Solution Brief

Telco Edge Cloud, MEC, Wireline Broadband Internet
Cloud Native BNG

Achieving 600 Gigabits of BNG Performance in a Standard 2U Form Factor Server

Authors

Sudhir Krishnan
Casa Systems

Tim Miskell
Intel Corporation

Paul Mannion
Intel Corporation

Andrew Duignan
Intel Corporation

In order to meet the continually increasing performance requirements for Software Defined Networking (SDN) and Network Function Virtualization (NFV) use cases in terms of throughput and power consumption Intel Corporation and Casa Systems*, have been working jointly to address the challenge.

The 3rd Generation Intel® Xeon® Scalable processors provide several feature enhancements specifically targeted for networking workloads. Such enhancements include Intel® Ethernet Adapter enhancements for processing broadband access traffic, support for PCIe Gen4, along with improved Intel® Advanced Vector Extensions 512 (Intel® AVX-512) performance. All greatly aid with the high-speed data plane processing requirements associated with BNGs.

The following solution brief describes the impacts on a virtualized Broadband Network Gateway (vBNG), a key VNF workload for broadband aggregation that requires not only high throughput and resiliency, but also efficiency in terms of power consumption.

Specifically, this document highlights the performance of the Casa vBNG AxyomTM workload on 3rd Generation Intel® Xeon® Scalable processors, generating up to 600 Gbps per server at less than 0.5% packet loss with variable power consumption based on the operating workload.

A Broadband Network Gateway (BNG) is typically deployed in a Central Office (CO) at edge locations close to end users and serves to aggregate multiple home and office locations while providing access to the core of the network. The portion of the network connecting to home and office locations is referred to as the access network, while the portion of the network connecting to datacenters is referred to as the core network. Packets emanating from the access network destined for the core network are collectively referred to as upstream (US) or uplink (UL) traffic. Packets emanating from the core network destined for the access network are collectively referred to as downstream (DS) or downlink (DL) traffic.

In terms of typical broadband access protocols, the network cards must be capable of supporting IP-over-Ethernet (IPoE) traffic patterns. Upstream traffic typically consists of Q-in-Q encapsulated IPv4 or IPv6 packets consisting of an outer service provider VLAN (S-VLAN) tag and an inner customer VLAN (C-VLAN) tag that allow routing to specific locations within the access network. The ratio of downstream to upstream traffic typically favors downstream traffic and oftentimes ranges from 8:1 to 9:1 for the DL:UL traffic mix.

Additional traffic shaping is typically performed by BNGs in order to implement specific Quality-of-Service (QoS) along with differentiated services for specific subscribers. Bandwidth rate limiting may be performed in order to differentiate specific tiers of subscribers, for example home versus office locations, customers that may opt for higher data rates, as well as higher priority traffic classes carrying VoIP data. BNG deployments also typically include support for firewalls and ACLs in order to limit not only which access network locations can connect to the BNG, but also which TCP/UDP ports are allowed within the access network or home.

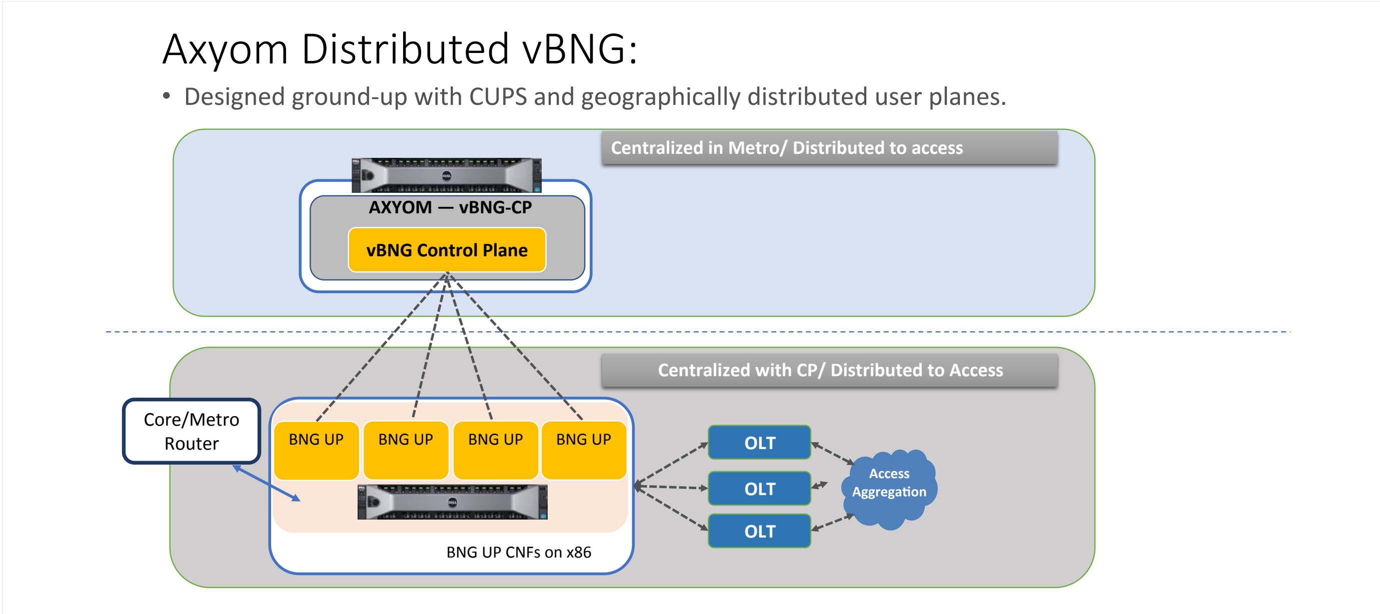

Figure 1. Casa Axyom Distributed vBNG – Control and User Plane Separation (CUPS). Source – Casa Systems

Historically BNGs have been deployed as fixed function hardware. However, with the advent of NFV, BNGs have been virtualized, i.e., as a VNF orchestrated and deployed in a Virtual Machine (VM), as well as containerized, i.e., as a CNF orchestrated and deployed in a container pod. Deploying virtualized or containerized appliances allows for running the VNF/CNF on standard dual socket IA server.

Casa Systems Axyom Virtual BNG solution enables network evolution towards a Cloud Native Distributed BNG with Control and User Plane Separation (CUPS) compliant with TR-459 spec from Broadband Forum. Casa is a pioneer in the field of CUPS based BNG starting as early as 2018 in our product development and commercialization. We offer unmatched subscriber scale and high throughput on dis-aggregated commercial off the shelf hardware with Intel® Xeon® CPUs. Our BNG is future proof with support for Wireline-Wireless Convergence (WWC) standards. Bringing in knowledge from our Wireless side of the house where the benefits of CUPS were well understood Casa built our vBNG solutions ground up with control and user planes as separate components that could run on the same server or separated in different locations. For more information, please visit us at: https://www.casa-systems.com/solutions/bng/

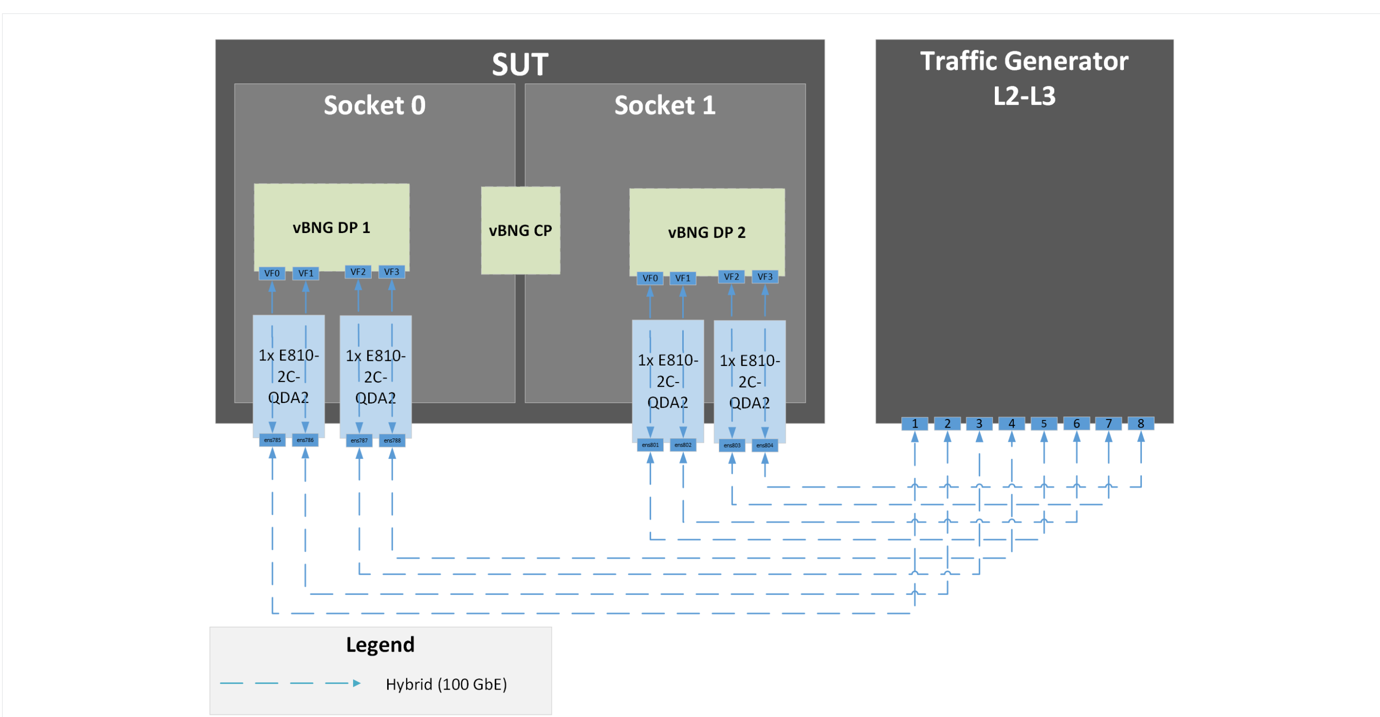

Figure 2. Network Topology Overview

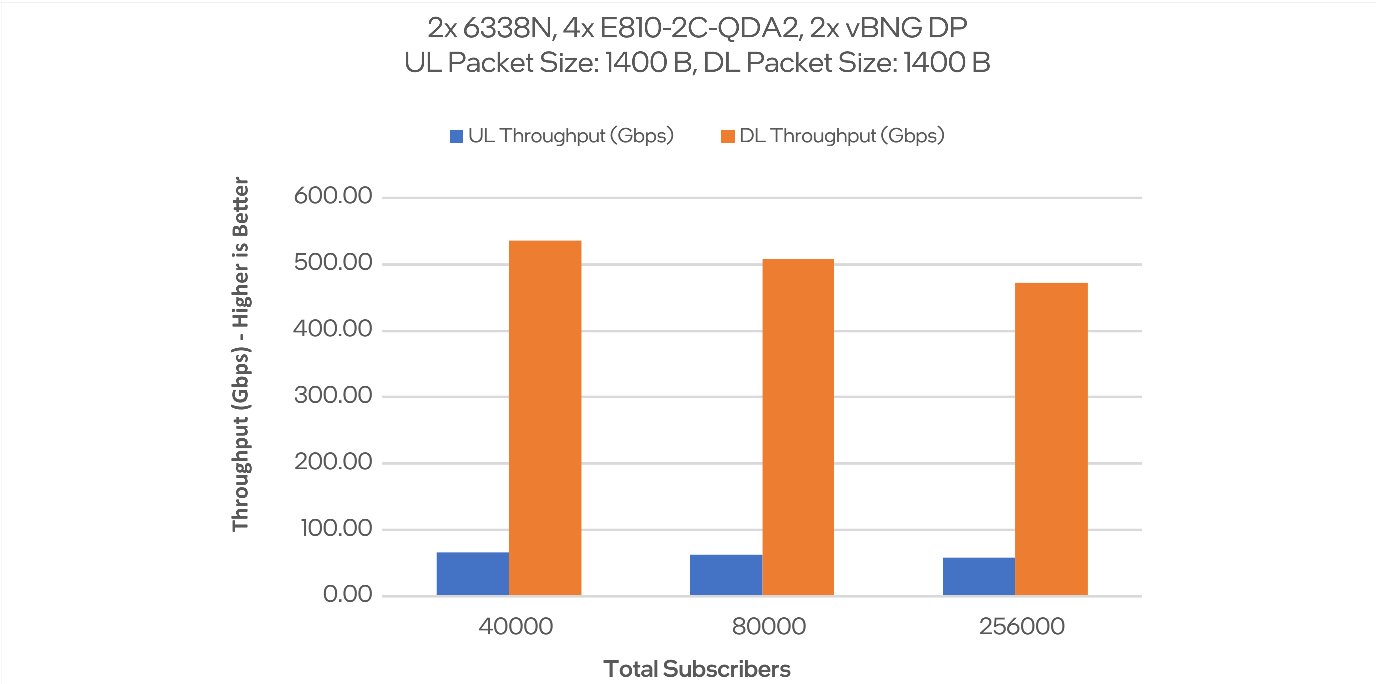

Figure 3. Uplink and Downlink BNG Performance tested by Intel on 9/29/2022

The server complex is composed of 2 x Intel® Xeon® Gold 6338N CPUs each of which terminates 400GbE of I/O provided by 2x200GbE Intel® Ethernet E810 Adaptor cards. The platform is capable of processing 800GbE of data.

Using a typical asymmetrical broadband access profile (9:1) Intel and Casa recently completed testing of the Casa Axiom BNG on a dual socket Intel Xeon Gold 6338N M50CYP 2UR208 CRB server showing a combined uplink and downlink throughput capability of up to 601.96 Gbps across both CPUs.

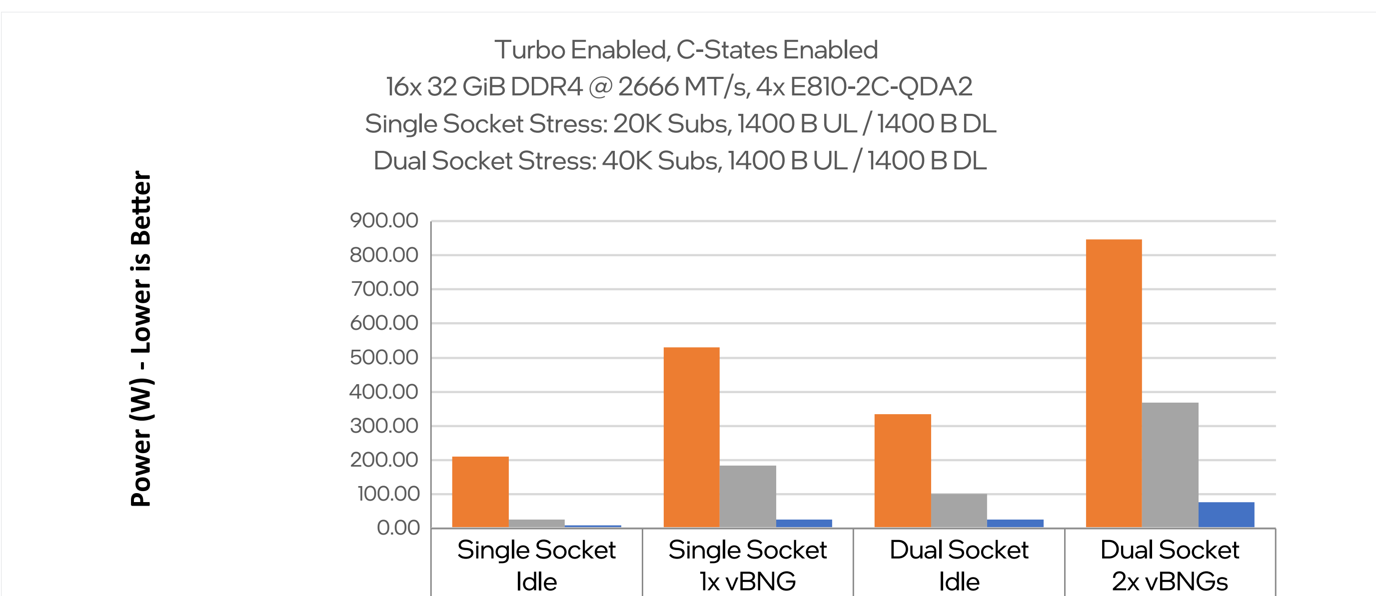

	Single Socket Idle	Single Socket 1x vBNG	Dual Socket Idle	Dual Socket 2x vBNGs
Total Wall Power (W)	210.98	531	337.72	846
CPU Power (W)	25.45	184.36	99.20	368.89
DRAM Power (W)	9.66	26.05	27.08	75.85

Figure 4. CPU, DRAM, and Server Power Consumption for Single and Dual Socket Topologies. Tested by Intel on 9/29/2022

While the performance improvement is very significant, the power savings is also just as important to service providers as they look to reduce their thermal footprint. The results show power consumption scaling from approximately 211 Watts to 531 Watts for single socket servers, and 338 Watts to 846 Watts for dual socket servers where the BNG is processing to its maximum capability. This provides an excellent opportunity for the network operators. As traffic varies in a broadband network, the network manager can realize significant energy savings by scaling compute capacity up or down based on predictable traffic models.

Conclusion

This solution brief demonstrates the performance of the Casa vBNG AxyomTM on 3rd Generation Intel® Xeon® Scalable Processors along with Intel® 800 Series Ethernet Controllers. In this case, the Casa vBNG AxyomTM is capable of generating up to 300 Gbps of throughput per NUMA node and 600 Gbps per server at less than 0.5% packet loss within a single server platform (see Figure 3 above).

Specifically, the server platform leverages x16 PCI express Gen 4 together with Intel® 800 Series Ethernet Controllers, which offers full 2x 100 GbE support within a single NIC. As a result, deploying the Casa vBNG workload on a server with 3rd Generation Intel® Xeon® Scalable Processors and E810-2C-QDA2 NICs provides an improved base level of performance and power utilization, VNF consolidation within a single server capable of up to 800 GbE near the edge, all with significant potential for increased rack server density.

1.1 Hardware and Software BOM

The table below details the Hardware BOM for the SUT.

COMPONENT	SPECIFICATION
Platform	Intel® M50CYP
Central Processing Unit (CPU)	2x Intel® Xeon® Gold 6338N CPU 32c 2.2 GHz 185W
Memory	16x 32 GiB (512 GiB) DDR4 2666 MT/s
Network Interface Card	4x 100 GbE Intel® Ethernet E800 Series E810-2C-QDA2
LAN On Motherboard (LOM)	1x Intel® Ethernet 700 Series X722-DA2
OS Drive	Intel® SSDSC2KB24 240 GiB Drive
Storage Drive	2x Intel® SSDSC2KB48 480 GiB Drive
BIOS	SE5C620.86B.01.01.0006.2207150335
Microcode	0xd000375

Table 1. Hardware Bill of Materials. Tested by Intel on 9/29/2022.

https://builders.intel.com/docs/networkbuilders/intel-select-solutions-for-nfvi-v3-with-red-hat-openstack-platform.pdf presents the software stacks for Red Hat Enterprise Linux* (RHEL*), which is based on Intel® Select Solution for Network Function Virtualization Infrastructure (NFVI) on Red Hat* Enterprise Linux* and Red Hat* OpenStack* Platform Reference Architecture, Document number 606324. With reference to https://builders.intel.com/docs/networkbuilders/intel-select-solutions-for-nfvi-v3-with-canonical-ubuntu.pdf , note that the SUT may also be deployed on Intel® Select Solution for Network Function Virtualization Infrastructure(NFVI) on Canonical Ubuntu* v3 Reference Architecture, document number 606325.

SOFTWARE STACK	RHEL8.6
Host OS	4.18.0-372.19.1.el8_6.x86_64
Libvirt	6.0.0
QEMU	4.2.0
NVMe	1.0
vBNG DP Version	3.3.1
vBNG CP Version	3.3.1
DPDK	21.11
Ice	1.9.11
i40e	4.18.0-372.19.1.el8_6.x86_64
iavf	DPDK 21.11
Ixgbe	4.18.0-372.19.1.el8_6.x86_64

Table 2. Software Stack Configuration for System Under Test. Tested by Intel on 9/29/2022.

1.2 vBNG Test Workload Topology (Test Configuration)

The following table presents the resource allocation for the vBNG CP VM.

COMPONENT	SPECIFICATION
vCPUs	2C4T
Memory	24 GiB
Storage	15 GiB
MGMT Interface	1
DP Interface	N/A

Table 3. Resource Allocation per vBNG CP VM. Tested by Intel on 9/29/2022.

The following table presents the resource allocation for each of the vBNG DP VMs.

COMPONENT	SPECIFICATION
vCPUs	24C48T
Memory	50 GiB
Storage	15 GiB
MGMT Interface	1
DP Interfaces	4x 100 GbE VFs (1x VF per PF)

Table 4. Resource Allocation per vBNG DP VM. Tested by Intel on 9/29/2022.

The following table presents the traffic profile configured on the traffic generator as part of the benchmark. In this case, the traffic profile consists of a range of subscribers, from a total of 40,000 subscribers up to 256,000 subscribers. Note that all ports are configured in Hybrid mode, in which case each Virtual Function may receive and process both access and core network traffic.

COMPONENT	SPECIFICATION
Total Access Network Subscriber Count	{40,000, 80,000, 256,000}
Core Network Endpoint Count	254
Access Network Packet Type	Q-in-Q Encapsulated DHCPv4/IPv4 IpoE
Core Network Packet Type	VLAN Encapsulated Ipv4
MGMT Interface	1
DP Port Configuration	Hybrid Mode (Access & Core)
Number of Network Topologies	8x 100 GbE Networks
Number of Ports per Network Topology	1
UL:DL Line Rate Ratio	8:1
UL/DL Traffic Mix	{1400 B/1400 B}

Table 5. Network Topology Configuration Settings. Tested by Intel on 9/29/2022.

The following figure presents the vCPU mapping for the CP VM along with both of the vBNG DP VMs. Note that in this case Hyper-threading (HT) is enabled on the platform, with vCPUs allocated to each of the vBNG DP VMs in HT sibling pairs. In addition, note that 4C8T are reserved for the hypervisor.

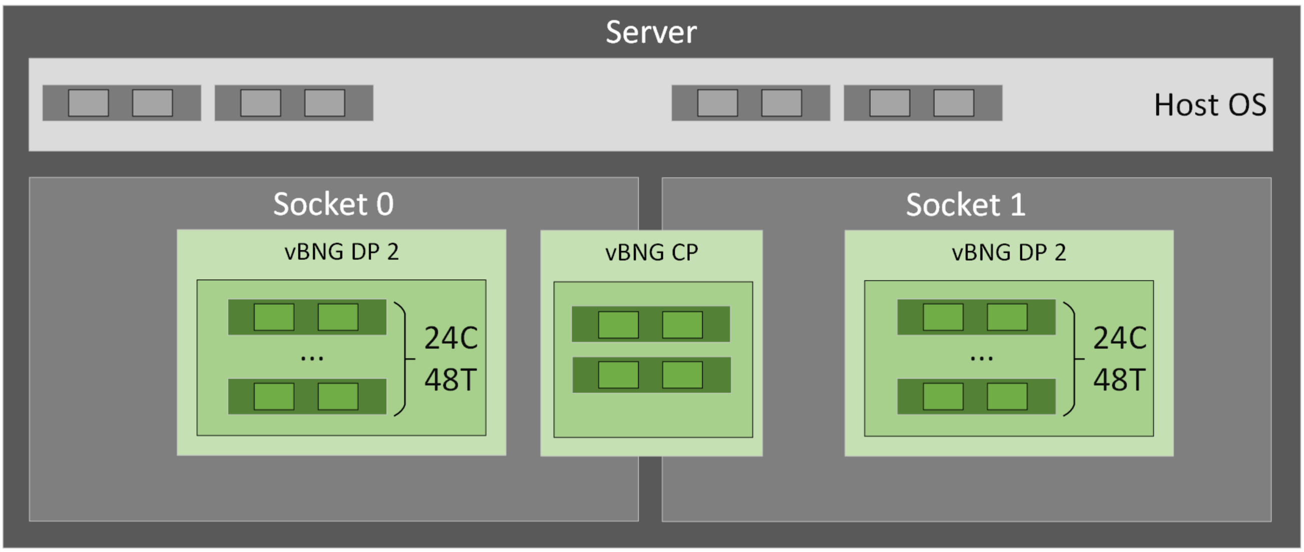

Figure 5. vBNG CP and DP vCPU Mapping. Tested by Intel on 9/29/2022.

The following table presents the Key Performance Indicators (KPIs) along with the boundary conditions for each of the corresponding KPIs where applicable.

KPI	UNITS / BOUNDARY CONDITION
Throughput	Gbps / RFC2544 at less than 0.5%
Packet Loss	% / Less than 0.5 %
Latency	µs / Less than 500 µs