNG architecture.

As broadband traffic, service requirements and infrastructure strategies evolve, that coupling can constrain how independently each of these dimensions can change.

Capacity is tied to the appliance

The service is dimensioned inside the capacity boundaries of a specific platform. Capacity tends to be purchased in platform-specific increments, frequently with headroom for future growth. Resources remain dedicated to the BNG even when the service does not fully use them.

Service evolution can depend on the hardware roadmap

When functionality relies on capabilities unique to a proprietary forwarding platform, service evolution can depend on development roadmaps, chipset capabilities, platform generations, line-card evolution and vendor qualification cycles. In addition, multi-supplier strategies are limited to the least common discriminator. This can couple innovation and time to service to the hardware roadmap rather than to the service roadmap.

Dedicated capacity limits mutualization

Compute or forwarding resources installed for a dedicated BNG are generally not available to another network function when BNG utilization is low. Across multiple function-specific silos, this can create stranded capacity and force each appliance to be sized independently.

3. Cloudification Changes the Infrastructure Model

Cloudification is not simply the same BNG implemented on a different processor. It changes the relationship between service workloads and infrastructure.

Traditional model: vertical network silos

Each network function runs on dedicated hardware, with its own capacity, lifecycle and management.




Figure 3: Traditional network architecture based on dedicated, function-specific hardware silos.

Cloud model: horizontal infrastructure

Multiple network functions run as software workloads on a shared, scalable infrastructure.




Figure 4: From dedicated network-function silos to a shared horizontal infrastructure model.

Network functions become software workloads consuming hardware and operational resources from a horizontal infrastructure layer. Service capacity and physical infrastructure can therefore be dimensioned more independently. The architecture creates the option to scale a function across instances and nodes, and to use the same infrastructure pool for different network functions subject to isolation, resilience and performance requirements.

From appliance sizing to workload sizing

*Subscriber demand → required vBNG capacity → compute and networking resources.
Capacity planning can progressively shift from selecting a fixed appliance envelope toward sizing and scaling software workloads.*

4. Four Sources of Value from BNG Cloudification

Cloudifying the BNG creates value beyond replacing a dedicated appliance with a software workload. It changes how broadband services can be operated, evolved and scaled by separating service capacity and software evolution from a fixed, function-specific hardware platform.

As illustrated below, this transformation can be viewed through four complementary dimensions: preserving service continuity, increasing lifecycle independence, enabling elastic scalability, and improving infrastructure mutualization. Together, these dimensions define the broader value proposition of a cloudified BNG and provide the framework for assessing whether the architecture can deliver greater flexibility without compromising the service characteristics expected from the broadband edge.




Figure 5. Four complementary sources of value from BNG cloudification.

4.1 Preserve the service layer

The first requirement is not an efficiency gain. It is the absence of an unacceptable service trade-off. A cloudified BNG must continue to deliver the expected throughput, subscriber scale, latency, QoS and broadband functionality.

4.2 Decouple the service lifecycle from the hardware lifecycle

Running broadband functions as software on generic compute and networking resources reduces dependency on functionality unique to a specific hardware platform. The goal is not to eliminate vendors, but to standardize the infrastructure boundary so that software and hardware

can evolve on more independent timelines.

This can reduce the dependency of service innovation on a proprietary hardware roadmap, with potential benefits for time to service and ultimately time to revenue.

Lifecycle independence also changes how operators can approach migration. In a vertically integrated environment, introducing a new BNG platform typically requires parallel hardware deployment, followed by the potentially complex migration of subscriber populations, service configurations and operational dependencies. A software-defined architecture provides greater flexibility to introduce new vBNG workloads progressively while existing BNG capacity continues serving selected subscribers. Residential and business services, for example, can be moved according to their respective requirements and operational readiness, rather than being forced into a common hardware replacement schedule. Combined with complementary subscriber-steering mechanisms, this creates opportunities to reduce migration complexity, preserve existing assets during transition and align service evolution more closely with business priorities.

4.3 Scale the service independently from a dedicated platform

A physical server does not need to correspond to one BNG. Multiple vBNG instances can provide horizontal subscriber scalability, allowing capacity to grow through additional software instances and COTS resources rather than primarily through larger appliance variants.

4.4 Mutualize the infrastructure

Subject to operational constraints, the same horizontal infrastructure can host different network functions. A server or server pool could, for example, host multiple vBNG instances or combine vBNG, CGNAT and IPsec VPN workloads to build higher-value service chains. These co-location examples are architectural options and were not validated by the benchmark described later in this paper.

The common infrastructure model also changes the investment required to introduce or evolve BNG services. New service environments do not necessarily require additional dedicated physical platforms: multiple vBNG instances, such as separate residential and business workloads, can share compatible COTS resources while retaining distinct service configurations, capacity allocations and lifecycle requirements. This enables operators to make greater use of available infrastructure and introduce capacity incrementally rather than committing to a new function-specific appliance for each service evolution. Together with progressive subscriber migration, infrastructure mutualization creates opportunities to protect existing investments, limit stranded capacity and align infrastructure expenditure more closely with actual service demand.

5. Efficiency Is a System Property

The efficiency of a cloudified BNG should not be reduced to a single throughput-per-core or watts-per-Gbps metric. It results from the interaction of software efficiency, hardware efficiency and the way cloud infrastructure resources are allocated and utilized over time.

This last dimension is particularly important. A dedicated BNG appliance is dimensioned for its function and typically remains powered and reserved even when traffic demand is below peak capacity. A software-based BNG running on COTS infrastructure creates the possibility to align active compute resources more closely with changing service demand. Resources that are not required by the BNG can potentially be power-managed, released for other workloads, or used to support additional network functions, while capacity can be increased again as demand grows.

Efficiency layer	Mechanism	Why it matters
Optimized software	Efficient packet processing and subscriber/QoS implementation	Converts bounded CPU resources into high service capacity
Power-aware COTS hardware	High-density compute, high-speed NICs and platform power-management capabilities	Provides mechanisms to manage compute and power according to workload
Elastic resource utilization	Adjust active resources and scale processing capacity as demand changes	Creates the potential to reduce resources dedicated to off-peak demand and restore capacity when required
Infrastructure mutualization	Resource pooling, workload co-location and horizontal scale-out	Reduces function-specific stranded capacity and improves utilization across aggregate demand

Efficiency also has a time dimension. Broadband traffic demand varies throughout the day, while a dedicated hardware platform remains allocated to the BNG irrespective of how much of its peak capacity is being used. In a cloud architecture, the objective is to make infrastructure consumption more elastic: allocate processing resources when the service requires them, reduce or power-manage resources when demand falls, and scale out when demand increases. This creates the potential to optimize both infrastructure utilization and energy consumption across the daily traffic cycle.

The extent to which resources can be dynamically reduced, placed into lower-power states or reassigned depends on the server platform, orchestration model, workload distribution and service-availability requirements. These mechanisms were not evaluated in the present benchmark and are therefore presented as an architectural opportunity rather than a measured energy-saving result.




Figure 6: From efficient packet processing to demand-aligned infrastructure utilization.

The tested server was configured with BIOS and Linux parameters optimized for power consumption. The supplied results do not contain numerical power measurements, so this paper does not quantify energy savings. Power measurement remains an important input to the wider production-dimensioning and TCO exercise.

Efficiency principle

Efficiency is not only about doing more with each CPU core. It is also about improving how much of the infrastructure is productively used over time and allowing capacity to follow service demand.

6. Making the Architecture Real: COTS + 6WIND vBNG

The architectural proposition was evaluated using a concrete implementation of a virtual BNG on standard server infrastructure.

Layer	Validated implementation
COTS server	Dell PowerEdge R660 [6]
Networking	2 × 100 GbE Intel E810 interfaces [3]
Virtualization / I/O	vBNG started as a VM with SR-IOV access to the NICs
Packet processing	6WIND vBNG using an accelerated data plane [1]
Available CPU allocation	36 physical cores (72 threads) for data plane; 4 physical cores (8 threads) for control plane
Traffic generation	Ixia with 2 × 100 GbE interfaces




Figure 7: Validation topology between subscriber-side and network-side traffic generation.

Ecosystem contribution

The validation brings together complementary software, compute, networking and infrastructure capabilities to demonstrate the feasibility of a cloudified broadband service edge.

6WIND contributes the software-defined vBNG [1] and high-performance packet-processing architecture. Its accelerated data plane enables broadband forwarding, subscriber processing and QoS functions to operate as a software workload on compatible COTS infrastructure.

Intel contributes the general-purpose compute and Ethernet technologies used in the validated implementation. Intel Xeon 6 SP processors [5] provide multicore processing resources and platform power-management capabilities, allowing compute capacity to be allocated according to workload requirements.

The Intel Ethernet Network Adapter E810 [3] complement the software data plane through hardware-assisted packet classification and distribution. In this implementation, SR-IOV provides the vBNG with direct access to network resources, while the E810's Device Config Function (DCF) enables hardware flow rules to be configured through a trusted virtual function. The SR-IOV eSwitch supports traffic steering between virtual functions, and Intel Dynamic Device Personalization (DDP) [4] enables more efficient PPPoE-aware packet distribution across processing resources.

By offloading packet classification and distribution functions to the network adapter, this approach reduces the need for dedicated software-based packet dispatching and helps preserve CPU resources for vBNG processing. It illustrates how general-purpose compute and programmable Ethernet capabilities can be combined to optimize a software BNG implementation without tying the broader architecture to a function-specific hardware appliance.

Dell Technologies provides the PowerEdge R660 [6] COTS server platform used to host the validated vBNG workload, illustrating how standardized server infrastructure can support production-relevant broadband processing requirements. While the PowerEdge R660 is a standard server for usual environments, the same functionality can be achieved on servers like the PowerEdge XR8000 Series designed for harsh edge environments, hence enabling vBNG deployments in practically every location.

Radisys contributes to the broader cloud-native broadband and migration ecosystem through its SDN controller and 3S management and control capabilities. These provide complementary mechanisms for subscriber placement, service differentiation and progressive migration between existing and virtualized broadband environments.

Deutsche Telekom AG contributes the service-provider perspective through production-relevant broadband requirements, operational constraints and use cases that inform the architectural evaluation. These requirements help frame the performance, subscriber scale, QoS and latency objectives considered in the paper. The validation results are laboratory measurements of the specified 6WIND vBNG implementation and are not presented as Deutsche Telekom AG test results or as a certification or endorsement of the solution.

The selected partner technologies demonstrate one optimized implementation of the architectural model. They do not redefine the vBNG as a hardware-specific solution: the underlying proposition remains the separation of broadband service software from dedicated, vertically integrated infrastructure.

7. Can the Service Layer Be Preserved?

Once the architectural benefits are established, one question remains: can those benefits be obtained without compromising the characteristics expected from a production BNG?

Answering that question requires more than demonstrating forwarding capacity. A production BNG must first provide the subscriber functions on which broadband services depend, including session establishment, IPv4 and IPv6 operation, authentication, subscriber and access identification, and the handling of access-network information. These functions must then coexist with the throughput, subscriber scale, latency and QoS characteristics required from the broadband service edge.

The validation program therefore addressed both dimensions. Functional testing established the broadband service baseline, while the subsequent performance testing evaluated how the vBNG behaved as subscriber density, traffic load and QoS processing increased.

Service requirement	Validation question
Broadband functionality	Can the software platform establish and operate the subscriber services expected from a BNG, including IPv4 and IPv6 operation, authentication and access information handling?
Performance	Can COTS deliver production-relevant forwarding capacity?
Subscriber scale	Can subscriber density increase without losing target throughput?
Latency	Can packet latency remain within the defined service requirement?
QoS	Can the system enforce service queues and congestion behavior while sustaining the workload?
Resource footprint	Can the workload be delivered without consuming the full server CPU capacity?

The validation evidence is presented accordingly in two stages. Section 8 establishes the functional baseline by examining PPPoE operation across IPv4 and IPv6, concurrent IPv4/IPv6 services, subscriber authentication and access information handling. Section 9 then evaluates production-relevant performance through subscriber scale, throughput, hierarchical QoS, oversubscription behavior and latency. Together, these two stages address whether the BNG can be cloudified while preserving the tested service layer.

Evidence, not objective

The validation provides the technical evidence supporting a prerequisite of the cloudification model: that the BNG can be implemented as a software workload on COTS infrastructure while preserving the broadband functionality and service characteristics exercised by the test program. The results are used to establish this production-relevant service baseline and to support the broader architectural analysis developed in the following sections.

8. Functional Validation of the Broadband Service Layer

Preserving the broadband service layer requires more than forwarding traffic at scale. The vBNG must provide the subscriber, addressing, authentication and access functions expected from the broadband edge. This section summarizes the functional validation performed to establish that foundation before evaluating performance under production-relevant load.

8.1 Establishing the functional baseline

The functional validation exercised representative subscriber services across different IP, VLAN and subscriber configurations. It covered PPPoE operation with IPv4 and IPv6, concurrent IPv4 and IPv6 services, subscriber authentication, NAS functionality, subscriber and access-circuit identification, and the processing of broadband access-loop attributes. All functional scenarios included in the validation passed.

The tests were exercised across several subscriber and C-VLAN configurations rather than a single access configuration. These included a single subscriber on a single C-VLAN, multiple subscribers sharing a C-VLAN, and multiple subscribers distributed across multiple C-VLANs. Traffic was generated in both directions and, where applicable, with different packet-size profiles.

These results establish the functional baseline for interpreting the subsequent performance validation. The scalability and QoS measurements should therefore be understood as the next stage of a broader vBNG validation, rather than as isolated packet-forwarding benchmarks.

8.2 PPPoE operation across IPv4 and IPv6

PPPoE subscriber establishment and forwarding were validated independently for IPv4 and IPv6. For PPPoEv4, the test program exercised subscriber sessions across the defined C-VLAN and subscriber combinations, with bidirectional traffic and multiple packet profiles.

The PPPoEv6 validation extended beyond basic session establishment. It included IPv6 network-address assignment to the subscriber and IPv6 prefix assignment for hosts behind the subscriber. Traffic was generated between hosts within the assigned prefix and the core-side network, exercising IPv6 forwarding through the vBNG. The tests were repeated across the different C-VLAN/subscriber configurations and packet-size profiles.

The IPv6 validation therefore covered both PPPoE session establishment and the subscriber-side addressing and forwarding functions required to carry IPv6 traffic through the vBNG.

8.3 Concurrent IPv4 and IPv6 operation

The validation also exercised PPPoEv4 and PPPoEv6 concurrently on the same vBNG, extending the independent protocol tests to a configuration in which both environments operated in parallel.

During the combined tests, IPv4 traffic was exchanged between subscriber addresses and

the core network while IPv6 traffic was simultaneously exchanged between prefix-assigned subscriber addresses and IPv6 core destinations. The tests covered single- and multi-subscriber configurations and multiple packet-size profiles, with packet loss and latency collected as traffic KPIs.

The successful completion of these scenarios demonstrates that the tested vBNG configuration could operate the two protocol environments concurrently within the same broadband processing environment.

8.4 Subscriber authentication and access information

The functional program also validated functions associated with subscriber authentication and access-network context. PAP and CHAP authentication were exercised together with NAS support. Circuit-ID and Remote-ID handling was validated to verify that the vBNG could process subscriber and access-circuit identification information supplied by the access network.

The test scope further included the processing of broadband access-loop attributes describing upstream and downstream line characteristics. These included actual, minimum, attainable and maximum data-rate information, together with selected low-power and interleaving-delay parameters.

Functional area	Scope exercised	Result
PPPoEv4	Session establishment and bidirectional IPv4 forwarding across defined C-VLAN / subscriber configurations	PASS
PPPoEv6	Session establishment, IPv6 address/prefix assignment and bidirectional IPv6 forwarding across defined C-VLAN / subscriber configurations	PASS
Concurrent IPv4 + IPv6	PPPoEv4 and PPPoEv6 operated concurrently on the same vBNG with bidirectional traffic	PASS
Authentication / NAS	PAP, CHAP and NAS support	PASS
Subscriber/access identification	Circuit-ID and Remote-ID handling	PASS
Access-loop attributes	Defined data-rate and interleaving-related attributes	PASS

With the functional broadband baseline established, the next stage of the validation addressed whether these service capabilities could be supported while meeting production-relevant requirements for subscriber scale, throughput, QoS and latency.

9. Validation Under Production-Relevant Conditions

With the functional broadband baseline established, the next stage of the validation examined the vBNG under increasing subscriber density, traffic load and QoS complexity. The scenarios were designed to exercise subscriber and hierarchical QoS behavior under representative broadband traffic conditions, rather than to maximize packet-forwarding throughput by removing service functions.

Two initial scenarios combined different subscriber densities, traffic profiles and service-queue configurations. A further extension increased the subscriber population to 60,000 sessions while maintaining 100G-class traffic and exercising both two-queue and three-queue QoS configurations. Together, these scenarios provide complementary views of service throughput, subscriber scale and QoS behavior.

Scenario A: 100 Gbps target, three service queues

Scenario A represents the more traffic-intensive service profile. It combines 15,000 PPPoE subscribers with a 100 Gbps target and three downstream service queues representing Internet, IPTV and VoIP traffic.

- ▶ **15,000** PPPoE subscribers.
- ▶ **100 Gbps** target throughput.
- ▶ **3** downstream QoS queues:
 - Internet 79% using IMIX (35×78-byte, 15×740-byte, 50×1500-byte packets);
 - IPTV 20% using 1500-byte packets;
 - VoIP 1% using 78-byte packets.
- ▶ **10%** oversubscription scenario to exercise QoS under contention.

The mix deliberately combines high-volume Internet and IPTV traffic with a smaller latency-sensitive VoIP class. The oversubscription case increases the offered load beyond the configured service capacity, allowing the validation to examine whether the QoS hierarchy enforces the intended service differentiation when contention occurs.

Scenario B: 60 Gbps target, two service queues

Scenario B increases subscriber density to 30,000 sessions while using a lower aggregate traffic target and a two-queue service profile. It provides a complementary operating point for examining subscriber scale and QoS processing.

- ▶ **30,000** PPPoE subscribers.
- ▶ **60 Gbps** target throughput.
- ▶ **2** downstream QoS queues:
 - Internet 99% using IMIX (35×78-byte, 5×740-byte, 60×1500-byte packets);
 - VoIP 1% using 78-byte packets.
- ▶ **10%** oversubscription scenario.

Rather than representing a higher per-subscriber traffic profile, this scenario increases the number of active PPPoE sessions and therefore complements Scenario A by exercising a different balance between subscriber density and aggregate throughput.

Subscriber-scale extension

The validation was subsequently extended to 60,000 PPPoE sessions at 95 Gbps, using both two-queue and three-queue QoS configurations. This extension increased subscriber density while retaining 100G-class forwarding and hierarchical QoS processing.

The 60,000-subscriber configuration used a VM assigned 14 physical CPU cores, of which 10 physical cores were allocated to the vBNG fast path. The QoS scheduler was shared across the fast-path threads rather than consuming a dedicated scheduler core.

10. Evidence: Preserving the Tested Service Objectives

The purpose of the validation was not simply to establish a peak forwarding number, but to determine whether the service characteristics expected from a production broadband edge could be preserved when the BNG is implemented as a software workload on COTS infrastructure. The following results therefore consider throughput, subscriber scale, QoS behavior, compute utilization and latency together. Particular attention is given to latency distribution, since average latency alone does not describe the occurrence or frequency of higher-latency packets.

10.1 Reaching the target data-center connectivity envelope

Scenario A delivered 95 Gbps with 15,000 PPPoE sessions. This result should be read against the approximately 100G target connectivity envelope of the validation and the target data-center fabric model. It should not be interpreted as the maximum throughput capability of the vBNG solution.

The optimized 15k/30k tests used a VM assigned 11 physical cores, with seven physical cores allocated to the fast path. This is important to the cloudification story: production-relevant network capacity was reached without allocating the full processing capability available on the server platform.

10.2 Scaling subscriber density and QoS complexity

Configuration	Sessions	Throughput	DP CPU avg.	Loss	Avg latency
Scenario A, 3 queues	15k	95 Gbps	95%	0%	22.4 μ s
Scenario A, oversubscribed	15k	95 Gbps	96%	7%*	22.5 μ s
Scenario B, 2 queues	30k	60 Gbps	76%	0%	11.7 μ s
Scenario B, oversubscribed	30k	60 Gbps	83%	4.5%*	12.8 μ s
60k Subs, 2 queues	60k	95 Gbps	70%	0%	13 μ s
60k Subs, 2 queues, oversubscribed	60k	95 Gbps	88%	2.6%*	41 μ s
60k Subs, 3 queues	60k	95 Gbps	90%	0%	17 μ s
60k Subs, 3 queues, oversubscribed	60k	95 Gbps	95%	2.6%*	50 μ s

* Packet loss in oversubscribed cases is attributed in the test results to expected QoS behavior under deliberate oversubscription.

CPU note: utilization refers to the CPU resources allocated to the vBNG fast path, not to utilization of the complete server CPU capacity.

10.3 Latency distribution and the 80 μ s objective

The validation defined an 80 μ s latency objective. Across the reported configurations, average latency remained below this objective, including at the highest tested subscriber density. At 60,000 sessions, average latency was 13 μ s with two queues and 17 μ s with three queues. Under deliberate oversubscription, average latency increased to 41 μ s and 50 μ s respectively, while remaining below the 80 μ s objective on an average basis.

Average latency, however, does not provide a complete view of packet behavior. The measurements also recorded individual packets above the 80 μ s threshold and rare maximum-latency excursions into the millisecond range. For this reason, the latency distribution is important when interpreting the results. A latency histogram shows the proportion of packets within different latency ranges and provides the statistical context needed to distinguish the predominant packet behavior from infrequent outliers.

The observed maximum should therefore not be interpreted as representative latency for the vBNG workload. At the same time, the outliers should not be hidden by the average. Both the distribution and the maximum values are reported to provide a transparent view of latency behavior. The cause of the higher-latency excursions remains under investigation.

11. What the Evidence Establishes

The value of the measurements lies in the architectural conclusions they support rather than in any individual benchmark number. In particular, the validation provides evidence that a software BNG can deliver a production-relevant combination of throughput, subscriber processing, QoS and latency using a bounded subset of the resources available on a COTS server. The approximately 100G result represents the connectivity envelope exercised by this validation, not an identified processing ceiling for the vBNG architecture.

A production-relevant traffic profile can be supported on COTS infrastructure

Scenario A is particularly relevant when relating the laboratory results to current broadband traffic profiles. The configuration delivered 95 Gbps across 15,000 PPPoE sessions while three QoS queues were active, corresponding to approximately 6.3 Mbps of aggregate measured throughput per subscriber. This is close to the 6 to 7 Mbps per-subscriber peak-hour traffic level used as a production reference in this paper.

The significance of this result is therefore not simply that the system approached a 100G interface rate. It did so while exercising subscriber processing and a multi-queue QoS model at a subscriber-to-throughput ratio representative of the production dimensioning assumption.

Subscriber scale does not inherently require a service-layer trade-off

The 60,000-session tests address a different dimension of the architecture: subscriber density. The vBNG maintained up to 95 Gbps while processing 60,000 PPPoE sessions in both two-queue and three-queue configurations. Average latency remained below the defined 80 μ s objective, although rare higher-latency excursions were observed as described in Section 10.

The 60k result should therefore not be interpreted as a model of 60,000 subscribers simultaneously consuming peak-hour bandwidth at 6 to 7 Mbps each. Rather, it demonstrates that high subscriber density, QoS processing and 100G-class forwarding can coexist within the tested software architecture.

100G is the validation envelope, not the architectural scaling limit

The approximately 100G throughput reached during the validation reflects the physical connectivity envelope of the test configuration. It should not be interpreted as the maximum forwarding capability of the vBNG.

The resource measurements are important in this context. The 60k configuration used 10 physical cores for the vBNG fast path on a server containing 96 physical CPU cores in total. This comparison does not imply that throughput scales linearly with core count, nor that all remaining cores can simply be converted into additional forwarding capacity. NIC bandwidth, NUMA placement, memory bandwidth, traffic characteristics, QoS processing and other platform constraints also contribute to system dimensioning. It does, however, demonstrate that reaching the tested connectivity envelope did not exhaust the compute capacity of the server platform.

The software architecture provides scale-up and scale-out paths

Moving beyond a single 100G workload can be approached along two complementary dimensions. Scale-up increases the processing resources available to a vBNG workload and the corresponding network I/O capacity. Scale-out distributes capacity across additional vBNG instances and COTS nodes, allowing the broadband service to grow horizontally rather than remaining bounded by the capacity envelope of one dedicated appliance.

For higher aggregate capacities, these two mechanisms can be combined. The validation establishes a 100G-class building block while demonstrating that the tested workload does not consume the full server resource envelope. Higher capacities can therefore be addressed through a combination of scale-up, by assigning additional processing and network resources to the vBNG workload, and scale-out, by distributing capacity across additional vBNG instances and COTS nodes. The results should not be extrapolated linearly, since actual system capacity depends on compute, networking, memory, NUMA placement, traffic characteristics and service configuration. Nevertheless, the validation provides the technical foundation for progressively scaling the architecture beyond the connectivity envelope exercised in this test.

Bounded resource use creates headroom for the cloud model

The same observation has implications beyond forwarding scale. A cloud architecture does not require every available server resource to remain permanently dedicated to a single BNG workload. Where sufficient compute and I/O capacity remain available, those resources can potentially support additional vBNG instances, increased BNG processing capacity, or compatible co-located network functions.

The benchmark does not itself demonstrate production co-location, dynamic resource allocation or a multi-instance 1 Tbps system. It does, however, establish the technical basis for evaluating those deployment patterns because the validated BNG workload operates within a bounded compute footprint rather than consuming the complete server resource envelope.

Technical conclusion

A properly engineered vBNG running on standard x86 COTS infrastructure can provide a credible foundation for cloudifying the broadband service edge while preserving the service characteristics exercised by the validation.

12. What Cloudification Enables at Infrastructure Scale

The validation establishes the service-layer prerequisite. The next question is what the cloud model can enable at infrastructure scale. A dedicated hardware BNG is primarily scaled within the capacity envelope of the installed platform, while a cloudified BNG introduces horizontal scale across software instances and COTS nodes.




Figure 8: Horizontal scaling changes capacity planning from appliance sizing to workload sizing.

Vertical appliance model	Horizontal cloud model
Capacity grows inside a platform envelope	Capacity can grow by adding vBNG instances and/or COTS nodes
Resources remain dedicated to the BNG	Resources can potentially be pooled across compatible workloads
One function is associated with one appliance domain	Multiple instances and functions can share a common infrastructure layer
Upgrade path follows platform generations	Software and infrastructure can evolve on more independent cycles

Horizontal subscriber scalability

A COTS node or cluster can host multiple vBNG instances serving different subscriber populations or scaling domains. The unit of growth becomes a software workload rather than necessarily a new dedicated BNG chassis. With the fast innovation cycles in COTS Hardware and the related processing power increase growing capacity needs can be served within the same footprint and without the need of adding new chassis and racks.

Higher-value service compositions

The same horizontal infrastructure can potentially combine vBNG with functions such as CGNAT or IPsec VPN. This creates the option to build service chains on a common compute layer rather than on separate function-specific appliances. These examples are architectural deployment options, not benchmarked configurations in the present validation.

13. Decoupling Service, Software and Hardware Lifecycles

One of the strategic benefits of COTS-based network functions is the ability to reduce the dependency of service evolution on a specific proprietary hardware roadmap.

Traditional coupling

Service evolution is tightly coupled to network software and hardware roadmaps.




Figure 9: Traditional coupling between service, software and dedicated hardware roadmaps.

When a service depends on hardware-specific capabilities, introducing new functions can require support from a particular chipset or platform generation. Hardware qualification and product cycles can therefore influence service innovation and deployment timing.

Cloud model: decoupled evolution

Services, software and hardware lifecycles evolve independently on a shared, standard infrastructure.




Figure 10: Decoupling service, software and hardware lifecycles.

A software-defined BNG on general-purpose compute and networking resources establishes a more standardized infrastructure boundary. It does not eliminate vendor dependencies, but it can allow service software and the underlying compute platform to evolve on more independent timelines. In conjunction with SDN-based control and subscriber-steering mechanisms such as 3S, this separation can extend to the lifecycle of the BNG workloads themselves.

Different vBNG instances can be assigned to distinct subscriber populations, for example retail and business services, while running on the same COTS infrastructure. This creates the possibility to introduce, upgrade, scale or migrate broadband services according to their own requirements rather than forcing all services to follow the lifecycle of a common dedicated hardware platform.

The objective is not necessarily to create a multi-vendor BNG environment. The architectural value is the ability to manage service lifecycles more independently. A common COTS resource pool can support multiple software workloads, while SDN and subscriber steering provide mechanisms to direct subscriber populations toward the appropriate service instance. This can support progressive service introduction and migration without requiring the underlying infrastructure to be replaced at the same time.

Operator value

The goal is to decouple the broadband service lifecycle from the dedicated hardware lifecycle, creating more freedom to innovate in software and potentially improving time to service and time to revenue.

14. From Infrastructure Efficiency to TCO

The economic hypothesis is broader than comparing the purchase price of a server with the purchase price of a hardware BNG. A meaningful TCO model must capture the system-level effects of cloudification.

TCO lever	What should be quantified
Software efficiency	Compute required per subscriber and per unit of traffic under defined service profiles
Hardware efficiency	Server/NIC capacity, utilization and refresh assumptions
Power management	Measured power under idle, representative and high-load conditions
Infrastructure utilization	Impact of pooling, co-location and reduced stranded capacity
Horizontal scalability	Granularity and timing of capacity additions
Lifecycle independence	Different refresh cycles for software and COTS hardware
Operations	Automation, orchestration, support and lifecycle-management effort
Migration	Coexistence, transition costs and residual value of existing assets




Figure 11: From technical feasibility to production dimensioning and TCO assessment.

The present validation primarily establishes technical feasibility and provides CPU/resource-utilization inputs. Numerical power results were not included in the supplied test results, so the paper should not claim quantified energy or TCO savings. Those measurements and the broader utilization model are the next inputs required for the economic assessment.

15. Progressive Transformation and Asset Protection

A technically and economically attractive target architecture does not imply a big-bang replacement of the installed BNG estate. Cloudification can be introduced progressively and aligned with FTTH rollout, subscriber migration, capacity growth and hardware end-of-life. The transformation is therefore not simply a choice between dedicated hardware and a software BNG. It also introduces greater flexibility in how subscriber populations are placed across BNG instances and how those instances are evolved over time.

Subscriber steering as a migration and service-placement enabler

Subscriber Session Steering (3S) is positioned here as a complementary control mechanism for decoupling subscriber placement from a fixed BNG assignment. During migration, this can allow existing BNG capacity and new cloud vBNG capacity to coexist, with subscriber populations progressively directed toward the appropriate environment rather than requiring all subscribers to migrate at once. The Broadband Forum describes 3S (currently in working text stage, WT-474) as enabling dynamic, service-aware placement of subscriber sessions across available service gateways [2].




Figure 12: Subscriber steering enables progressive migration and service-specific vBNG placement on common COTS infrastructure.

The same mechanism can support a broader cloud operating model. Multiple vBNG instances can run on a common COTS infrastructure and serve different subscriber populations or service requirements. For example, residential and business subscribers could be assigned to separate vBNG instances, allowing each service environment to be dimensioned, upgraded and lifecycle-managed according to its own requirements without requiring dedicated physical hardware for each BNG. This illustrates an important aspect of cloudification: the infrastructure can remain common while the lifecycle of the service workloads becomes more independent.

This is also consistent with the broader 3S concept. Broadband Forum material describes policy-based subscriber placement and identifies enterprise and other specific market segments among the potential applications of session steering.

3S was not part of the laboratory benchmark, and no 3S performance or functionality conclusion is derived from the test results presented in this paper. It is included here as a complementary architectural mechanism for progressive migration, subscriber placement and service-specific lifecycle management.

Asset-protection principle

Cloudification can provide an incremental transformation path in which existing BNG assets and new vBNG capacity coexist, while subscriber steering enables services and subscriber populations to migrate and evolve according to their own lifecycle requirements.



CONCLUSION

Cloudify the Function, Preserve the Service

The BNG transformation should not be characterized simply as replacing an ASIC with an x86 processor. It is a transition from a network function vertically integrated with dedicated hardware toward a software-defined broadband workload running on standardized, horizontal COTS infrastructure. The objective is to establish a more hardware-independent infrastructure model, in which broadband service capabilities are not tied to a particular server, processor or proprietary hardware lifecycle.

That transition creates the potential to dimension service capacity more independently from a dedicated hardware platform, scale subscriber capacity horizontally, mutualize COTS infrastructure across instances and network functions, reduce dependency between service innovation and proprietary hardware roadmaps, and allow service, software and hardware lifecycles to evolve more independently.

The prerequisite is that these benefits must not come at the expense of the broadband service itself. The laboratory validation addressed this prerequisite at two levels: first by establishing the functional broadband baseline, and then by evaluating the service characteristics under production-relevant subscriber and traffic conditions. The tested implementation combined 6WIND vBNG software with Intel compute and Ethernet technologies on a Dell PowerEdge COTS server, demonstrating how complementary partner capabilities can support this more hardware-independent architectural model.

The functional validation demonstrated successful operation of the broadband capabilities exercised by the test program. PPPoE operation across IPv4 and IPv6, concurrent IPv4 and IPv6 services, subscriber authentication, NAS functionality, subscriber and access-circuit identification, and defined broadband access-loop attributes all passed the functional scenarios. The subsequent performance validation then demonstrated that these results were complemented by production-relevant throughput, subscriber scale, multi-queue QoS and

average latency within the defined service objective.

The approximately 100G connectivity envelope exercised by the validation was reached using only a subset of the compute resources available on the server platform. This demonstrates that the tested 100G-class workload did not exhaust the server's compute capacity. The validated workload therefore provides a building block for a modular $n \times 100\text{G}$ scaling model. Higher aggregate capacity can be addressed vertically by allocating additional processing and networking resources to a vBNG workload, horizontally by deploying multiple vBNG instances on the same COTS platform where resources permit, or by extending those instances across additional COTS nodes.

These approaches can also be combined according to service demand and infrastructure capacity. Actual aggregate throughput depends on compute, network I/O, memory resources, workload isolation and service configuration, and must be validated for the intended deployment.

Technical evidence for the cloudification model

The validation is not the destination of the analysis. It provides the technical evidence that makes the broader cloudification model credible: the tested BNG functionality and service characteristics can be preserved while moving the function from dedicated infrastructure to a software workload on COTS.

The next step is to translate technical feasibility into production dimensioning and a wider TCO model that quantifies software efficiency, power-aware hardware behavior, horizontal scaling, resource pooling, network-function co-location, operations and migration economics.

References and Further Reading

- [1] 6WIND, Virtual BNG product documentation.
Product documentation and deployment guidance for the 6WIND Virtual Broadband Network Gateway used as the software BNG architecture discussed in this paper.
[6WIND product documentation](#)
- [2] Broadband Forum, "Steer Smarter. Deliver Better. 3S Overview."
Overview of the Subscriber Session Steering (3S) architecture and its role in dynamic subscriber placement, resource optimization and subscriber experience.
[Broadband Forum 3S Overview](#)
- [3] Intel® Ethernet Controller E810.
Technical information for the Intel® E810 Ethernet controller family, including support for up to 100GbE interfaces and SR-IOV capabilities relevant to the validation infrastructure.
[Intel® Ethernet E810 Controllers](#)
- [4] Intel® Ethernet Controller E800 Series Dynamic Device Personalization (DDP) for Telecommunications.
Technical guide describing how Intel DDP extends Ethernet controller packet-processing capabilities to recognize telecommunications protocols, including PPPoE, and support efficient hardware-assisted packet classification and distribution.
[Intel® Ethernet E810 DDP technology guide](#)
- [5] Intel® Xeon 6700 with P-cores Product Brief.
Product brief outlining the performance, multicore scalability and power-efficiency capabilities of Intel® Xeon 6700 processors with P-cores, supporting the dimensioning of demanding software-defined network functions on COTS infrastructure.
[Intel® Xeon 6700 Product Brief](#)
- [6] Dell Technologies, PowerEdge R660 Technical Guide.
Platform architecture and technical specifications for the Dell PowerEdge R660 server family used as the COTS server platform in the validation.
[Dell PowerEdge R660 Technical Guide](#)