

Application Note: Server and VM Sizing for vBNG

netElastic vBNG is a software system consisting of a Control Plane and a highly-scalable Forwarding Plane. To get the desired performance from the vBNG, adequate server resources must be provided. Furthermore, if you plan to take advantage of the power of virtualization to run other Virtual Network Functions (VNFs) in the same server, you need to plan for that as well. This document will serve as a guide on how to configure such a server so you can get all the functionality and performance with a minimal capital outlay.

Understanding Sizing Factors

netElastic vBNG has two main components, the Control Plane (CP) and the Data Plane (DP). The Data Plane moves packets in and out, applying QoS and other policies along the way. It's sizing is based on how much bandwidth is needed and uses CPU Cores, network interfaces, and RAM. The Control Plane communicates with management tools, manages policies, establishes and updates routing tables for the Data Plane, handles AAA and most other functions that aren't involved with actual packet forwarding. The Control Plane sizing is generally based on the complexity of the use case and number of subscribers expected on the vBNG.

Deployment Options

vBNG can be deployed in several different manners described below.

Host Mode on Bare Metal – Entire vBNG running directly on a server. This is the most common option and uses the least amount of resources, avoiding virtualization layer overhead.

Host Mode in a Virtual Machine – Entire vBNG running in one VM. This is a good option for smaller vBNGs going into environments with existing virtualization capabilities and eliminates the need for a dedicated server.

CP/DP Split Mode – Two virtual machines, one for the Control Plane processes and one for the Data Plane processes. This option is generally reserved for very large-scale use-cases on two-socket servers. In this deployment mode, the Data Plane VM can be pinned to one of the CPUs and all CPU cores used for forwarding. The Control Plane can be put on the other CPU along with any other workloads.

Since the CP/DP Split Mode deployment is not common, the rest of sizing recommendations below will only focus on Host Mode.

Sizing Table

Product	Targeted Max Subscribers¹	Interfaces (N²)	CPU Cores for Host OS + Control Plane	CPU Cores for Data Plane (2 + N)	Total CPU Cores Required	RAM³	Storage⁴
BNG-64	4000	6x 1GbE	2	3	5	16GB	128GB
BNG-202	16000	2x 10G	4	4	8	32GB	256GB
BNG-402	32000	4x 10G	5	6	11	32GB	256GB
BNG-602	32000	6x 10G	5	8	13	32GB	256GB
BNG-802	64000	8x 10G	6	10	16	64GB	256GB
BNG-1202	128000	12x 10G	6	14	20	64GB	256GB
BNG-1602	128000	16x 10G	6	18	24	64GB	256GB

CPU Cores mentioned refer to physical cores, not hyper-threaded or virtual cores.

25Gbps and 40Gbps interfaces are also supported. To size for those, use the same table replacing 6x 10G for 2x 25G and 4x 10G for 1x 40G keeping the same CPU cores (e.g., 8 DP Cores for 2x 25G, 10 DP Cores for 2x 40G).

1 – Max Subscribers listed is based on the Cores and RAM listed.

2 – N is the number of 10Gbps equivalent interfaces. 1 for each 10GbE, 4 for 40GbE, etc.

3 – RAM should ideally be small (4GB+) increments able to fill all memory slots. While not always practical, it does have an impact on DPDK performance. (See [link](#) for details).

4 – Storage can be SSD or HDD. It is primarily used for system files and logging and not performance impacting. RAID 1 is recommended for HDD.

Future Sizing

It is often desirable to over-size the server to allow for future expansion. One can start with 8 or 12 cores for a minimal vBNG, but if you want to add more interfaces later, you should choose more CPU cores up front. Then add resources like memory and NIC cards in the future to scale up, minimizing your initial capital outlay. Multi-core CPUs get progressively more expensive so it's important to balance initial investments and growth plans.

You can also increase server resources to accommodate other workloads. The recommendations above are dedicated to the vBNG.

The NUMA Bus and NIC/PCI Considerations

For dual socket systems, the NUMA and NIC placement needs to be considered. Since the NUMA bus that connects the processors in a dual socket server is slower than the on-chip

pathways, it is best to keep all the Forwarding Plane (DP) cores assigned to the same 'socket', and use the other socket for the CP and other VNFs. The NICs need to be on PCI slots that share the same NUMA with the DP CPU.

Processor Speed

In order to maximize performance, especially in the DP, a minimum processor speed of 2.4 GHz is necessary. A higher clock speed will come in handy when you're processing mostly small packets or doing heavier packet processing such as ACLs and QoS. Slower clock speeds can have a linear impact on theoretical maximum bandwidth (e.g., 2.4GHz = 10Gbps, 2.2GHz is 10% less thus expected to max out at 9Gbps, etc.)

Memory

A minimum of 32GB of RAM is required to run BNGs with up to six 10G interfaces with 64GB recommended for larger use cases. Make sure to consider the requirements of any other VMs you plan to run on the server. As mentioned above, RAM should be installed in smaller DIMMs to fill memory channels to achieve advertised performance.

Disk storage

You can get away with very little storage – 256GB of storage is plenty. The vBNG automatically rotates and compresses internal logs. We recommend enabling Syslog for longer term log management and storage on an dedicated logging server.

NIC Cards

We currently require Intel branded NICs for all forwarding ports. Here are the recommended NICs:

- 10G SFP+/Copper X710 or X520/540 (82599 controller)
- 25G SFP+/Copper XXV710
- 40G SFP+/Copper XL710

VM Considerations

When running in a virtualized environment, either a single VM or the CP/DP Split Mode, the host OS will require its own core on top of the requirements listed above. And in CP/DP Split Mode, the DP VM will need an additional core for the guest OS on that VM.

Another important consideration for virtualized environments is how the physical NIC will be presented to the Data Plane processes. For maximum performance, you'll want to use PCI Pass-through so the NIC is directly controlled by the DP. SRIOV and even OVS are also supported, but since these modes are designed to allow NIC sharing with other workloads, they will be slower (significantly slower in the case of OVS). Packet latency will also rise in the shared NIC modes as individual BNG packets need to wait for their turn to use the interface.

Reference Servers

While netElastic prefers to not endorse any particular platform, we offer these Lanner Server Appliances as references:

netElastic Product	Lanner PN	Cores	Clock	RAM	Disk	Processor	Ports
BNG-64	L-1515A-8C-16E-128M-NT1	8	2.2 GHz	16GB	128 GB	Intel Denverton C3758	4 x GbE RJ45; 2 x SFP
BNG-202L	L-5520B-8C-32R-256-4XF-NT1	8	2.5 GHz	32GB	256 GB	Intel Xeon Silver Cascade Lake 4215	4x GbE RJ45; 4 x 10GbE SFP+ (2 useable)
BNG-402L	L-5520B-12C-32R-512-4XF-NT1	12	2.4 GHz	32GB	512 GB	Intel Xeon Silver Cascade Lake 4214R	4x GbE RJ45 4 x 10GbE SFP+
BNG-802L	L-5520B-16C-64R-512-8XF-NT1	16	2.9 GHz	64GB	512 GB	Intel Xeon Gold Cascade Lake 6226R	4x GbE RJ45 8 x 10GbE SFP+
BNG-1202L	L-5520B-16C-64R-512-12XF-NT1	16	2.9 GHz	64GB	512 GB	Intel Xeon Gold Cascade Lake 6226R	4 x GbE RJ45 12 x 10GbE SFP+
BNG-1602	L-5520B-24C-64R-512-8XF-2QF-NT1	24	2.4 GHz	64GB	512 GB	Intel Xeon Gold Cascade Lake 6240R	4x GbE; 8 x 10GbE SFP+ 2 x 40 GbE SFP+

[Lanner NCA-5520 Datasheet](#)