

Virtual Broadband Network Gateway (vBNG)

Subscriber Management for the Software-Defined Era

netElastic vBNG is a high performance and cost-effective broadband network gateway that enables service providers to deliver network services faster and scale more efficiently. vBNG also includes a set of rich and flexible advanced subscriber management features to help service providers develop customized services based on customer needs.

vBNG is deployed at the subscriber network edge, in the same fashion as traditional BNG or BRAS. It controls subscriber access to the network, while performing management functions such as authentication, authorization, and accounting (AAA), user IP address management, subscriber management, access security, policy management, and Quality of Services (QoS). The diagram below shows a standard deployment of vBNG at the first aggregation point in the subscriber access network.

HIGHLIGHTS

- Up to 160 Gbps throughput per node in 2RU form factor
- Multi-node scalability up to 3 Tbps in a full-rack 42RU
- Rich feature set with IPoE, PPPoE, and L2TP capabilities
- 128K concurrent sessions, with horizontally scalable architecture for millions of sessions
- 2 million maximum NAT sessions
- 100K new NAT sessions per second
- NAT throughput at line rate
- Maximum 128K QoS queues
- Bare metal and virtual platforms
- NetConf and YANG
- CLI, SSHv2, CLI 5 privilege levels, CLI logging, and SNMP

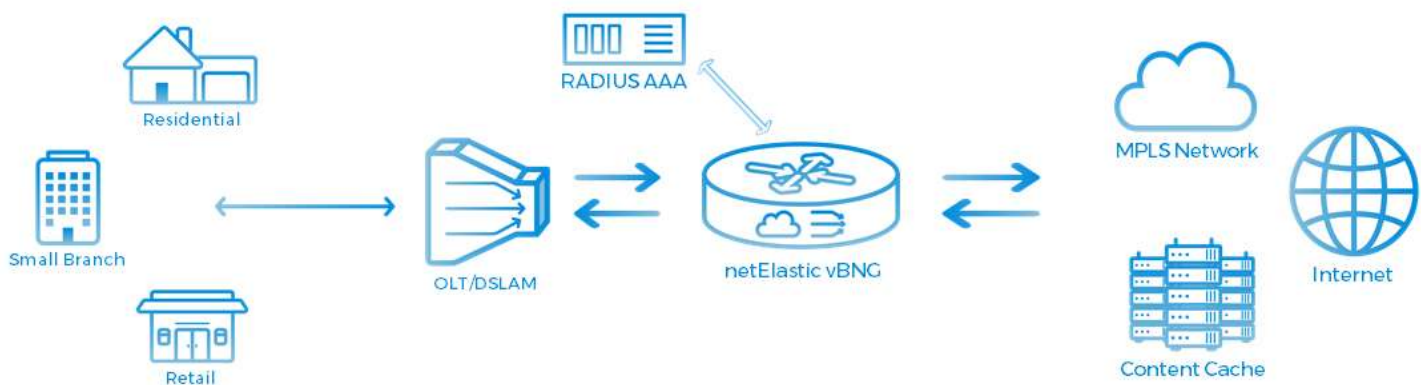


Figure 1 - netElastic Virtual BNG



FEATURES AND BENEFITS

FLEXIBILITY AND SCALABILITY

Service providers require flexibility in their networks to meet ever-changing business and customer requirements, and to pursue new market opportunities. As the central point for subscriber and service management, broadband network gateways are crucial to service differentiation, revenue generation, and customer experience. Unfortunately, most traditional BNGs are built on dedicated hardware platforms that come with high costs and limited scalability. Hardware-based BNGs also encourage over-provisioning, which results in large upfront costs that aren't matched by customer revenue.

netElastic vBNG is designed to deliver unmatched performance and scalability with a highly optimized data plane and patented packet processing techniques. vBNG separates the control plane and data plane, with the control plane running the network operating system on a virtual machine containing a comprehensive suite of network service functions. The data plane runs the packet-forwarding engine, enabling line-rate data transfer.

Both the control plane and data plane can scale up/down to meet performance needs, and they can run in different hosts to take advantage of CPU cores to maximize performance and scalability. The netElastic vBNG data plane can also be expanded across multiple hosts to obtain the performance of dedicated hardware BNGs, while still benefiting from the greater flexibility and lower-costs of a virtualized solution.

ROBUST SUBSCRIBER MANAGEMENT CAPABILITY

vBNG's advanced subscriber management capabilities lower operational costs and accelerate time to market by providing flexibility for access, authentication, and QoS offerings. vBNG is designed to support high concentrations of subscribers and supports all common BNG access features such as PPPoE and IPoE, and also supports subscriber IP address management functions such as DHCP, DHCP Proxy, DHCP Relay, and DHCP Option 60. Each single control plane can support up to 128K users concurrently and throughput can be scaled by expanding the data plane across additional nodes as necessary.

COMPREHENSIVE PROTOCOL SUPPORT

The core of the netElastic vBNG is a powerful routing engine that supports full Layer 2 and Layer 3 functions. These protocols include VxLAN, GRE, QinQ, L2TP for Layer 2, and OSPF, ISIS, BGPv4, MP-BGP, Route Policy, Dual Stack, and Graceful Restart for Layer 3. In addition to IPv4, vBNG also supports various IPv6 protocols, such as IPoEv6, PPPoEv6, OSPFv3 and BGP4+. This comprehensive list of protocols makes vBNG an ideal technology for numerous use cases.

HIGH PERFORMANCE AND THROUGHPUT

netElastic vBNG leverages Intel's Data Plane Development Kit (DPDK) and Single Root IO Virtualization (SR-IOV) coupled with netElastic optimized code to enhance forwarding performance. vBNG delivers up to 160Gbps throughput on a single server and can be scaled to support multi-terabit services. The data plane can also be scaled-out across multiple servers to increase total bandwidth capacity.

CARRIER GRADE SOLUTION

vBNG's comprehensive subscriber management includes Carrier Grade Network Address Translation (CG-NAT) that handles 2 million sessions per instance and can initiate 100K NAT sessions per second. This high throughput makes vBNG an ideal choice for high density network environments that require high performance.



netElastic designed vBNG with carrier-grade reliability at its core. Its full management and orchestration layer (MANO) initiates and monitors vBNG, and combined with support for Stateful Switchover (SSO) ensures high reliability and high availability. Incrementally adding physical compute and network resources (as you need it) results in lower, linear cost scaling that aligns with business growth. vBNG's open APIs allow for OSS/BSS integration.

TRAFFIC MANAGEMENT

netElastic vBNG has a fully capable traffic classification system based on TOS, IP Precedence, DSCP, VLAN, as well as rate limiting and congestion management at the subscriber ingress. By applying Hierarchical QoS, vBNG can enable per subscriber queues with multiple services, or a single service for multiple subscribers. This more effectively manages Committed Information Rate (CIR) and Peak Information Rate (PIR) for subscriber bandwidth control.

H-QoS multi-level scheduling supports traffic classification policy, scheduling policy, and congestion management. The traffic classification policy is based on VLAN, IP Precedence, and ACL. The scheduling policy can be static, using Committed Access Rate (CAR) to assign bandwidth, or can be dynamically assigned using Class Based Queuing (CBQ). Class Based Queuing uses IP Precedence (or DSCP), ingress interface, 5-tuples, and MPLS EXP values. This effective traffic management enables operators to allocate the appropriate bandwidth for each specific service that subscribers use.

SIMPLE ADMINISTRATION AND MANAGEMENT

Using Command Line Interface (CLI), an administrator can leverage familiar CLI syntax to manage vBNG and can assign up to 5 privilege levels for other users. In addition, vBNG comes equipped with modern provisioning tools such as YANG and NetConf to remotely provision the system. All network management sessions can be securely managed using SSHv2 and RADIUS, and vBNG supports SNMPv1, SNMPv2 and SNMPv3.

REDUCE OPERATING COSTS

netElastic vBNG is a high-performance Virtual Network Function (VNF) that can be deployed on a bare-metal commercial-off-the-shelf (COTS) server to lower costs per subscriber and costs per megabit. vBNG can run on KVM, OpenStack, and Wind River Titanium, among others. A standard implementation provides up to 160Gbps throughput and up to 70% in cost savings compared to traditional BNG vendors.

MONITORING AND REPORTING

vBNG provides complete logging with Syslog functions, and CLI commands at different privilege levels can be logged as well. The system also provides detailed logging and record keeping capability for NAT users' activities which can be logged to an independent server for regulatory compliance and lawful inspection.

BNG FEATURE SUPPORT

AAA

- Local authentication/authorization
- RADIUS authentication/authorization/accounting
- RADIUS server group
- RADIUS load balancing
- Web redirect templates
- Access Domain
- RADIUS Attributes supports
- Calling Station ID attribute

SUBSCRIBER MANAGEMENT

- IPoEv4
- IPoEv6
- PPPoEv4
- PPPoEv6
- Web Authentication
- Circuit Authentication
- Lease-Line Access
- DHCPv4
- DHCPv6
- DHCP Option 60
- DHCP Proxy
- DHCP Relay
- Local IP Address Pool

LAYER 2 AND TUNNEL PROTOCOL

- VLAN – IEEE 802.1q
- QinQ – IEEE 802.1ad
- L2TP
- LAC
- LNS
- VxLAN
- GRE

QOS

- Committed Access Rate (CAR)
- Traffic Classification (based on TOS, IP Precedence, DSCP, VLAN, ACL (L3/L4))
- Traffic Policing
- Rate Limiting
- Hierarchical Policing
- QoS Template
- Class-map
- Policy-map
- QoS and Traffic Policy

LAYER 3 PROTOCOLS

- Static Route
- RIPv1
- RIPv2
- OSPF v1
- OSPFv2
- OSPFv3
- OSPF Grace Restart (GR)
- ISIS
- ISIS GR
- BGPv4
- BPG4+
- MP-BGP
- BPG GR
- Route Policy
- IPv4
- IPv6
- Dual Stack

IP TRANSITION (NAT)

- Network Address Translation (NAT)
- Port Address Translation (PAT)
- NAT Log when new and when trigger by user flow
- Static NAT
- CG-NAT (Large Scale NAT)
- Detailed logging and record keeping in Large Scale NAT

MPLS

- MPLS (RFC-3031, RFC-3031)
- Label Edge Router (LER), Label Service Router (LSR)
- Label Distribution Protocol (LDP)
- Static and dynamic Label Switch Path (LSP)
- Local and remote LDP session
- LDP GR
- MPLS L2VPN – VPWS
- MPLS L2VPN – VPLS
- MPLS L3VPN
- L3VPN CE
- L3VPN PE
- VPLS LDP Signaling (RFC 4762)

IP MULTICAST

- IGMPv1
- IGMPv2
- IGMPv3
- PIM-SM
- PIM-SSM
- Multicast VLAN

SECURITY

- L3 ACL
- L4 ACL
- ACL6
- IPSEC
- IPSEC AH
- IPSEC ESP
- IKEv1
- uRPF
- Applying ACL to subscriber session

HIGH AVAILABILITY

- Bidirectional Fault Detection (BFD)
- BFD for Static Routes
- BFD for OSPF
- BFD for ISIS
- ECMP
- Non-Stop Forwarding

MANAGEMENT

- CLI
- SSHv2
- CLI privilege levels (5 levels)
- Allow commands based on privilege level
- CLI command logging
- SNMPv1
- SNMPv2
- SNMPv3
- Syslog
- NetConf
- YANG
- NTPv4
- Error message notification
- Ping & Traceroute



PERFORMANCE AND SPECIFICATIONS

USER ACCESS

IPv4 / IPv6 Dual stack
IPoE: 128K (1500 sessions/sec)
PPPoE: 128K (1500 sessions/sec)
L2TP: 64K maximum

INTERFACES

802.1Q Interfaces: 4096
QinQ Interfaces: 4096
Total Layer 3 Interfaces: 4096

FORWARDING TABLES

FIB 1M (IPv4) 50K (IPv6)
MRIB 4K
LIB 64K

ACL & HQOS

Max ACL: 1,000
Max Rules per ACL: 4,000
Max Rules: 32,000
Max Class Mappings: 1,000
Max Policy Mappings: 256
Max HQoS Queues: 32,000

TUNNELS & OVERLAYS

Max L3VPN: 1,000
Max VPLS: 1,000
Max VPWS: 1,000
Max VXLAN: 4,000
Max IPSEC SA: 4,000
Max IPSEC Tunnels: 8,000

MULTICAST

Multicast Groups: 4,000
Max User Side: 128,000
Max Member Egress: 8

THROUGHPUT PER CHASSIS (SINGLE CP & DP)

Pure Layer 3 Forwarding: 100 Gbps (≥ 256 byte/packet)
With Standard Services: 80 Gbps

CGNAT

NAT44 Forwarding: 100 Gbps (≥ 256 byte/packet)
With Standard Services: 80 Gbps
100,000 NAT44 sessions/sec
2 Million NAT44 Sessions Total

SOFTWARE REQUIREMENTS

Operating System: CentOS 7.4, Red Hat Enterprise 7.4, and Ubuntu 16.04
Hypervisor: QEMU-KVM 2.0.0
Virtualization Platforms: Wind River Titanium Server, KVM, OpenStack

MINIMUM HARDWARE REQUIREMENTS

Processor: Intel Xeon E5-26XX with 8 cores
Memory: 32GB
Hard Drive: 120 GB
Network:
2+ DPDK compatible 10Gb NIC (*Data Plane*)
2+ 1Gb NIC (*Control Plane & Management*)

RECOMMENDED PLATFORM

Single server platform to support 128,000 subscribers and 80 Gbps of total bandwidth





Processor: Dual Intel Xeon E5-26XX with 12 cores
Memory: 128 GB
Hard Drive: 120 GB
Network:
10 DPDK compatible 10Gb NIC (*Data Plane*)
5 1Gb NIC (*Control Plane & Management*)



MBUZZ Europe

Your netElastic Systems Partner in Europe

Contact Us

-  contact@mbuzzeurope.com
-  www.mbuzzeurope.com
-  [linkedin.com/company/mbuzzeurope](https://www.linkedin.com/company/mbuzzeurope)
-  twitter.com/MBUZZEurope

ABOUT NETELASTIC SYSTEMS

netElastic Systems is an innovative software company providing end-to-end solutions for network function virtualization and carrier infrastructure. Built on our extensive experience in the open source communities along with our patented data-plane technology, we deliver customized solutions that are optimized for networking performance, high availability, security, and operator efficiency. Visit us at www.netelastic.com