



HARDWARE ARCHITECTURE GUIDE FOR HPE VIRTUALIZED NONSTOP ON VMWARE



CONTENTS

Introduction.....	3
Scope.....	3
HPE Integrity NonStop X architecture.....	3
HPE Virtualized NonStop.....	4
HPE Virtualized NonStop on VMware.....	4
HPE Virtualized NonStop architecture guide.....	5
Sharing server hardware between multiple HPE vNS systems.....	11
Hardware implementation guide.....	11
Connectivity.....	13
Rack.....	14
VMware requirements.....	14
HPE Virtualized NonStop and support for hardware product-lines.....	14
Implementing HPE vNS Architecture Guide compliant hardware using HPE Synergy.....	15
Appendix A: HPE vNS hardware support matrix.....	16
Section 1: Server models.....	16
Section 2: Fabric NICs.....	16
Section 3: NICs supporting SR-IOV for network interface in IP and Telco vCLIMs.....	16
Section 4: NICs supporting PCI passthrough for network interface in IP and Telco vCLIMs.....	16
Section 5: Storage products usable with HPE vNS.....	16
Section 6: Ethernet switches.....	16
Appendix B: Storage considerations.....	17
Multipath access between CLIM and storage volumes.....	17
Appendix C: System requirements.....	19
Appendix D: HPE Virtualized NonStop—System configurations.....	20
Appendix E: Building hardware Bill of Materials (BoM)—A REAL-WORLD example.....	21
Step 1: Determine the number and type of physical servers needed.....	21
Step 2: Build the storage.....	24
Step 3: Networking the servers and storage.....	26
Step 4: Build the rack.....	28
Step 5: Software licenses.....	28
Connections.....	28
Fabric network.....	28
Storage Network.....	28
Maintenance Network.....	29
Management LAN.....	29
Enterprise LAN.....	29
Backup LAN.....	29
Rack Layout.....	31
Bill of materials.....	32
Fabric Connectivity for the HPE vNS system.....	33
Storage Networking Connectivity for the HPE vNS system.....	34
Small hardware footprint to host two HPE vNS systems!!.....	34
References.....	36



INTRODUCTION

The HPE Virtualized NonStop (vNS) system introduces a whole new way of implementing HPE NonStop solutions in today's enterprise IT. It allows the ability to deploy an HPE NonStop system as a guest in a virtualized IT infrastructure or in a private cloud environment. This opens up the implementation choices for HPE NonStop solutions to a wide variety of hardware products available in the market.

To support HPE NonStop fundamentals of high availability, scalability, and security, Hewlett Packard Enterprise requires the virtualized hardware environment to meet a set of rules so that the HPE vNS system offers the same features and benefits as available in the HPE NonStop converged system (HPE NonStop X).

This document describes the requirements and rules for deploying an HPE vNS system in a virtualized environment. The document is intended to help customers prepare the underlying environment and deploy HPE vNS systems in compliance to these rules and guidelines.

Scope

The scope of the document is to:

1. Specify the hardware components (such as server, storage, networking, and connectivity) of the infrastructure eligible to host an HPE vNS system in a generic fashion.
2. State the rules governing the distribution and configuration of HPE vNS virtual machines (VMs) on the virtualized hosts.
3. Provide information about hardware configurations that HPE vNS has been implemented on. This is essentially a reference to the readers of an HPE vNS system implementation in order to help them design systems for their specific requirements.
4. This document covers only the VMware® based virtualization environment.

This is a live document and will get updated periodically. The latest version of this document is available for download at hpe.com/info/nonstop-ldocs.

HPE INTEGRITY NONSTOP X ARCHITECTURE

Architecturally, HPE vNS closely follows the architecture of HPE Integrity NonStop systems. Hence let's first look at the architecture of an HPE Integrity NonStop X system at a high level.

HPE NonStop X is a system comprising servers running independent instances of software stacks in a "shared nothing" architecture. The servers run either an instance of the HPE NonStop OS (also called NonStop Kernel or NSK) or Cluster I/O Modules (CLIM) software. CLIM is an appliance built using HPE's version of Debian Linux®. They are connected together with a high-speed, low latency system interconnect (called the HPE NonStop fabric) implemented using remote direct memory access (RDMA) over an InfiniBand physical network. Storage is provided by drives in SAS JBOD drive enclosures or in an HPE XP Storage array connected by a SAN. The system provides a shared nothing, massively parallel architecture, which forms the foundation for the high availability and scalability features of the HPE NonStop systems. While the system is a cluster of servers running independent OS instances, to the outside world, it presents a single system image. In other words, it can be seen and addressed by outside client as one HPE NonStop system!

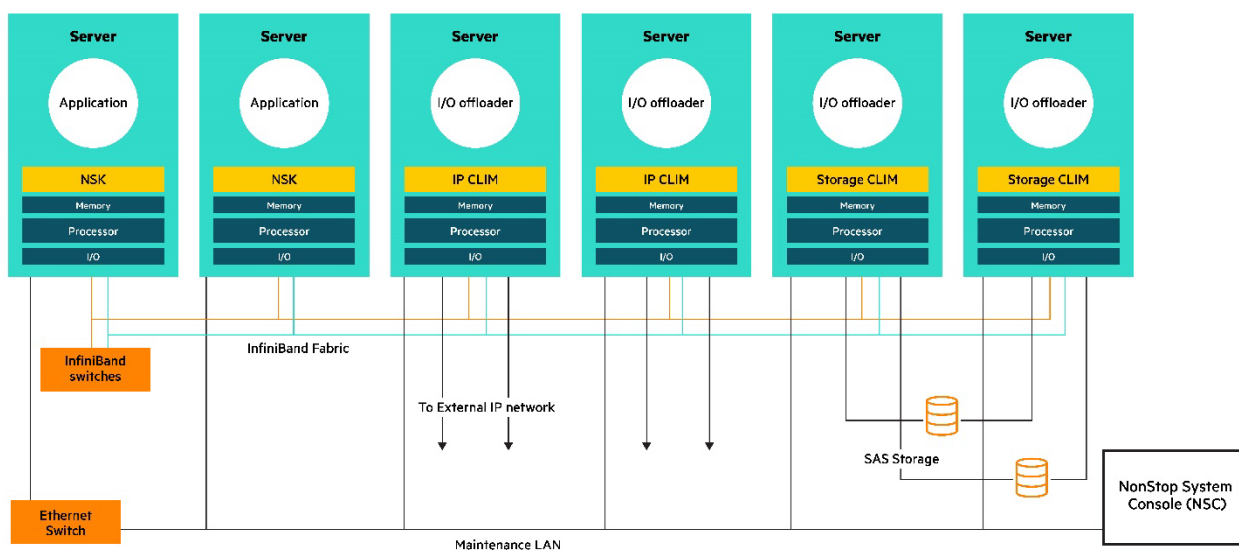


FIGURE 1. HPE Integrity NonStop X architecture diagram

A key feature of the HPE NonStop X architecture is redundancy against a single point of failure. The system interconnect consists of two independent physical fabrics. Each storage volume is provisioned in two mirrored drives, each of which is connected to two CLIM storage controllers to protect against failure of a single drive or a single CLIM storage controller. Network interfaces can be configured as failover pairs to provide continuous connectivity against failure of one of the interfaces. The HPE NonStop software too is highly fault tolerant. Two processes can be run on two separate logical processors in a primary-backup mode (called “process-pair”) with the primary process sending regular status updates to the backup process (a method called “check-pointing”). Such an architecture is the cornerstone of the near-continuous availability of HPE NonStop systems.

HPE VIRTUALIZED NONSTOP

HPE vNS is an implementation of the HPE NonStop X architecture in a virtualized and cloud environment. The HPE NonStop X architecture is equally applicable to an HPE vNS system except that the system fabric is implemented over Ethernet.

In a virtualized environment, the hypervisor software creates VMs, virtual storage volumes, and virtual network interfaces from a collection of physical server, storage, and networking hardware that comprise the cloud.

An HPE vNS system is a collection of VMs that provide processor, storage, and network functions that work in tandem and present a single system image to the outside world. The VMs belonging to an HPE vNS system are logically bound to each other by a high-speed, low latency system interconnect based on RDMA over Converged Ethernet (RoCE) v2. Figure 2 provides a logical view of an HPE vNS system in a virtualized environment.

Creating an HPE vNS system involves

- Creating VMs for HPE vNS CPUs, Storage vCLIMs and Network vCLIMs
- Establishing fabric connectivity between the VMs
- Provisioning the physical storage and networking resources required by the VMs

This cluster of VMs and its associated resources (but not the hardware they are hosted on) is brought under the management of one or a pair of HPE NonStop System consoles. The role of an orchestrator is critical for a clean HPE vNS system creation and eventual shutdown in an intuitive and user-friendly manner. An orchestrator is a tool available in cloud environments, which helps administrators to automate the tasks of VM definition, configuration, provisioning, sequencing, instantiation, connectivity, and such others through simple workflows aided by a powerful graphical interface.

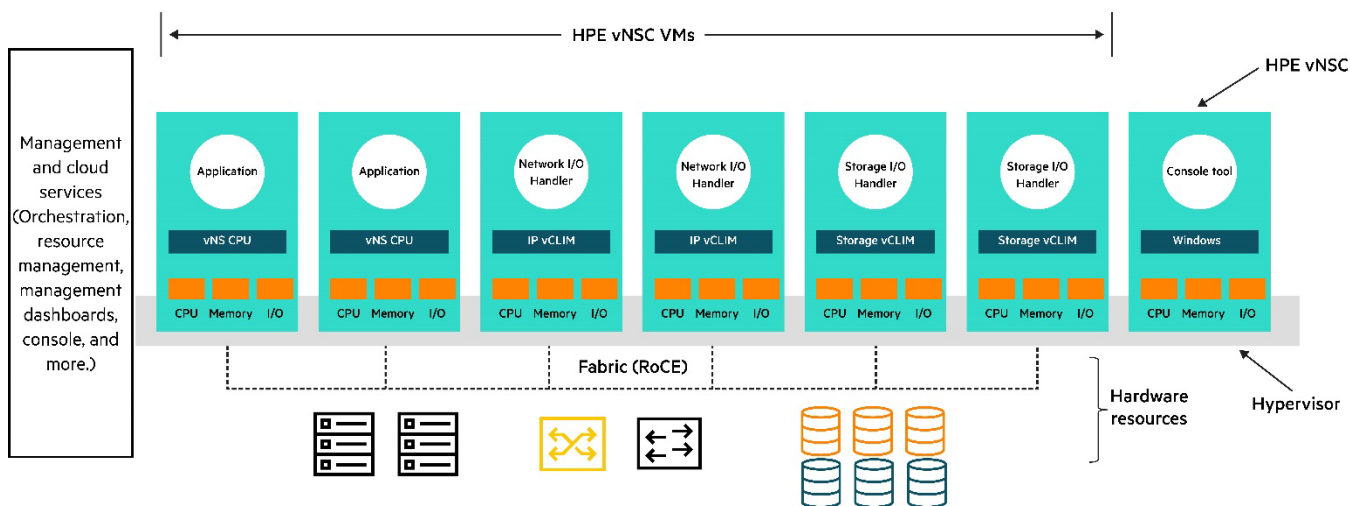


FIGURE 2. HPE NonStop deployed on a virtualized hardware—HPE Virtualized NonStop

HPE Virtualized NonStop on VMware

VMware is a popular virtualization and cloud software vendor in the IT industry. It has a dominant position in the virtualization market today, especially among large enterprises. VMware offers a wide set of software products under various suites such as VMware vSphere®, VMware vCenter®, VMware vRealize®, and VMware vCloud®. The foundational piece of VMware software is vSphere, which contains the VMware ESXi™ hypervisor and does the basic tasks of virtualizing physical servers and VM administration.

HPE vNS requires the following two VMware products:

1. VMware vSphere® Enterprise Plus Edition™: Of the various products bundled into this, HPE vNS requires VMware ESXi and VMware vRealize® Orchestrator™.
 - a. ESXi is the hypervisor that runs on physical servers and virtualizes them to create VMs. An HPE vNS system requires several VMs running different guest operating systems (HPE vNS CPU, vCLIM, and HPE vNS Console [HPE vNSC]) and works in tandem to present as a single system to the outside world.
 - b. VMware vRealize Orchestrator helps the task of creation, deployment, and configuration of these VMs, thereby easing the task of creating HPE NonStop systems and relieving the user of the involved tasks.

2. vCenter Standard Edition

vCenter is the VM manager in a VMware environment. It performs tasks such as VM configuration, resource assignment, monitoring, and several others for all the VMs under its span of control. HPE vNS does not require a vCenter instance to be dedicated to manage its VMs and resources. A HPE vNS system can be deployed and managed within an existing vCenter environment.

With this basic introduction to HPE vNS and VMware, let us now look at the guidelines for building an HPE NonStop system in a VMware environment.

HPE Virtualized NonStop architecture guide

Why the Architecture Guide (AG)?

HPE NonStop systems have a unique value proposition in the industry. HPE NonStop's high availability and linear scalability features have made it the platform of choice for mission-critical computing in today's enterprises. These benefits are made possible by the unique HPE NonStop architecture consisting of hardware and software elements and the tight coupling between them. HPE vNS offers the same advantages to a HPE NonStop system implementation in a virtualized environment built using industry-standard hardware.

An HPE NonStop X system comes to you with hardware and software integrated by the HPE factory from HPE factory. It has all the components correctly installed, wired, provisioned, configured, and ready to be deployed for the purpose you ordered it for. In contrast, an HPE vNS system is a software-only product offering. The deployment of the HPE vNS system takes place at the customer's site and in customer's IT environment.

There are a wide variety of choices available in the industry today to build a virtualized IT infrastructure and the layout of the elements (such as servers, VMs, fabric, and I/O) can be done in several different ways. Hence, it's important that:

1. The underlying hardware meets the requirements that would enable an HPE vNS system to be provisioned and configured correctly.
2. The HPE vNS system is deployed in a correct manner so that the system as a whole provides true HPE NonStop capabilities and value.

The objective of this AG is to ensure that the HPE vNS system implementation is able to offer the same benefits to the customer as an HPE Integrity NonStop system shipped from the factory.

While the HPE Pointnext Services organization offers professional services to do these tasks, it's important that the rules governing the implementation of HPE Virtualized NonStop systems are clearly understood by customers.

The HPE vNS AG specifies the requirements and guidelines that its implementation should meet for it to be supported by HPE. An HPE vNS system built-in compliance with this AG should provide the same availability and scalability advantages as does an HPE NonStop X system. The AG covers each component of an HPE vNS system such as HPE vNS CPUs (or interchangeably referred to as CPUs in this document), storage vCLIMs, network, system interconnect fabric, and HPE vNSCs.

The following sections explain the requirements for each of these HPE vNS elements.

Compute servers

As represented in the Figure 2, in an HPE vNS system, the HPE vNS CPUs and vCLIMs are VMs created in physical servers virtualized by hypervisors—ESXi in the VMware environment. The HPE vNS CPUs, vCLIMs, and HPE vNSCs run as guest operating systems inside these VMs and are distributed over several physical servers. The VMs for HPE vNS CPUs and vCLIMs are connected together by a system interconnect using Ethernet switches and adapters that support RoCE v2 protocol for low latency data transfer. Using Ethernet for the HPE vNS fabric enables HPE vNS to be deployed in standard IT environment where Ethernet is pervasive.

The AG for servers for deploying an HPE vNS system are:

1. The servers require one or more Intel® Xeon® x86 processors from one of the following processor families:
 - a. E5 v3 or E7 v3 (Haswell)
 - b. E5 v4 or E7 v4 (Broadwell)



- c. Scalable processor (Skylake)
- d. 2nd Gen Scalable processor (Cascade Lake)

HPE vNS uses Intel® Virtualization Technology available in these processor families to improve the performance and security of virtualization. For additional information, see intel.in/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html.

2. The processors used in the servers must support hyper-threading.
HPE vNS VMs use hyper-threading to improve performance. Hence, it's required that the processors used in these servers support it.
3. The servers should have adequate number of cores and memory to meet the requirements of all the HPE vNS VMs (HPE vNS CPUs and vCLIMs) to be deployed in the server, in addition to the requirements of the ESXi hypervisor. If supporting virtual appliances such as HPE vNSC, vCenter, vRO, and hardware management appliances (such as HPE 3PAR Virtual Service Processor and HPE OneView) are to be deployed on the same servers, consider the resource requirements of these as well.

Today's microprocessors have a many more cores than required by a single HPE NonStop CPU or CLIM. This feature combined with hardware support for virtualization allows multiple independent VMs to be deployed in a physical server. The VMs for the CPUs and CLIMs in an HPE vNS system can be deployed over a smaller number of physical servers thereby reducing the overall hardware footprint of the system.

NOTE

Throughout this document, when the term "core" is used in the context of a VM, a full processor core is implied (and not a hyper-thread).

4. HPE vNS VMs cannot share the processor cores with other VMs.
These VMs are highly sensitive to latency and require that the cores be dedicated to the HPE vNS VMs and not shared with others.
5. HPE vNS VMs require physical memory to be dedicated to the VMs.
For the same reasons as mentioned in (4) and for security, HPE vNS VMs require physical memory to be dedicated and not shared with other VMs.
See [Appendix A Section 1: Server models](#) for the list of Intel processor models and HPE servers that have been used for HPE vNS qualification by HPE.

System interconnect fabric

The system interconnect fabric provides high-speed and low-latency connectivity between the HPE vNS CPUs and vCLIMs in an HPE vNS system. The AG for the fabric are:

1. RoCE v2
In HPE NonStop X systems, the fabric is implemented using RDMA over InfiniBand. In an HPE vNS system, the fabric is implemented using RoCE v2. This allows the system to be deployed in today's data centers where Ethernet is ubiquitous.
2. Fabric connectivity is supported on networks with Ethernet speeds between 25GbE and 100GbE. One 100GbE network interface card (NIC) can support up to four HPE vNS VMs of which at most two can be HPE vNS CPUs. One 25GbE NIC can support up to two HPE vNS VMs of which only one can be a HPE vNS CPU.
100GbE provides sufficient bandwidth for a server fabric adapter to be shared by up to four VMs (CPUs and vCLIMs) of which two can be CPU VMs. Likewise, 25GbE provides sufficient bandwidth to support fabric traffic of two HPE vNS VMs. HPE tests indicate that a VM requires at least 10 Gb link speed for fabric connectivity to avoid a performance bottleneck.
3. The servers should use one of the fabric adapters listed in [Appendix A Section 2: Fabric NICs](#).
To provide high performance and low latency connections over the fabric, the CPU and CLIM VMs require direct access to the server fabric NICs. Single Root-I/O Virtualization (SR-IOV) is an I/O technology that provides a VM with shared direct access to an I/O card. This feature is dependent on specific NIC models and firmware versions. HPE has qualified a select set of NIC models for use as an HPE vNS server fabric adapter and only NICs from this set are supported by the HPE vNS software. Apart from the use of specific NIC models, the specific OpenFabrics Enterprise Distribution (OFED) driver versions and manufacturer firmware versions are required. Please refer to HPE Virtualized NonStop deployment and configuration guide available at HPESC hpe.com/info/nonstop-ldocs for more information.
4. Ethernet switches used for the fabric connectivity require support for data center bridging (DCB) protocols, and IEEE 802.3x Global Pause.
IEEE 802.3x Global Pause is used by HPE vNS for flow control to handle fabric congestion.



- Two independent Ethernet fabrics (X and Y) are required.

Two separate fabrics provide independent paths for the HPE vNS VMs to ensure redundancy and to protect against single point failure in fabrics. This is similar to the X and Y InfiniBand fabric in an HPE NonStop X system.

Implementation notes—fabric

- One fabric NIC for each physical processor in a server is recommended for production systems. Such a configuration provides an HPE vNS CPU or vCLIM hosted on any processor in a server with direct access to the system interconnect fabric instead of indirect access through other processors in the server. Indirect fabric access through another processor in the same server can have up to 10% performance penalty versus a direct fabric access.
- To implement X and Y fabrics, use two separate NIC ports. Connect one port to the X fabric and the other port to the Y fabric.

Figure 3 illustrates the logical view of a sample HPE vNS system deployed on virtualized server hardware. In this example, one HPE vNS system (with a unique system number) consisting of four HPE vNS CPUs, four IP vCLIMs, four storage vCLIMs, and two HPE vNSCs are deployed on four physical servers based on Intel Xeon processors.

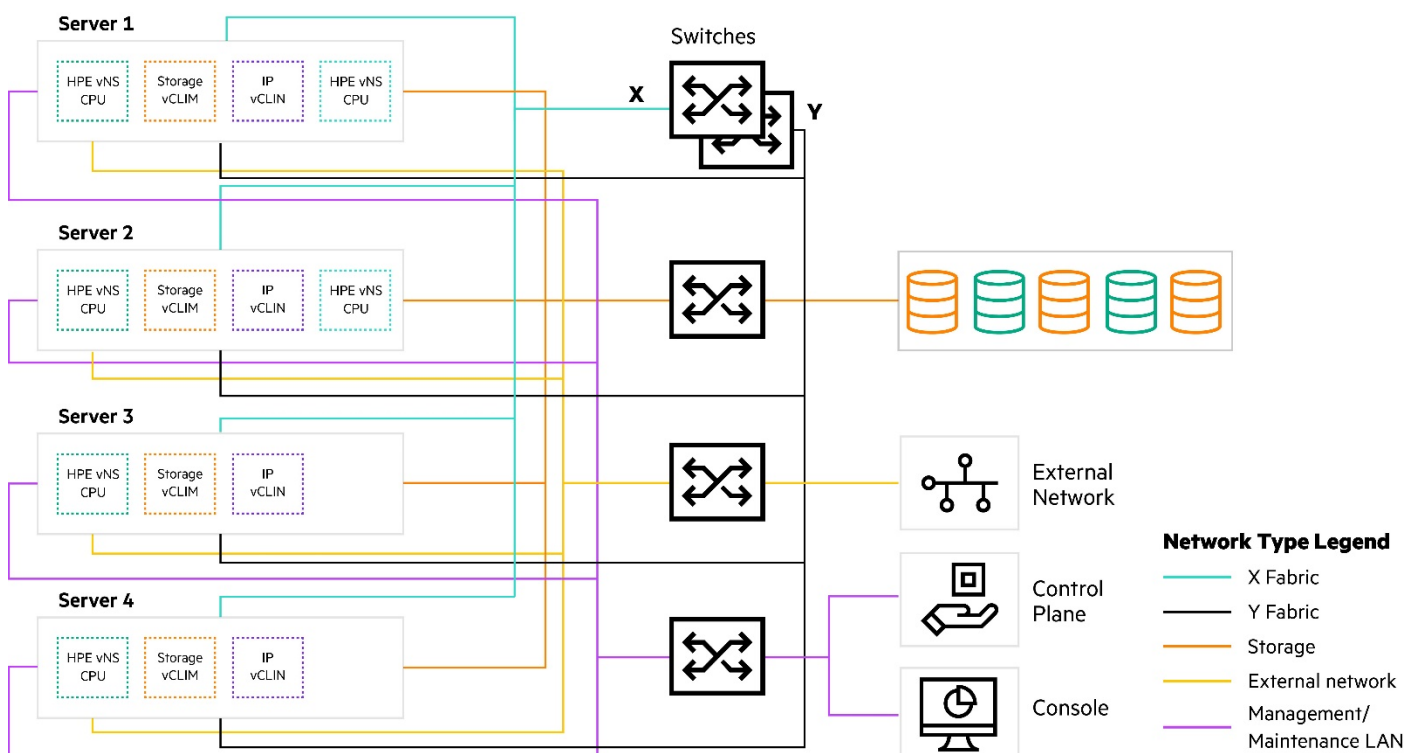


FIGURE 3. Logical view of an HPE vNS system deployment and connectivity

HPE vNS CPU

A HPE vNS CPU runs the HPE NonStop OS and the application stack. In an HPE NonStop X system, the CPU runs the HPE NonStop OS natively on a server blade. In HPE vNS systems, a CPU is a HPE vNS CPU VM, which runs the HPE NonStop OS as a VM guest OS. A HPE vNS CPU typically shares a physical host server with vCLIM VMs belonging to the same or a different HPE vNS system or with HPE vNS CPUs belonging to other HPE vNS systems or with other VM types such as HPE vNSC, vCenter, vRealize Orchestrator, or any other general purpose VMs used in a customer's environment.

The AG for HPE vNS CPU are:

- Only one CPU VM of an HPE vNS system can be deployed in a physical server. While more than one CPU VM can be deployed in a physical server, each of those CPU VMs should belong to a different HPE vNS system.

While virtualization provides opportunities to reduce the hardware footprint of an HPE NonStop system, it should be protected against single point of failure such as a failure of one physical server. Two HPE vNS CPUs from one HPE vNS system may be running processes belonging to a process pair. If both of these CPUs are hosted on the same physical server, a server failure will take down both leading to the loss of functionality provided by the process-pair. This is considered a system outage and hence should be prevented.

2. An HPE vNS CPU requires the physical processor cores and memory to be dedicated to itself and not shared with any other VMs. Due to its sensitivity to latency and performance, a HPE vNS CPU VM should have dedicated cores.
3. See the section [“Selecting processor SKUs for servers”](#) later in this document for guidance on processor selection for servers hosting HPE vNS CPUs.

IP or Telco vCLIMs

An IP or Telco CLIM offloads network I/O processing from the CPUs in an HPE NonStop system. It terminates TCP/IP sessions between external entities and an HPE NonStop system. The IP or Telco CLIM function is provided by the respective vCLIMs in an HPE vNS system. Similar to high-end HPE NonStop X systems, a high-end HPE vNS system supports between 2 and 54 IP or Telco vCLIMs. Similarly, an entry-class HPE NonStop X system and an entry-class HPE vNS system support between 2 and 4 IP/Telco vCLIMs.

Physical IP and Telco CLIMs in a HPE NonStop X system provide failover features to handle the failure of hardware ports (intra CLIM failover) and failure of the entire CLIM (CLIM-to-CLIM failover).

The AG for IP and Telco vCLIMs are:

1. An IP or Telco vCLIM requires the physical processor cores and memory to be dedicated to itself and not shared with any other VMs.

Due to its sensitivity to latency and performance, an IP or Telco vCLIM should have dedicated cores and memory and the underlying processor should have hyper-threading enabled.

2. IP and Telco vCLIM VMs require 4 or 8 dedicated cores. All IP or Telco vCLIMs from the same HPE vNS system should have the same number of dedicated cores.

The default configuration for IP and Telco vCLIMs has eight dedicated cores. If the HPE vNS system is not expected to have heavy network traffic, four cores may be dedicated to IP or Telco vCLIMs instead of eight. This flexibility eases deployment of IP or Telco vCLIMs in development and test systems as they require fewer cores from the underlying server.

3. IP and Telco vCLIMs belonging to the same failover pair should not be deployed in the same physical server. More than one IP or Telco vCLIM may be deployed on the same physical server if those IP or Telco vCLIMs belong to different failover pairs or are from different HPE vNS systems.

If two vCLIMs belonging to the same failover pair are deployed in a physical server, should that server fail, both the primary and backup vCLIM will fail simultaneously, leading to an outage.

4. IP or Telco vCLIMs can be configured to provide one of the following three types of network interfaces:

- a. VMXNET3—this is the VMware paravirtualized network interface, which allows any network I/O card supported by VMware to be used by an IP or Telco vCLIM. For the list of I/O cards supported by VMware, refer to vmware.com/resources/compatibility/pdf/vi_io_guide.pdf.

Of the three network connection types, VMXNET3 provides the lowest network throughput for a given Ethernet wire speed due to the virtualization overhead. Network interfaces in IP and Telco vCLIMs that use VMXNET3 do not support CLIM failover features.

- b. SR-IOV—in this type of interface, a physical port in a NIC is directly accessed and shared by multiple network interfaces belonging to one or more VMs. It requires a NIC with sets of virtual functions and registers to allow such an access. As with the VMXNET3 interface, network throughput in a SR-IOV connection is divided between the network interfaces sharing the NIC. The aggregate throughput of the network interfaces sharing the NIC using SR-IOV is closer to the wire speed of the NIC due to their direct access to the NIC port.

Network interfaces of IP vCLIMs using SR-IOV based NIC access support CLIM to CLIM IP-address failover but do not support intra-CLIM failover.

IP or Telco vCLIM support for SR-IOV-based network interface has dependency on specific device drivers and hence is limited to specific NIC models. See [Section 3: NICs supporting SR-IOV](#) in [Appendix A](#) for more information.

- c. PCI passthrough—this provides an IP or Telco vCLIM with exclusive direct access to a physical port in a NIC which it uses to provide one network interface. Such a network interface offers the highest throughput compared to VMXNET3 or SR-IOV based network interface types because the entire NIC port is dedicated to that interface.

PCI passthrough supports both intra-CLIM and CLIM-to-CLIM failover. Of the three network connection types, PCI passthrough provides the closest match to the feature-set available in physical CLIMs.

IP or Telco vCLIM support for PCI passthrough network interface is limited to specific NIC models. See [Section 4: NICs supporting PCI passthrough for network interface in IP and Telco vCLIMs](#) in [Appendix A](#) for more information.

IP and Telco vCLIMs—Implementation notes

1. An IP or Telco vCLIM provides up to four network interfaces with PCI passthrough using two 2-port Ethernet NICs.
2. For IP or Telco vCLIM eth0 connection to management LAN, one of the embedded LOM ports in the server can be used with VMXNET3.



3. See the section [Selecting processor SKUs for servers](#) later in this document for guidance on processor selection for servers hosting vCLIM VMs.

Storage vCLIMs

A storage CLIM offloads low-level storage I/O processing from the CPUs in an HPE NonStop system. The storage CLIM function is provided by the VMs (vCLIMs) in an HPE vNS system. Similar to a high-end HPE NonStop X system, a high-end HPE vNS system can have between 2 and 54 storage vCLIMs. Likewise, as in an entry-class HPE NonStop X system, an entry-class HPE vNS system can have between 2 and 4 storage vCLIMs. The number of storage vCLIMs in an HPE vNS system must be an even number.

Storage drives can be connected to physical servers hosting storage vCLIMs as internal drives in the server or as external drives in one or more network storage systems.

In a virtualized environment, the hypervisor intermediates between a VM and physical hardware resources. For VM access to physical storage drives, VMware provides the means to create virtual disks from the physical storage drives and presents the virtual disks to the VM.

VMware vSphere provides several virtual SCSI storage controllers for VMs to access the virtual disks. For the HPE vNS storage vCLIM, the VMware Paravirtual SCSI (PVSCSI) controller provides the best storage performance and is the recommended controller for the vCLIM.

The AG for storage vCLIMs are:

1. On the HPE Virtualized NonStop (vNS) systems, the storage vCLIMs should be assigned dedicated processor cores and memory.
For similar reasons as stated earlier for CPU VMs and IP vCLIMs, storage vCLIMs also require dedicated processor cores and memory that are not shared with other VMs by the hypervisor.
2. A storage vCLIM can be provisioned with either 4 or 8 processor cores. All storage vCLIMs should have the same number of cores assigned.
This flexibility helps with making efficient use of the available hardware resources. Use of eight processor cores is the default configuration and is required if volume level encryption (VLE) is implemented. Use of four processor cores supports systems with lower storage requirements.
3. Storage vCLIMs belonging to the same failover pair should not be deployed on the same physical server. As a corollary, each of the storage vCLIMs deployed on a physical server must belong to separate failover pairs.
Similar to the explanation for the IP vCLIMs of a system, if two storage vCLIMs belonging to the same failover pair are hosted on a physical server, an outage of that server will cause an outage of the system.
4. A pair of storage vCLIMs can connect between one and 50 mirrored storage devices (up to a total of 100 LUNs).
5. If VLE is used, storage vCLIMs require connectivity to an Enterprise Secure Key Manager (ESKM). This is an IP connection, which can be provisioned over a VMXNET3 interface of the storage vCLIM.
6. Storage CLIMs require storage I/O cards supported by VMware as specified in vmware.com/resources/compatibility/pdf/vi_io_guide.pdf.
HPE vNS uses VMware PVSCSI controller to connect to storage volumes. Hence any storage I/O card supported by VMware for block level storage access will work for HPE vNS.
7. For connecting to external SAN storage a physical server is recommended to have one storage NIC for each storage vCLIM deployed on it.
8. HPE vNS requires block storage devices that are supported by VMware. For external storage options, refer to vmware.com/resources/compatibility/pdf/vi_san_guide.pdf.
9. If you are using internal drives, multiple paths to storage volumes are not supported. HPE vNS systems using external storage may be configured with multiple paths to storage volumes in either 2 CLIM-2 Disk (2C2D) or 4 CLIM-2 Disk (4C2D) configurations. See the section [Multipath access between CLIM and storage volumes](#) in [Appendix B: Storage considerations](#) for more information.
10. For external storage connectivity, you may use either iSCSI (Ethernet) or Fibre Channel (FC) networks.
Since HPE vNS uses VMware PVSCSI controller, the deployment of storage volumes could use any storage networking technology supported by VMware. Historically, the use of FC was popularized as a faster alternative to Ethernet networks for storage access. However, the advent of faster Ethernet technologies coupled with the ubiquitous nature of Ethernet networks in enterprise data centers has led to the increasing adoption of Ethernet networks for storage when compared to FC. You may implement either of these storage networking options for connecting your servers to external storage.
11. For backup requirements, HPE vNS only supports virtual tapes and requires HPE NonStop BackBox.
HPE vNS does not support physical tapes. For backup needs, HPE vNS supports virtual tapes and requires either a virtual BackBox or a physical BackBox VTC. You can connect multiple HPE vNS or converged HPE NonStop systems to a virtual BackBox or to a physical BackBox VTC.



Storage vCLIMs—Implementation notes

1. If you are using internal storage drives for an HPE vNS system:
 - a. Use a server with higher number of drive bays in order to accommodate more disks and hence implement a larger storage configuration if required.
2. A 10GbE converged network adapter (CNA) with offloads FCoE/iSCSI protocol processing from the host server processor.
3. For better storage I/O performance of HPE vNS disk volumes, consider using SSDs.
4. The following table lists the typical storage requirements for the HPE vNS system.

TABLE 1. Storage requirements for a typical HPE vNS system

Volume	Size	Remarks
\$SYSTEM	100 to 600 GB	In increments of 1 GB
\$SWAP	100 to 2000 GB	Use the formula $1/2 \times \text{memory per CPU} \times \text{number of CPUs}$
\$AUDIT	100 to 600 GB	
\$DSMSCM	100 to 600 GB	
\$OSS	100 to 600 GB	
\$DATA volumes	1 to 600 GB	Based on user requirements
Storage vCLIM OS for first pair of storage vCLIMs	300 GB	
Storage vCLIM OS for additional storage vCLIMs	100 GB	
IP and Telco vCLIM OS	100 to 300 GB	Use larger size to support longer TCP/IP monitor dumps
HPE vNSC	250 GB	
HPE Virtual BackBox	300 GB	

These storage requirements are in addition to the storage requirement for VMware products such as vSphere, vRealize Orchestrator, and vCenter which will be applicable on servers they're deployed on. For information on storage requirements for VMware products go to: docs.vmware.com.

5. For storage that mandates RAID configuration to protect against drive failures, consider RAID 5 unless your storage vendor and/or storage architect have a different recommendation. RAID 5 protects against a single drive failure, provides higher write performance than RAID 6 and uses physical storage capacity more efficiently than RAID 1.
6. Network storage systems can have storage overhead for high availability beyond RAID, which reduces the usable storage capacity. The network storage system storage sizing tool must be used to determine usable storage capacity (such as NinjaSTARS for HPE 3PAR and HPE Primera storage systems).
7. A network storage system can distribute the logical storage volumes across the entire set of drives in the storage system by default. The primary and mirror HPE vNS volumes can thus be provisioned on the same set of drives. To protect against two drive failures, the primary and mirror HPE vNS volumes should be provisioned on mutually exclusive set of drives (preferably inside separate enclosures) and supported by separate controllers.

Refer to the storage system vendor documentation to understand the implications of points 5 to 7.

See the section "[Selecting processor SKUs for servers](#)" later in this document for guidance on processor selection for storage vCLIMs.

HPE vNSC

An HPE Virtualized NonStop System Console (HPE vNSC) should be hosted on a Windows VM. The HPE vNSC is a set of applications and not an appliance. Customers need to install it on a VM with separately licensed Windows Server 2012 or Windows Server 2016 or Windows Server 2019.

The AG for HPE vNSC are:

1. An HPE vNS system should be managed by an HPE vNSC or a pair of HPE vNSCs.

An HPE vNSC is required to perform installation, configuration, and management tasks for an HPE vNS system. In an HPE NonStop X system, NSC provides the HPE NonStop Halted State Service (HSS) software image to network boot a CPU before the HPE NonStop OS is loaded on it. Two NSCs provide high availability for the HSS network boot server function. In an HPE vNS system, the HSS and the HPE NonStop OS images are hosted in an independent management plane and critical functions such as HPE vNS CPU VM reloads do not require access to the HPE vNSC. Hence one HPE vNSC instance is sufficient to manage a HPE vNS system.
2. One instance of HPE vNSC can manage up to eight HPE vNS systems.



Sharing server hardware between multiple HPE vNS systems

One of the main advantages offered by virtualization is workload consolidation. Such an advantage may also be leveraged to consolidate the HPE vNS systems over a smaller hardware footprint. The deployment guidelines described here provide the guiding principles for achieving such a consolidation. The sections above describe the rules governing the deployment of HPE vNS VMs over virtualized hardware. The deployment guidelines for sharing hardware between multiple HPE vNS systems, described in this section, shall be consistent with those rules:

1. A physical server can host more than one HPE vNS CPU where each such HPE vNS CPU belongs to different HPE vNS systems. This is to ensure that, should a server have an outage, no more than one HPE vNS CPU belonging to the same HPE vNS system is impacted.
2. A physical server can host more than one IP vCLIM where each such IP vCLIM:
 - a. Belongs to different HPE vNS systems
 - Or
 - b. Belongs to the same HPE vNS system but belongs to a different failover pair

This is to ensure that, should a server have an outage, the network path accessed through the impacted IP vCLIM will failover to its backup IP vCLIM running on a different physical server.
3. Each IP vCLIM requires two 2-port Ethernet NICs. If a physical server is hosting more than one IP vCLIM, the total number of Ethernet NICs in the server should equal twice the number of IP vCLIMs deployed on it. However, as described in the section on IP vCLIMs, if an HPE vNS system does not require four passthrough Ethernet interfaces or uses VMXNET3- or SR-IOV-based interfaces, the required number of Ethernet NICs can be lower.
4. A physical server can host more than one storage vCLIM where each such storage vCLIM either:
 - a. Belongs to different HPE vNS systems
 - Or
 - b. Belongs to the same HPE vNS system but belongs to a different failover pair

This is to ensure that, should a server have an outage, the storage path accessed through the impacted storage vCLIM will failover to its backup storage vCLIM running on a different server.
5. If a server is hosting multiple storage vCLIMs which in turn access SAN storage, it is recommended to use separate storage NICs for each of these vCLIMs in the server in order to avoid starving of I/O bandwidth (the noisy neighbor effect).
6. Running multiple IP and/or storage vCLIMs belonging to the same HPE vNS system on a physical server will broaden the fault-zone of the system. Consider your system availability requirements carefully while designing such a system.

Hardware implementation guide

Selecting processor SKUs for servers

Intel offers a large number of x86 processor SKUs. When selecting the processor SKU for the servers hosting HPE vNS systems, adhere to the following recommendations:

1. The choice of processor SKU is determined by the VM that is most sensitive to the processor performance (i.e., the HPE vNS CPU) to be deployed on the host server.
2. Servers hosting HPE vNS CPUs of an HPE vNS system should have the same processor SKU. However, during a hardware upgrade, you may do a rolling upgrade of the processors while the system is online. During this period of transition, you will have a system with a mix of old and new processor SKUs.
3. If a server is hosting only vCLIMs, such a server may use processor SKUs different from the ones used by the servers hosting HPE vNS CPUs. Such servers may use lower (and less expensive) processor SKUs to save on cost.
4. The processor core frequency is the primary factor that influences HPE vNS CPU performance. Processors with higher core frequency and faster memory bus provide higher HPE vNS CPU performance. For best system performance, typically required for production systems, use processors with higher core frequency (≥ 3.2 GHz) and faster memory bus in servers which host HPE vNS CPUs.

High core frequency processors have higher cost than lower core frequency processors. For servers hosting HPE vNS CPUs of systems not having demanding performance requirements (such as development systems) or for servers hosting only vCLIMs (storage or IP), you may use processors with lower core frequency. Even for such servers, it's recommended to use processors with a core frequency ≥ 2 GHz.

Intel Xeon Scalable processor family uses names of precious metals to indicate processor performance. Accordingly, the names **Platinum** and **Gold** are used along with the processor model numbers to identify faster processors and these are good candidates for servers hosting HPE vNS CPUs for production systems. The next processor tier called **Silver** is a good candidate for servers hosting HPE vNS CPUs for development systems or for servers hosting only vCLIMs (of a production or a development system). **Bronze** processors are not recommended for the servers hosting any of the HPE NonStop VMs.



As an example, you may refer to the [HPE ProLiant DL380 QuickSpecs](#) to see the list of Intel Xeon SKUs orderable with HPE ProLiant DL380 servers. Based on your target configuration, you may select the appropriate processor SKU by referring to this document.

5. In servers with two processors, installing all of the memory required by the HPE vNS CPU VM in the DIMM sockets of one of the two processors provides higher performance than splitting the same memory between the two processors as in a usual balanced memory configuration. The amount of memory installed in the DIMM sockets of the second processor could be reduced to lower the server hardware cost. This unbalanced memory configuration provides higher performance by allowing the HPE vNS CPU VM to access all of its memory without having to access the second processor. For example, in a server with two processors that hosts a HPE vNS CPU VM with 256 GB memory and a vCLIM with 16 GB memory, 256 GB + overhead could be installed in DIMM sockets of one processor while 16 GB + overhead could be installed in DIMM sockets of the second processor.
6. The total number of cores required in a server should be equal to or greater than the sum of the cores required by all the VMs hosted in the server and the cores required by ESXi. See [Appendix C: System requirements](#) for information on the number of cores required by various constituents of a server hosting a HPE vNS system.
7. Use the following guidelines to arrive at the number of cores required by ESXi:
 - a. Compute 21% of the total number of cores required by the VMs hosted in the server that require dedicated cores and round up to next higher whole number.
 - b. For example, since HPE vNS CPUs and vCLIMs require dedicated cores, if those VMs consume 20 cores in a server, take 21% of 20, which is 4.2. Round up to 5 and add to the 20 cores dedicated to HPE vNS CPU and vCLIM VMs for a total minimum of 25 cores that would be required by ESXi in order to deploy the HPE vNS CPU and vCLIM VMs.
8. In general, faster processor SKUs and higher core count may have higher power rating and non-linear price increase for the processor SKU.
9. Add free cores to the number of cores required to support future expansion or if you plan to use NSDC to dynamically scale HPE vNS CPU cores.
10. For optimum performance, all the cores used by a HPE vNS CPU or a vCLIM VM should be deployed in the same processor instead of being split across two or more processors. Splitting the cores of a HPE vNS CPU or a vCLIM VM across two or more processors will have a performance penalty associated with the data transfer between the processors.

Server memory

As mentioned in earlier sections, HPE vNS VMs require dedicated memory which cannot be shared with other VMs. For such VMs, ESXi reserves an additional 20% of memory. To arrive at the total memory required in a physical server, mark up the memory required by HPE vNS VMs by an additional 20% and add up the memory required by other VMs deployed in the server. Server vendors provide recommendation on allowed DIMM configuration in the server models. HPE servers provide documentation on recommended DIMM configuration in the QuickSpecs. It is recommended to comply with these guidelines.

Laying out HPE vNS VMs in servers

The following is a sample layout of the VMs in an HPE vNS system with two HPE vNS CPUs (2-cores each), 2 IP vCLIMs, and 4 storage vCLIMs. In this example, the vCLIMs have been assigned with 8 cores each. Servers 1 and 2 could be loaded with a higher performance Intel® Xeon® Gold processor SKU (for example, Intel Xeon Gold 6246R processor) since it hosts a HPE vNS CPU. Servers 3 and 4 could be loaded with a lower cost Intel® Xeon® Silver processor SKU (for example, Intel Xeon Silver 4214 processor) since it does not host a HPE vNS CPU.

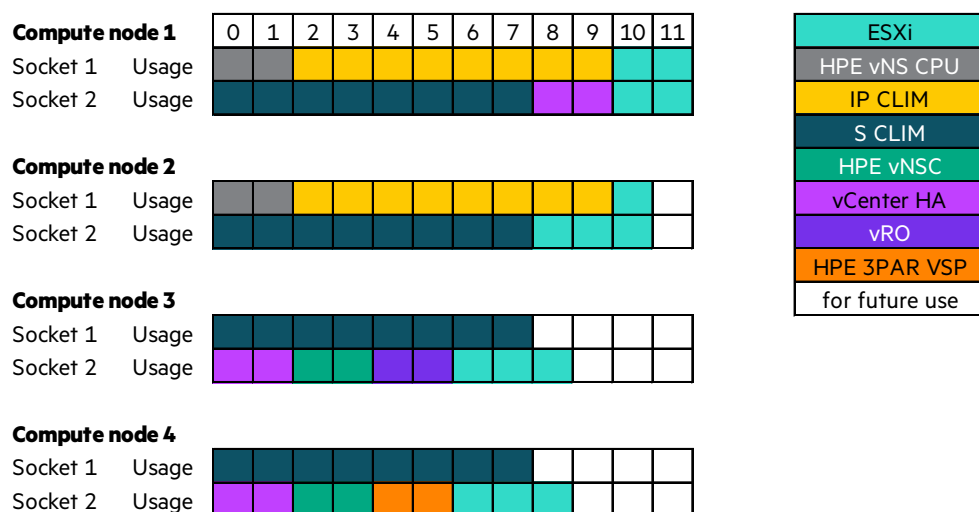


FIGURE 4. Sample core layout of VMs in physical processors of a server

NOTE

The layout is only a logical representation. The actual core assignment is determined by ESXi at the time of VM deployment. The HPE vNS deployment tools have limited influence over it.

Connectivity

An HPE NonStop system contains several elements, spanning technologies such as servers, storage, and networking, which work in tandem to present as a “single system” to the outside world. Proper connectivity between these elements becomes highly critical for the correct operation of the system. Following sections explain these connection setups and best practices.

Fabric connection between HPE vNS CPUs and vCLIMs

HPE vNS CPUs and vCLIMs connect over the high-speed RoCE v2 fabric. The architecture guide for HPE vNS fabric is explained in an earlier section. Only specific NIC models are supported for the fabric connection as explained in [Section 2: Fabric NICs](#). These NICs are used in the servers hosting HPE vNS CPU VMs and vCLIM VMs. The fabric switches should support the fabric speed requirement (25GbE to 100GbE) and support data center bridging (DCB) and 802.3x Pause frames for flow control.

Two separate switches are required to support X and Y fabric for redundancy. In other words, the X and Y fabric cannot share the same physical switch. On servers having a single 2-port fabric NIC, connect one of the ports to X and the other to Y. On servers having two 2-port fabric NICs, connect both ports of one of the NICs to X and those of the other NIC to Y. This provides better availability characteristic. The fabric switches should have adequate number of ports to support all the fabric NICs in the system.

[Figure 15](#) illustrates the fabric connections for the sample BoM which may be used as a reference.

External Storage network

For external storage connectivity, you may use either iSCSI (Ethernet) or FC networks. You may use a dedicated storage network for your HPE vNS system or connect to your enterprise SAN to meet the storage needs of your HPE vNS system. The latter offers cost benefits through storage consolidation. These are standard storage connectivity options and no special considerations are necessary for HPE vNS.

For redundancy, the paths to storage devices hosting primary and mirror volumes should be separated to ensure availability in case of single point failures. If you're using external storage arrays, it's recommended to:

- a) Use separate storage arrays for hosting primary and mirror volumes
- b) Have separate, redundant, connection paths between the storage vCLIMs and storage devices. See [Figure 14](#) for a possible connection diagram.

If you are using dedicated external storage and use iSCSI connection, you may share the fabric switch for storage networking as well. This is achievable if the switch model chosen supports 10GbE (for iSCSI) and the fabric speed (25GbE to 100GbE). This saves the cost of dedicated 10GbE switches for the storage network. However, in such a configuration, it is highly recommended to isolate the fabric and iSCSI network using VLANs for security purposes.

Maintenance LAN

The HPE vNS maintenance LAN is used by HPE vNSC to connect to the HPE vNS CPUs, and the vCLIMs, to securely perform commissioning and configuration tasks. This LAN is also used to perform CLIM administration tasks from HPE NonStop host using CLIMCMD from TACL. A dedicated 1GbE switch is used for this LAN.

The OSM tools in HPE vNSC only manage the software in HPE vNS systems.

Redundancy is not a requirement for management or maintenance LAN.

Management LAN

If the data center has a management network for managing all hardware in the data center, the hardware for hosting the HPE vNS system may be integrated into the same network. For managing the hardware, you must use the hardware management application from your vendor (for example, HPE OneView).

An alternate approach would be to connect the hardware management ports to the same switch as the maintenance LAN switch. The management appliances connect to this hardware (e.g., HPE iLO ports of DL380 servers) through this switch. If such an approach is adopted, isolate the two maintenance and management networks using a VLAN configuration for better security.

Additional Security considerations

As is evident from the earlier sections, a number of datapaths exist between VMs of an HPE vNS system and external entities. The fabric traffic, in particular, needs to be protected against eavesdropping since it contains information exchanged between processes running on



HPE vNS CPU VMs and the vCLIMs. Therefore, it is highly recommended to use a dedicated pair of fabric switches for one HPE vNS system. These switches should not be shared with any other VM or node within the network and should also be protected against unauthorized logical or physical access.

The other network paths associated with the HPE vNS VMs are:

1. Storage (iSCSI or FC) network—if external storage is used
2. Maintenance LAN
3. External network (for IP and Telco vCLIMs)

Separate VLANs should be used to isolate these networks from one another. This not only logically isolates the traffic between the VMs, it can also be used to implement QoS for throughput and latency-sensitive traffic such as storage access.

Rack

HPE vNS does not require a dedicated rack. Depending on the target environment considerations (security and space), you may host the HPE vNS hardware in a dedicated rack or share it with other hardware. 2U high servers are recommended as they provide more PCIe slots for I/O cards.

VMware requirements

HPE vNS requires following three VMware products:

1. VMware vSphere 6.7 and above
2. VMware vCenter 6.7 and above
3. VMware vRealize Orchestrator 7.3 and above

ESXi is the virtualization software that should be run on each physical server. HPE vNS requires vSphere Enterprise Plus Edition, which supports SR-IOV for VM I/Os. For deploying a system, vRealize Orchestrator is needed, which is an appliance available in the vSphere Enterprise Plus bundle.

VMware vCenter is required for managing and administering the VMware environment. HPE vNS does not require a dedicated vCenter. You may integrate the hardware running an HPE vNS system into an existing vCenter managed environment. For production use, vCenter High Availability (HA) is required, which involves running three vCenter instances (active, standby, and witness) on separate physical servers. One license of vCenter Standard allows you to run the vCenter HA configuration consisting of the three instances.

While arriving at the HPE vNS hardware configuration, consider the resource requirements for the three VMware software. They are documented at docs.vmware.com/.

HPE Virtualized NonStop and support for hardware product-lines

A common question that gets asked to HPE NonStop Product Management is—Does HPE Virtualized NonStop support “XYZ” hardware product? The answer to this question lies in understanding the relationship that HPE vNS has with hardware.

A HPE Virtualized NonStop System is hardware agnostic to a large part. It depends on the underlying virtualization layer (vSphere) to provide it with the required computing resources. However, it should be deployed in accordance with the rules described in the earlier sections of this document. Having said that, HPE vNS does have three hardware specific requirements which are described in earlier sections of this document, summarized as below:

(R1) The servers hosting HPE vNS VMs should use processors from specific Intel Xeon x86 processor families.

(R2) Each server hosting HPE vNS VMs should have one or two Mellanox ConnectX host adapters for the system interconnect fabric. (see Appendix A, section 2).

(R3) Each server hosting HPE vNS IP vCLIMs should have one or two Ethernet NICs based on either Intel X550 or X710 network processors or Marvell (QLogic) 57810S processor (this is in addition to R2) see Appendix A, Section 3.

Apart from the above, all hardware should be compatible with the VMware release being used in the environment. VMware publishes the hardware compatibility information at vmware.com/resources/compatibility which should be referred to for this purpose.

Hence the more appropriate question to ask is: Can a HPE vNS architecture guide compliant implementation be achieved using “XYZ” hardware products? Referring to this guide, your hardware specialist should easily be able to answer this question.



Implementing HPE vNS Architecture Guide compliant hardware using HPE Synergy

Synergy offers two types of Gen10 compute modules: 1) HPE Synergy 480 Gen10 (2-socket) 2) HPE Synergy 660 Gen10 (2–4 socket). Both these compute modules offer Intel Xeon Scalable processors of Gold and Silver class. Hence this meets the first requirement (R1) of HPE vNS reference architecture.

HPE Synergy compute modules support HPE Synergy 6410C 25/50Gb Ethernet Adapter which is based on Mellanox ConnectX-4 family of NICs. While it is not the exact SKU specified in Appendix A, Section 2, it's functionally equivalent and hence meets the requirement (R2).

HPE Synergy compute modules support HPE Synergy 2820C 10GbE converged network adapter which uses the Cavium (QLogic) 57840S network processor and the same firmware version as the one supported by HPE 530SFP+/530T NIC. Hence this meets the requirement (R3) of HPE vNS reference architecture.

In summary, it's possible to build a HPE Synergy based infrastructure that complies with the HPE vNS hardware architectural guidelines. Such an infrastructure must use HPE Synergy 6410C 25/50Gb Ethernet Adapter for HPE vNS system fabric and HPE Synergy 2820C 10GbE converged network adapter for Ethernet interfaces in IP vCLIMs.

HPE Virtualized NonStop and composability

Hardware composability is one of the defining features HPE Synergy. It accomplishes this using two management appliances—HPE Synergy Composer and HPE Synergy Image Streamer. An obvious question is—does HPE Virtualized NonStop support (more appropriate phrase would be “leverage”) these capabilities. Before answering the question, it's important to understand the applicability of these concepts to a typical HPE NonStop workload.

HPE NonStop has abilities to scale out and scale up but does not support dynamic assignment of hardware resources to a large extent. In the context of a HPE vNS system:

- The cores assigned to HPE vNS VMs are dedicated to the VMs. These cannot be shared with other VMs.
- Number of “active cores” should be identical among a class of VMs. All HPE vNS CPUs should have the same core count, all storage vCLIMs should have the same core count and all IP vCLIMs should have the same core count. Any change to the active core count of vCLIMs requires shutting down the entire system and redeploying it (a process known as “cold-load” in the HPE NonStop parlance).
- Physical resources such as cores, processors, and NICs assigned to HPE vNS VMs cannot be altered or reassigned without redeploying VMs.
- The storage devices (or the LUNs) can be added or reassigned online to a HPE vNS system but require a specific workflow to be followed.
- New VMs (in pairs) can be added to a HPE vNS system online. For example, new CPUs, new IP vCLIMs and new storage vCLIMs can be added to a HPE vNS system online (subject to license rights of course!) and without requiring any cold-load.
- Updating resource profiles (e.g., IP addresses, firmware) of a HPE vNS system cannot be done online.

As you can see from the above, a HPE NonStop system offers the ability to “recompose” the system but within strict boundaries. Customers use a HPE NonStop system to serve mission-critical computing needs. These are workloads that need “continuous computing” while the user traffic itself may vary within bounds. The primacy in these deployments is given to the availability of HPE vNS systems and typical customer deployments enforce very strict rules on modification of system parameters. In other words, a HPE vNS workload is fairly static in its resource requirement. Hence the “composability” paradigm has limited applicability for a mission-critical workload that typically runs on HPE NonStop systems.

HPE Virtualized NonStop needs to be deployed using the HPE Virtualized NonStop deployment tool. It is a vCenter plug-in and views computing resources in a generic way simply because it's agnostic to the profile of the hardware running underneath (e.g., vendor, model, form-factor). It has no knowledge of special characteristics of HPE Synergy such as its ability to dynamically compose the hardware resources or its ability to align a set of software images, ready to be deployed on a newly composed hardware. The HPE Virtualized NonStop deployment tool does not interface with either the HPE Synergy Composer or the HPE Synergy Image Streamer.

A HPE vNS system needs to be deployed on HPE Synergy in a way that is similar to how it is deployed on rackmount servers i.e., using the vCenter plug-in. This is typically done by HPE Pointnext Services specialists at the time of implementing the HPE vNS installation and startup service at the customer site. The HPE Synergy hardware resources required to deploy a HPE vNS system should be provisioned by the user prior to the deployment, likely using the HPE Synergy Composer.

Management of the hardware and the virtualization layer is outside of the scope for HPE vNS. This is as true for deployment on HPE Synergy as it is for deployment on any other hardware environment. OSM and other HPE NonStop manageability products can be used to manage only the HPE NonStop software.



APPENDIX A: HPE VNS HARDWARE SUPPORT MATRIX

Section 1: Server models

Processor model	Tested with
2nd Gen Intel Xeon Scalable processors (Gold and Silver)	HPE ProLiant DL380 Gen10
Intel Xeon Scalable processors (Gold and Silver)	HPE ProLiant DL380 Gen10
Intel Xeon Broadwell (E5-nnnn v4 and E7-nnnn v4)	HPE ProLiant DL380 Gen9
Intel Xeon Haswell (E5-nnnn v3 and E7-nnnn v4)	HPE ProLiant DL380 Gen9

Section 2: Fabric NICs

NIC models	Vendor SKU	Tested with
Mellanox ConnectX-6 VPI or Mellanox ConnectX-6 EN	MCX65nnnnA-AAAA or MCX61nnnnA-AAAA (where 'n' is a numeral and 'A' is an alphabet)	HPE IB HDR100/EN 100G 2p 940QSFP56 Adptr
Mellanox ConnectX-4 2p—VPI	MCX456A-ECAT	HPE 840QSFP28 IB EDR/Ethernet 100 Gb 2-port
Mellanox ConnectX-4Lx Adapter—EN	MCX4121A-ACUT	HPE 640SFP28 25GbE 2p ConnectX-4Lx Adapter—EN

Section 3: NICs supporting SR-IOV for network interface in IP and Telco vCLIMs

NIC processor	Servers	Tested with NICs	Tested on servers
Intel 82599	2nd Gen Intel Xeon Scalable	HPE Ethernet 10 Gb 2-port 560SFP+	HPE ProLiant DL 380 Gen9
	Intel Xeon Scalable	HPE Ethernet 10 Gb 2-port 530T	
Cavium (QLogic) 57810S	Intel Xeon Broadwell	HPE Ethernet 10 Gb 2-port 530SFP+	
	Intel Xeon Haswell		HPE ProLiant DL 380 Gen10
Intel X710		HPE Ethernet 10 Gb 2-port 562SFP+	

Section 4: NICs supporting PCI passthrough for network interface in IP and Telco vCLIMs

NIC processor	Servers	Tested with NIC	Tested on servers
Intel 82599	Intel Xeon Scalable Intel Xeon Broadwell Intel Xeon Haswell	HPE Ethernet 10 Gb 2-port 560SFP+	HPE ProLiant DL 380 Gen9 HPE ProLiant DL 380 Gen10
		HPE Ethernet 10 Gb 2-port 560T	
		HPE Ethernet 10 Gb 2-port 530T	
Cavium (QLogic) 57810S		HPE Ethernet 10 Gb 2-port 530SFP+	
		HPE Ethernet 10 Gb 2-port 562SFP+	
Intel X710		HPE Ethernet 10 Gb 2-port 562SFP+	
Intel X550		HPE Ethernet 10 Gb 2-port 562T	

Section 5: Storage products usable with HPE vNS

Refer VMware SAN guide [vmware.com/resources/compatibility/pdf/vi_san_guide.pdf](https://www.vmware.com/resources/compatibility/pdf/vi_san_guide.pdf). HPE vNS supports only block storage devices.

Section 6: Ethernet switches

The table below lists the switches used by HPE Quality Assurance (QA) teams for HPE vNS validation.

NOTE

This table is given as an example. You may use any other switch that supports the fabric speed (25GbE to 100GbE) and DCB (802.3x).

Switch model	SKU
HPE FlexFabric 5930 Switch Series	See product QuickSpecs
HPE FlexFabric 5950 Switch Series	See product QuickSpecs
HPE StoreFabric M-Series SN2100M Ethernet Switch	See product QuickSpecs



APPENDIX B: STORAGE CONSIDERATIONS

Multipath access between CLIM and storage volumes

HPE NonStop high-availability architecture for storage provides several different ways to protect against disruptions to storage I/O upon hardware failures. This section provides a high-level overview of the HPE NonStop storage availability architecture and how it applies to HPE vNS. This section does not intend to be an extensive description of this topic. The reader is referred to the section “Configuration for storage CLIM and SAS disk enclosures” in the “HPE NonStop X NS7 Planning Guide” available at HPESC (hpe.com/info/nonstop-docs) for more information.

The HPE NonStop storage volumes are created and provisioned as primary and mirrored pairs. These volumes are connected to storage CLIMs through which HPE NonStop processes perform I/O operations. At any given point, a CLIM may have access path to either the primary volume or to its mirror or to both. This leads to three different storage path configurations:

1. Two CLIMs, configured as primary and backup (also called failover pair), are connected to a primary volume and its mirror respectively. The CLIM connected to the primary volume does not have access to its mirror and vice versa. All write I/Os are directed to the primary and mirror volumes through the two CLIMs.

Upon a CLIM failure, only the surviving CLIM continues with the write and hence only the primary or backup volume (and not both) gets updated. Once the failed CLIM comes back up, the disk that it's connected to undergoes a revive operation wherein it's synced to its peer. The revive is initiated through the HPE NonStop I/O stack and hence consumes CPU cycles.

In an HPE vNS system that uses internal drives but without any storage virtualization (VMware vSAN™), this is the only supported storage. Internal drives are visible only to the server they are attached to. Since the primary and backup storage vCLIMs should be hosted on two separate physical servers, a disk is only accessible to the vCLIM hosted on that physical server. This configuration is depicted in Figure 5.

2. Two CLIMs, configured as primary and backup, with each CLIM connecting to both the primary and the mirrored storage volumes. This is called 2c-2d configuration. This configuration provides redundant I/O paths to redundant storage volumes to protect against single point of failure of one of the paths and/or one of the storage volumes. In other words, should one of the CLIMs fail, the surviving CLIM continues to write to both primary and mirror volumes to keep them in sync. During this period, the surviving CLIM will experience double the I/O load. After the failed CLIM comes back up, there is no need to revive disks since the primary and the backup volumes are already in sync. This configuration is depicted in Figure 6.
3. Four CLIMs where one pair of CLIMs connect to primary storage volumes and another pair of CLIMs connect to mirror storage volumes. This is called 4c-2d configuration. This configuration provides redundant I/O paths to redundant storage volumes to protect against single point of failure of one of the paths and/or one of the storage volumes using four CLIMs. The principles are similar as that of 2c-2d configuration described above except that separate CLIMs serve the I/O paths to primary and mirror volumes in both the normal operation and failover operation. This configuration requires twice the number of CLIMs. During a failure scenario, the backup CLIM experiences normal I/O load (and not double the I/O load, as in the case of a 2c-2d configuration). This configuration is depicted in Figure 7.



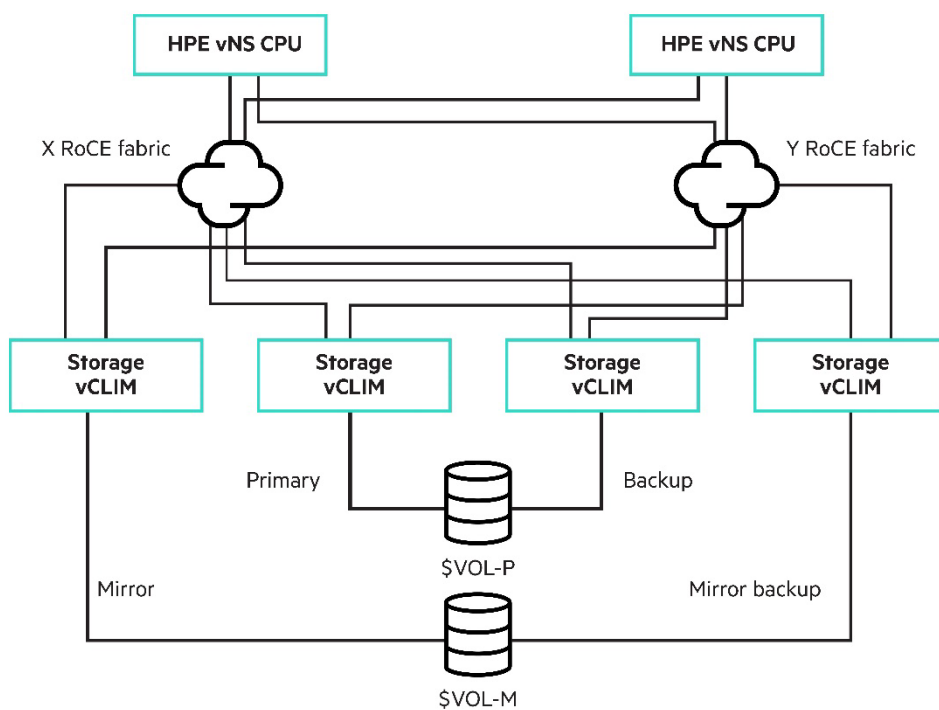


FIGURE 5. Storage redundancy—One path from CLIM to storage



2c-2d configuration with storage arrays

FIGURE 6. Storage redundancy 2c-2d configuration

4c-2d configuration

FIGURE 7. Storage redundancy 4c-2d configuration

APPENDIX C: SYSTEM REQUIREMENTS

A HPE Virtualized NonStop system consists of a number of VMs of different types, each of which have specific system requirements. Hence the requirement for the hardware is an aggregate of the system requirements for all the VMs. The table below lists the system requirements for the VMs that constitute an HPE vNS system.

TABLE 2. System requirements for an HPE vNS system

VM type	Cores	Memory	Remarks
HPE vNS CPU—Entry class	1 or 2	32 GB or 64 GB	In 1 GB increments
HPE vNS CPU—High End	2, 4, or 6	64 GB to 256 GB	In 1 GB increments
IP or Telco vCLIM	4 or 8	16 GB	In 1 GB increments
Storage vCLIM	4 or 8	8 GB	Use 8 cores if VLE is in use Memory in 1 GB increments
HPE vNSC	2	8 GB	Cores need not be dedicated
Virtual BackBox	2 or more	8 GB	Cores need not be dedicated
vCenter ¹	2	12 GB	For up to 10 servers and up to 100 VMs cores are not dedicated
	4	19 GB	For up to 100 servers and up to 1000 VMs cores are not dedicated
vRealize Orchestrator ²	4	12 GB	Cores are not dedicated
ESXi (hypervisor) ³	See Remarks	8 GB	Take 21% of the total core count required by HPE vNS VMs (don't include cores of other VMs such as vCenter and vRO) and round up to next higher integer, for example, if the VMs requires 20 cores and 21% of 20 cores is 4.2; rounding up to 5 cores for the number of cores held in reserve by ESXi

For storage requirements, please refer to Table 1.

¹ Hardware Requirements for the VMware vCenter Server® Appliance™ (docs.vmware.com)

² Hardware Requirements for the VMware vRealize® Orchestrator Appliance™ (docs.vmware.com)

³ ESXi hardware requirements (docs.vmware.com)



APPENDIX D: HPE VIRTUALIZED NONSTOP—SYSTEM CONFIGURATIONS

HPE Virtualized NonStop is available in two configurations—Entry Class and High End. See the table below for details:

TABLE 3. HPE vNS System configurations

Parameter	HPE vNS Entry Class	HPE vNS High End	Remarks
Allowed CPUs	2 and 4	2 to 16 (even counts only)	
Memory per CPU	32 GB to 64 GB	64 GB to 256 GB	In 1 GB increments
Number of IP/Telco CLIMs supported	2 or 4	2 to 54 (even counts only)	Sum of IP vCLIMs and storage vCLIMs cannot exceed 56
Number of storage vCLIMs	2 or 4	2 to 54 (even counts only)	Sum of IP vCLIMs and storage vCLIMs cannot exceed 56
Support for native clustering	No	Yes	
Support for Expand (over IP)	Yes	Yes	



APPENDIX E: BUILDING HARDWARE BILL OF MATERIALS (BOM)—A REAL-WORLD EXAMPLE

Having understood the theory behind the HPE vNS architecture, the deployment rules and the requirements from the underlying hardware, let us now take a real-world example of a HPE vNS system and build the bill of materials for the hardware that is going to host such a system. Before we begin, a caveat to note is that there is no “one way” to put together the hardware for such a system. What we are doing is to apply our learnings to the given example. This will help hardware specialists to build a system using products that are suitable for their IT environment.

Let's build the hardware to host a production class High End HPE vNS system with the below configuration:

HPE NonStop CPUs—4

Cores per CPU—2

Memory per CPU—128 GB

Number of Storage CLIMs—4

Number of Network CLIMs—2

Storage for the system and user data (NSK volumes)—20 TB (mirrored)

HPE vNSC—1

vBackBox (vBB)—1

For those from the HPE NonStop world, it is evident that this is one of the common configurations used in production deployments. Such a system will have a total of 12 Virtual Machines (VMs)—4 CPU VMs, 6 CLIM VMs, and 2 Windows VMs (for HPE vNSC and vBB). Hence the first task would be to arrive at a distribution of the VMs into different physical servers complying with the deployment rules.

We will build the hardware in four steps: 1) Server hardware 2) Storage hardware 3) Networking hardware 4) Rack and peripherals.

Step 1: Determine the number and type of physical servers needed

For the sake of this discussion, we shall consider 2U rack mount servers as candidates to host these VMs. HPE vNS systems can require a large number of I/O resources such as adapters and NICs, and 2U servers provide an adequate number of I/O slots required to host those resources.

Step 1a: How many servers are required?

The HPE vNS deployment rule states that:

1. Each CPU belonging to a HPE vNS system should be hosted in a separate server
2. Two CLIMs belonging to a failover pair should be hosted on separate servers
3. All servers hosting CPUs should have identical processor models

Since there are 4 CPUs in this configuration, we need at least 4 physical servers. Apart from the CPU, we can use these servers to host:

1. One Storage CLIM
2. One Network CLIM

In addition to the above, we need resources to host one HPE vNSC and one vBB. We will use two of the four servers to host these VMs.

Note that increasing the number of VMs that are deployed in a host server increases the fault zone of that server. If the server encounters a fault and goes down, the fault will bring down all the VMs deployed on the server. Hence, even though more than one Storage or one Network CLIM can be hosted on a server (as long as they don't belong to the same failover pair), such a configuration should generally be avoided.

With the preceding considerations, a possible server configuration is as follows:

1. Two servers hosting one HPE vNS CPU, one storage CLIM, one Network CLIM each and either a HPE vNSC or vBB. The resource requirements for these servers are as below:
 - a. **Cores:** Each of the servers will require a total of 18 cores dedicated for HPE vNS VMs (2 for HPE vNS CPUs, 8 each for CLIMs) and 2 additional cores for HPE vNSC or vBB. If the HPE vNS CPUs are expected to be upgrade online at a future date, the CPU VM will need to be supported with 2 additional cores, which will make the total core count to 22. ESXi will require an additional 21% of the



dedicated cores i.e., $18 * 0.21 = 4$ (rounded up). Thus, a total of 24 cores (or 26 cores if CPUs are provisioned with 2 additional cores as stated above) will be required in each server.

- b. **I/O slots:** These servers will require at least 4 I/O slots. One for system interconnect fabric NIC (100GbE), one for Storage CLIM NIC (10GbE), two for Network CLIM NICs (10GbE). The HPE vNSC and vBB require 1GbE ports. If the server chassis does not have an embedded LOM with 1GbE ports, an additional 1GbE NIC would be required. Since both vBB and storage vCLIM are hosted on the same physical server, a dedicated NIC for connectivity between vBB and the storage vCLIM is not required.
 - c. **Memory:** Each server requires 152 GB of memory dedicated for HPE vNS VMs (128 GB for CPUs, 8 GB for storage CLIM, and 16 GB for network CLIM). We recommend 20% additional memory for effective ESXi operation. In addition, 8 GB each is need by ESXi, vBB and HPE vNSC, adding to a total of $152 * 1.20 + 24 = 207$ GB.
2. Two servers hosting one CPU and a storage CLIM each. The resource requirements for these servers are as below:
- a. **Cores:** These servers will require 10 (or 12 for future upgrade of HPE vNS CPUs) dedicated cores for HPE vNS VMs (2 or 4 for the HPE vNS CPU and 8 for the vCLIM) and an additional $10 * 0.21 = 3$ cores for ESXi i.e., a total of 13 or 15 cores.
 - b. **I/O slots:** These servers will require at least 2 I/O slots. One for system interconnect fabric NIC (100GbE) and one for Storage CLIM NIC (10GbE).
 - c. **Memory:** Each server requires 136 GB of memory dedicated for HPE vNS VMs (128 GB for CPUs and 8 GB for storage CLIM). We recommend 20% additional buffer for dedicated memory. This and 8 GB of memory need by ESXi adds to a total of $136 * 1.2 + 8 = 172$ GB.

NOTE

Server vendors provide rules to populate DIMMs in the servers for optimum performance. Follow those rules in addition to the above to determine the final configuration of the DIMMs.

Step 1b: Select the processor model

Since the target purpose of this server is for production deployment, an Intel Xeon Gold or Platinum processor with high core frequency and large Last Level Cache is recommended for higher performance. In the Intel Xeon Cascade Lake family of servers, the Intel Xeon Gold 6246R with 16 cores, 35.75 MB Last Level Cache, and 4.1 GHz maximum boost frequency is a suitable SKU for this example.

Step 1c: How many sockets to fill in the servers?

2U servers (such as [HPE ProLiant DL380](#)) come in single socket (1s) or dual socket (2s) configurations. Two key factors guide the selection between 1s and 2s or a combination of the two:

1. Number of cores required to host the VMs that need to be deployed on the server and the availability of processor SKUs of high core frequency and cache specifications that provide the required number of cores.
2. Availability of I/O slots to host the resources required to support the VMs deployed on the server
3. Licensing rules for vSphere ESXi. Each installed processor requires an ESXi license. You may also be using other third-party software that uses total socket or total core count in determining the licensing cost.

As described above, two types of servers are needed. One requires 24 cores and the other requires 13 cores. The server requiring 24 cores also requires at least 4 I/O slots. To meet these two requirements, the first server type needs two processors whereas the second server type needs only one processor. An HPE ProLiant DL380 server with two processors can have up to 8 I/O slots and an HPE ProLiant DL380 server with one processor can have up to 3 I/O slots. In addition, these servers have a chassis option with an embedded LOM with four 1GbE ports. This configuration of two servers with 2 processors and one server with one processor provides the required resources to host all the VMs in the HPE vNS system.

Step 1d: Identify the NICs required in these servers

Fabric NICs

The fabric network and the associated requirements are described in section [System Interconnect Fabric](#). The first task is to select the fabric speed. The fabric network speed can range between 25GbE and 100GbE. For this example system, 100GbE will be used.

The next task is to select the fabric NIC. HPE vNS supports a limited set of NICs for the fabric that are listed in [Appendix A, Section 2](#). For 100GbE, the HPE IB HDR100/EN 100G 2p 940QSFP56 adapter will be used. In a server with two processors that each host one or more HPE vNS VMs, a separate fabric NIC connected to each processor is recommended to support HPE vNS VM access to the fabric without requiring the VM to access the other processor. Thus, for best performance, the two servers with two processors will have two fabric NICs each, and the two servers with one processor will have one fabric NIC each.



iSCSI NIC

Each server will need to access external storage, either to access NSK volumes (for storage CLIMs) or to access OS disk of the CLIM VMs or Halted State Service (HSS) software for the HPE vNS CPUs. Thus, each server will require a NIC dedicated for iSCSI storage traffic. A NIC with hardware iSCSI offload does not require the server processor to handle the iSCSI protocol and can possibly improve the server processor performance for the HPE vNS VMs. In this example, an HPE FlexFabric 10Gb 2P 534FLR-SFP+ converged network adapter (CNA) which has hardware iSCSI offload is used.

Ethernet NIC

Network vCLIMs can support up to five Ethernet interfaces. In this example, four Ethernet interfaces will be provided for each Network vCLIM. Network vCLIMs support a limited selection of Ethernet NICs due to the need for CLIM drivers for PCI Passthrough access to the NIC. From the list in [Appendix A, Section 4](#), HPE Ethernet 10Gb 2P 530T Adapter is selected for the example HPE vNS system. Each 530T NIC supports two Ethernet Passthrough interfaces. The example HPE vNS system Network vCLIM requirement is for four PCI Passthrough interfaces, so two of the 530T NICs are needed in each server that will host a Network vCLIM. The example HPE vNS system has two IP vCLIMs which will be hosted in the two servers with two processors. Hence we shall have two 530T NICs in each of these two servers.

NOTE

If the interface bonding feature for the Ethernet interfaces in the Network vCLIMs to support NIC failover is not required, then any Ethernet NIC supported by VMware may be used.

Ethernet interfaces for other types of network connections:

In addition to the fabric, storage and network connections, HPE vNS has additional requirements for network connections (see Step 3) for maintenance LAN, management LAN, and Backup network. These can be provisioned over 1GbE ports. The HPE ProLiant DL380 server chassis with the embedded four 1GbE port LOM supports the required 1GbE ports. If your server does not have these resources, another Ethernet NIC could be added to the server to support the required 1GbE ports.

Outcome summary of step 1.

The HPE vNS system will be hosted on four HPE ProLiant DL380 servers. Two of these servers will each have two processors and the remaining two servers will each have one processor. The Intel Xeon Gold 6246R processor SKU will be used in all four servers.

The following diagram shows the processor cores in each of the four servers used by the VMs in the HPE vNS system:

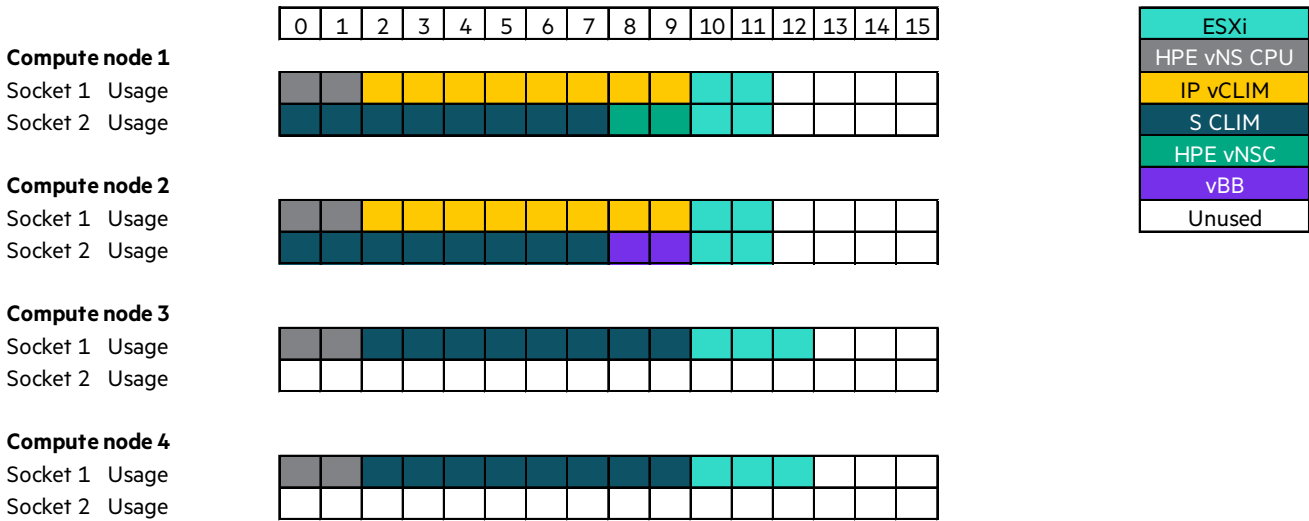


FIGURE 8. Layout of a possible core assignment to HPE vNS VMs and ESXi



The I/O slots of the DL380 Gen10 server will be filled as shown below.




FIGURE 9. Layout of the NICs in the I/O slots of DL380 Gen10 servers

NOTE

The above figure depicts the DL380 Gen10 system having LOM with four 1GbE ports

Step 2: Build the storage

After building the server configuration in step 1, the step 2 will be to build the storage configuration. A HPE NonStop system requires significant amount of storage. An entry level HPE NonStop system will typically require about 4 TB of raw storage. Hence defining storage hardware is the next important step in creating the BoM. There are many decision points while selecting your storage hardware. The two key requirements for HPE vNS from the storage hardware are:

1. It should be a block level storage
2. LUNs should be accessible from at least two host servers running CLIM failover pairs

From a technology perspective, three broad choices are available:

1. Internal drives of the hosts running storage vCLIMs
2. Internal drives of the hosts running HPE vNS VMs, accessed through a hyper-converged storage technology such as vSAN
3. External storage arrays

While a detailed discussion of these three choices is beyond the scope of this document, it is appropriate to note here that option (1) does not provide the 4-paths access to NSK volumes as discussed in [Appendix B: Storage Considerations](#). However, this may be an attractive option for development systems since it provides all HPE NonStop storage functions except path switching and is a fair trade-off between cost, performance, and availability characteristics.

2a. Selecting the storage product

Selecting the “right storage” for your HPE NonStop system requires careful considerations. It involves making several decisions, based on factors that may conflict with each other. For this discussion, we will use storage array-based technology. Storage arrays come in different levels of sophistication. HPE NonStop has its own storage performance and latency requirements, and implements host based mirrored disks analogous to RAID 1 arrays for availability. Hence a basic storage array product such as HPE MSA 2060 is sufficient for HPE vNS storage, particularly if the storage array is dedicated to the HPE vNS system. For best storage performance in production systems, consider using SSDs. For our example HPE vNS system, the [HPE MSA 2060 Storage Array](#) will be used.



2b. Arriving at target storage size

For the HPE NonStop system under consideration, the total storage requirement for NSK volumes is 20 TB. Since HPE NonStop implements host based mirrored disks, the required raw storage capacity is double that, i.e., 40 TB. In addition, various HPE vNS VMs need storage for OS images, as described in TABLE 1. Storage requirements for a typical HPE vNS system. The total storage requirement is calculated as (all in TB):

Purpose	Size
NSK volumes = 20 x 2	40.0
First pair of storage CLIMs = 0.3 x 2	0.6
Remaining CLIMs = 0.1 x 4	0.4
NSC	0.25
BackBox	0.3
Total	41.55

A large part of the storage requirement comes from NSK volumes (40 TB in our example) where the volumes exist in pairs of primary and mirror drives. According to the rules of deployment explained in the section [External Storage network](#), primary and mirror volumes should be provisioned on physically separate hardware for availability. Hence the total storage requirement should be grouped into two halves (20.755 TB in our example), and provisioned using physically separate hardware.

2c. Build one half of target storage

Due to Storage Mirroring, the storage can be specified for half of the total storage requirement (20.755 TB in the example) and then duplicated for the other half. While the HPE NonStop system requires a total of 20 TB of usable storage, the storage volumes used by a HPE vNS system are typically of the order of a few hundred GB. Hence the HPE NonStop system in our example uses a large number of such volumes.

The storage I/O performance depends on several factors, one of which is the number of I/Os to a physical drive. If a large number of storage volumes are provisioned on a storage drive, the I/O performance may be throttled due to a large number of I/Os simultaneously accessing a specific drive. By using a larger number of smaller capacity drives instead of a smaller number of larger capacity drives for equivalent total storage capacity, the same number of I/Os can be spread across more physical drives and improve storage performance.

The other significant factor affecting I/O performance is the storage drive hardware—HDD vs. SSD. The HPE vNS example system shall use SSDs for their superior price vs performance.

In MSA arrays, the smallest drive size available at the time of this writing is 960 GB. In order to provision 20.755 TB of storage, a minimum of $20.755/0.960 = 22$ drives are required. Storage arrays use RAID technology for availability. For our example implementation, we shall use RAID 5 and divide the 22 drives into four groups of 6 drives each to evenly balance the drives across two MSA arrays. This configuration requires $4 \times 6 = 24$ drives for data storage. Each RAID 5 group requires one additional drive for parity. The final implementation requires four RAID 5 groups with 7 drives per group (6 drives for data and 1 drive for parity) for a total of $4 \times 7 = 28$ 960 GB drives which will provide 23.04 TB of storage.

A single MSA array can host up to a total of 24 drives. For our example implementation, two MSA 2060 arrays will be used with each array containing two of the four RAID 5 groups.

2d. Build the other half of target storage

The two halves of storage should be identical to each other in every respect. Hence, to build the other half, duplicate the configuration that was specified for the first half.

The schematic figure below illustrates the storage hardware configuration for primary and mirror arrays.






FIGURE 10. Layout of storage drives in the four MSA 2060 arrays

Step 3: Networking the servers and storage

In this step, we will build the network required by the system. A HPE NonStop system requires following networks:

1. Fabric—This is a secure, high-speed network used for the system interconnect between the HPE vNS CPU and vCLIM VMs in the HPE vNS system.
2. Storage Area Network (SAN)—If the HPE vNS system uses external storage, it will connect to the storage resources (datastores in VMware parlance) over this network.
3. Maintenance LAN—This is a secure, low speed network used by the NonStop System Console (NSC) to perform configuration and system bring-up tasks on the HPE NonStop VMs.
4. Enterprise LAN—This network provides interfaces to the HPE NonStop system for the customer data center production LAN.
5. Management LAN—This is a low-speed network used by system management applications to access and manage the hardware hosting the HPE vNS system.
6. Backup LAN—If the HPE NonStop system backs up data using BackBox, this network provides a datapath to the backup storage solutions such as HPE StoreOnce.

The following sections describe these networks and the specification of the network hardware.

Step 3a: Build the fabric network

The key task in this step is the selection of the switch. As per the section [System interconnect fabric](#) we need two physically isolated fabric networks X and Y and hence two fabric switches are required for these. The next task is to count the number of fabric ports required on the switch which mainly depends on the number of fabric ports required for the system. In Step 1, six fabric NICs were specified for the four servers. Each NIC has two 100GbE ports—one for X fabric and one for Y fabric. Hence the X and Y fabric switches each require at least six 100GbE ports.

The HPE vNS system requires the fabric switches to support IEEE 802.3x Pause frames and DCB protocols for flow control.

For the example HPE vNS system, we will use a pair of [HPE SN2410bM 48SFP+ 8QSFP28 P2C](#) switches. This switch model has eight 100GbE ports and, in addition, has 48 10GbE ports. In step 3b, the 10GbE ports will be used for the SAN.

Step 3b: Build the Storage Area Network (SAN)

In Step 2, we built the external storage hardware to provision storage for the system. These need to be connected to the servers over an iSCSI or Fibre Channel (FC) network. The example HPE vNS system will use iSCSI for the SAN since it can be implemented over Ethernet switches instead of requiring separate Fibre Channel switches.

For availability considerations, a pair of switches are required in order to have an alternate access path to storage arrays in case of a switch failure. In Step 1, each server was configured with an HPE 534FLR-SFP+ NIC to be used by the storage vCLIM, and each NIC has two ports. The two ports of the storage NIC are connected to the two switches for high availability. The example HPE vNS system has four host servers with one storage NIC per server for a total of 8 ports to be connected to the two switches for the SAN. Thus, each switch requires 4 ports to implement these connections.

Each of the four MSAs from Step 2 has four ports for a total of 16 ports to be connected to the two switches for the SAN. Thus, each switch requires 8 ports to implement these connections.



Together, the four host servers and four MSAs require 12 ports on each of the two switches for the SAN. The SN2410bM 48SFP+ 8QSFP28 P2C switch selected in step 3a has 48 10GbE ports which supports the HPE vNS storage requirement for switch ports. Thus, the same pair of SN2410bM 48SFP+ 8QSFP28 P2C switches will be used to implement both the fabric and the SAN switch requirements of the example HPE vNS system.

Please note, as per section [Security Considerations](#), a switch that is used for the HPE vNS system interconnect fabric shall not be used to support any other network usage besides the HPE vNS system. In this example HPE vNS system implementation, the storage arrays are dedicated to the HPE vNS system and hence comply with this requirement. The SAN and the fabric traffic should be isolated using separate VLANs.

Similar availability and security considerations as above may be applied if you're building an FC network for storage or access storage from a central pool available in your data center.

Step 3c: Build the maintenance LAN

The maintenance LAN is required to establish a private, dedicated network for the HPE vNSC to communicate with HPE vNS VMs. This is a low speed 1GbE LAN connected to the servers hosting HPE vNS VMs and HPE vNSC. In our configuration, all these VMs are hosted on the four physical servers and one such 1GbE interface from each of these servers can be used for this purpose. If the server you use has 1GbE LoM ports, you may use that as the interface for maintenance LAN. Else, a dedicated 1GbE NICs should be included in the servers.

We need a switch that has at least four 1GbE ports to connect to the four servers and we will use [Aruba 2930F 24G Switch](#) for this purpose.

Step 3d: Connectivity to the enterprise network

In steps 3a and 3c, we created networks that are typically dedicated for the HPE vNS system whereas in step 3b, you may have either created a dedicated storage network for your HPE NonStop or connected the servers to a dedicated storage pool using an enterprise-wide SAN. The next set of networks we create allow the HPE NonStop system to connect to the various types of datapaths running in a typical enterprise and their connections to resource pools such as for backup.

The connection to a HPE NonStop system from Enterprise LAN is through IP CLIMs. The HPE NonStop system under consideration here has two IP CLIMs, as they're always in pairs, and each IP CLIM has four 10GbE interfaces. We will connect these eight interfaces to a patch panel in the rack so that the IT administrators can connect them to the required Enterprise network to allow external applications connect to the HPE vNS host.

Step 3e: Connectivity for virtual BackBox (for data backups)

The HPE vNS system has a virtual BackBox (vBB) which is an application that runs on a Windows VM on one of the host servers. It does data backup and restore functionality for the system. It presents backup volumes to the storage vCLIMs connected to it. If vBB is not co-located on the same host as the storage vCLIM, you need to establish an iSCSI path between the two.

For data backup, vBB connects to Enterprise backup solutions such as HPE StoreOnce. Typically, the connectivity to backup devices is established through a separate network. For this purpose, a 1GbE network is sufficient. As mentioned earlier, we will use one of the remaining three 1GbE LoMs for the purpose. The network port used for this connectivity will be connected to the patch panel to allow the administrator to connect it to the backup network, similar to how that is done for Ethernet ports of IP vCLIMs in Step 3d above.

Step 3f: Connectivity to the management network

The hardware hosting HPE vNS resources (servers, switches, storage arrays) need to be managed using customer's system management solutions (e.g., HPE OneView). vSphere ESXi running on hosts also need to be managed using a vCenter. All these management typically happens from an operations center which is a part of the same data center hosting the hardware or may be remote from it. Each of these hardware products provide a separate management port for this purpose (e.g., HPE iLO of HPE ProLiant servers). There are two ways to achieve this:

1. Connect the management ports of hardware products and HPE iLO of DL380 servers to the maintenance switch. The external management appliances reach this hardware through the maintenance switch. For security reasons, isolate the maintenance and management networks at layer 2 using VLANs. The BoM in this document uses this approach.
2. Use a separate 1GbE switch to connect the management ports of hardware products and HPE iLO of DL380 servers. The external management appliances reach the hardware through this (dedicated) switch. The physical separation in this option offers superior security compared to option (1) at the cost of an additional hardware.

This completes the networking configuration for the hardware. See the connectivity diagrams [Figure 15](#) Fabric Connectivity for the HPE vNS system and [Figure 16](#) Storage Networking Connectivity for the HPE vNS system for illustration.



Step 4: Build the rack:

From steps 1 to 3, we have identified all the hardware required to host the example HPE vNS system. This hardware may be housed in a rack that is dedicated for the HPE vNS system or shared with hardware used by other systems. If the rack is shared with hardware for other systems, the hardware components should be clearly identified to avoid mistakes in servicing hardware.

For our example HPE vNS system, the hardware will be installed in a dedicated 42U rack.

Note the example HPE vNS system does not specify an Uninterrupted Power Supply (UPS) to cover power outages. If the data center infrastructure supplying power to the HPE vNS system hardware does not have an UPS, then a local UPS can be added to the rack.

Step 5: Software licenses

To host a HPE vNS system, each server requires vSphere Enterprise Plus edition. Two of the servers hosting HPE vNSC and vBB require Windows 2019 standard edition. Each of the DL380 servers also require license for HPE iLO advanced.

The complete bill of materials arrived at from step 1 through to step 4 is provided in the section [Hardware Bill of Materials for the target system](#) below. Please look for the column titled “Purpose” to determine the hardware purpose. Some of the parts are specific to the North America region (e.g., the power socket specification) and may need to be customized for your region. The rack layout for the hardware is given in the section [Hardware rack layout for the target system](#). Customers may follow the example rack layout or may follow any other order as they prefer.

CONNECTIONS

An important aspect of putting the hardware together for HPE vNS is the connectivity among the various elements. A HPE vNS system requires a large number of cable connections. The following sections describe the cabling for the HPE vNS system hardware.

Fabric network

A HPE vNS system requires two fabric networks, X and Y, dedicated to the HPE vNS system and not shared with any other use. The switches and cables of X and Y networks should be physically separate.

We will connect the NICs to the switches in the following two ways:

1. The NIC in servers with one processor will be connected such that one port connects to the “X Switch” and the other port connects the “Y switch”.
2. The two NICs in servers with two processors will be connected such that both the ports of one NIC connect to the “X switch” and those of the other NIC connect to the “Y switch”. Availability of two NICs in a server gives us the option to physically isolate X and Y traffic at the NIC level rather than at the port level. This provides higher availability by supporting the fabric traffic of all VMs in the host server on one fabric if the NIC for the other fabric should fail.

Storage Network:

The storage network in the example HPE vNS system establishes a redundant path between the host servers running HPE NonStop VMs and the storage arrays providing redundant storage drives. The example HPE vNS system also shares the switches between the system interconnect fabric and storage network. This is allowed because the storage network is dedicated to the HPE vNS system and is physically isolated from any external network. Please note that the specified switch should have enough switch processing capacity to support both HPE vNS fabric traffic and storage traffic. Separate pairs of switches, one dedicated to the system interconnect fabric and the other dedicated to the storage network, can be specified to overcome switch processing capacity limitations.

Each host server uses a 534FLR-SFP+ two-port adapter to support storage iSCSI traffic. For high availability, one of the two adapter ports is connected to the “X switch” and the other adapter port is connected to the “Y switch”. This storage network connection configuration ensures connectivity between the storage vCLIMs in the host servers and the storage arrays if one of the two switches should fail.

Each MSA has two controllers and each controller has four SFP+ ports. The HPE vNS system configuration does not require all 8 SFP+ ports in each MSA storage array to be connected to support HPE vNS storage traffic. For high availability, one SFP+ port in each MSA controller is connected to the “X switch” and a second SFP+ port in each MSA controller is connected to the “Y switch”. The remaining two SFP+ ports in each MSA controller are left unconnected. This storage network connection configuration maintains access between both storage pools in the MSA storage arrays and the storage vCLIMs in the host servers if one of the two switches should fail.

Overall, the storage network connection configuration supports continued access between the storage vCLIMs and the MSA storage arrays after a NIC adapter port failure, storage network switch failure, or MSA controller failure.

Refer to the [HPE MSA 2060 1060/2060/2062 Installation Guide](#) section for an illustration of a SAN cable connection configuration that could be used for the network between the HPE vNS system host servers and the MSA storage arrays.



Since the fabric switches in the example HPE vNS system are used for both system interconnect fabric and storage networks, the two networks should be logically separated using VLANs in the switches.

Maintenance Network

Having built the fabric and storage network, let's now build the maintenance LAN. Remember that this network is used to establish a secure, private, and physically isolated network between HPE vNSC and HPE NonStop VMs hosted on servers. As mentioned earlier, we will use one of the LoM ports of each of these servers to the switch for this purpose

Management LAN

As described in Step 3f, this setup will share the switch between maintenance network and management network. The HPE iLO port of DL380 servers, the management ports of the SN2410bM and Aruba 2530 switches, and the management ports of each of the MSA controllers (two per MSA) are connected to the Aruba switch which in turn will be connected to the data center management LAN to establish connectivity between management applications and the hardware.

Enterprise LAN

The Enterprise LAN connectivity is used for provisioning external connectivity to the HPE vNS system. The Ethernet ports of IP CLIMs are used for the purpose. The network resources are provided by the data center. The Ethernet ports of network CLIMs are connected to a patch panel to which the data center's network cables can be connected.

Backup LAN

Last but not the least is the connectivity required to connect the virtual BackBox to the data center backup network. This could be done by connecting one of the remaining two 1GbE LoM ports of the server hosting the vBB to the patch panel. From there, it can be connected to the backup network through a separate cable.

The following schematic diagrams illustrates how all the connectivity requirements are implemented:

For servers with one processor







FIGURE 11. Connectivity of I/O card ports in 2s servers

For servers with two processors**FIGURE 12.** Connectivity of I/O card ports in 1s servers**For MSA****FIGURE 13.** Connectivity of MSA 2060 I/O ports

Rack Layout

**FIGURE 14.** Rack layout of the HPE vNS system hardware

Bill of materials

Qty	Product #	Product Description	Purpose	Remarks
1	P9K10A	HPE 42U 600x1200mm Adv G2 Kit Shock Rack	Rack	
1	P9K10A#001	HPE Factory Express Base Racking Service	Rack	
2	P9Q53A	HPE G2 Basic 3Ph 8.6kVA/C19 NA/JP PDU	Rack	This is for U.S. and Japan. Change this to your geography and power drop connectors.
2	P9Q53A#0D1	Factory Integrated	Rack	
1	AF630A	HPE LCD8500 1U US Rackmount Console Kit	Rack	
1	AF630A#0D1	Factory Integrated	Rack	
4	868703-B21	HPE ProLiant DL380 Gen10 8SFF CTO Server	Server	
4	868703-B21#0D1	Factory Integrated	Server	
4	P24472-L21	Intel Xeon-G 6246R FIO Kit for DL380 G10	Server	1 FIO processor for each of the CPU servers
2	P24472-B21	Intel Xeon-G 6246R FIO Kit for DL380 G10	Server	1 add-on processor in socket #2 of servers 0 and 1
2	P24472-B21#0D1	Factory Integrated	Server	
60	P00920-B21	HPE 16GB (1x16GB) Single Rank x4 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit	Server	12 DIMMs socket #1 of servers 2 and 3. 9 DIMMs in each sockets of servers 0 and 1
60	P00920-B21#0D1	Factory Integrated	Server	
2	870548-B21	HPE DL Gen10 x8 x16 x8 Rsr Kit	Server	1 riser kit each for 2 CPU servers
2	870548-B21#0D1	Factory Integrated	Server	
4	656596-B21	HPE Ethernet 10Gb 2P 530T Adptr	Server	In slots 4 and 6 of servers 0 and 1
4	656596-B21#0D1	Factory Integrated	Server	
4	700751-B21	HPE FlexFabric 10Gb 2P 534FLR-SFP+ Adptr	Server	In the FLR slot of all servers
4	700751-B21#0D1	Factory Integrated	Server	
6	P06251-B21	HPE IB HDR100/EN 100G 2p 940QSFP56 Adptr	Server	In slot 2 of all servers and in slot 5 of servers 0 and 1
6	P06251-B21#0D1	Factory Integrated	Server	
8	872477-B21	HPE 600GB SAS 10K SFF SC DS HDD	Server	2 SAS drives for each of the servers as the local disk for vSphere ESXi
8	872477-B21#0D1	Factory Integrated	Server	
4	804331-B21	HPE Smart Array P408i-a SR Gen10 Ctrlr	Server	In each of the Servers
4	804331-B21#0D1	Factory Integrated	Server	
8	865414-B21	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	Server	2 in each of the servers
8	865414-B21#0D1	Factory Integrated	Server	
4	733660-B21	HPE 2U Small Form Factor Easy Install Rail Kit	Server	In each of the servers
4	733660-B21#0D1	Factory Integrated	Server	
4	P01366-B21	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit	Server	In each of the servers
4	P01366-B21#0D1	Factory Integrated	Server	
2	Q6M28A	HPE SN2410bM 48SFP+ 8QSFP28 P2C Swch	Fabric + Storage Switch	Fabric + Storage Switch
2	Q6M28A#0D1	Factory Integrated	Fabric + Storage Switch	
4	R0Q76A	HPE MSA 2060 10GbE iSCSI SFF Storage	Storage	MSAs
4	R0Q76A#0D1	Factory Integrated	Storage	
56	R0Q46A	HPE MSA 960GB SAS 12G Read Intensive SFF (2.5in) M2 3yr Wty SSD	Storage	Each MSA to have two RAID 5 groups of 6 + 1 each
56	R0Q46A#0D1	Factory Integrated	Storage	



Qty	Product #	Product Description	Purpose	Remarks
19	C7535A	HPE Ethernet 7ft CAT5e RJ45 M/M Cable	Maint./Mgmt. LAN	
19	C7535A#0D1	Factory Integrated	Maint./Mgmt. LAN	
16	JD096C	HPE X240 10G SFP+ SFP+ 1.2m DAC Cable	Storage	Cable for connecting MSAs controller ports to the switch
16	JD096C#0D1	Factory Integrated	Storage	
8	JD097C	HPE X240 10G SFP+ SFP+ 3m DAC Cable	Storage	Cable for connecting SFP+ ports of 534FLR NIC to the switch
8	JD097C#0D1	Factory Integrated	Storage	
12	845406-B21	HPE 100Gb QSFP28 to QSFP28 3m DAC	Fabric	Cable for connecting QSFP56 ports of ConnectX-6 to the switch
12	845406-B21#0D1	Factory Integrated	Fabric	
1	JL253A	Aruba 2930F 24G Switch	Maint. LAN	
1	JL253A#0D1	Factory Integrated	Maint. LAN	
1	J9583B	Aruba X414 1U Universal 4-post Rack Mount Kit	Maint. LAN	Rack-mount kit for the Aruba switch
6	BD714AAE	VMw vSphere EntPlus 1P 1yr E-LTU	Server	1 ESXi license per processor socket
4	E6U59ABE	HPE iLO Adv Elec Lic 1yr Support	Server	One per each server
4	E6U59ABE#0D1	Factory Integrated	Server	
1	P11060-B21	MS WS19 (16-Core) Std FIO Npi en SW	Server	Windows license for servers 0 and 1
1	P11064-DN1	MS WS19 (16-Core) Std Add Lic AMS SW	Server	Windows license for servers 0 and 1 (to match the total core count)
1	P11073-DN1	MS WS19 RDS 5USR CAL AMS LTU	Server	
1	P11077-DN1	MS WS19 5USR CAL en/fr/es/xc LTU	Server	

Fabric Connectivity for the HPE vNS system

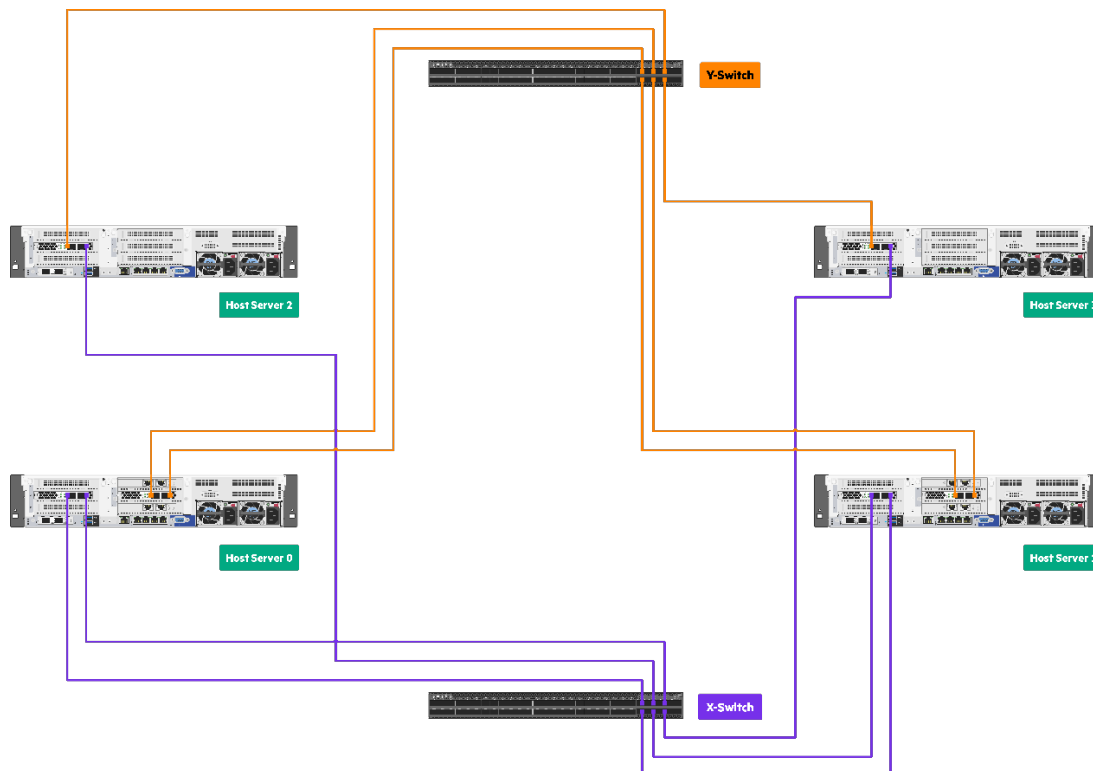


FIGURE 15. Connection diagram for the fabric network

Storage Networking Connectivity for the HPE vNS system

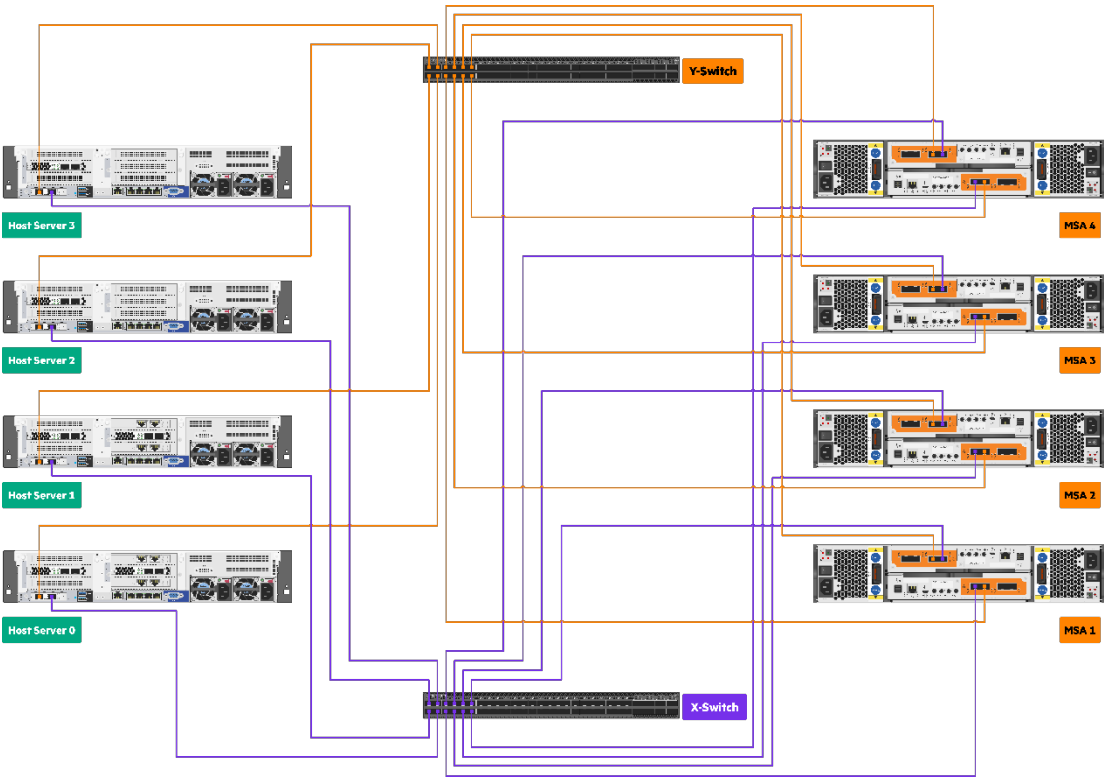


FIGURE 16. Connection diagram for SAN

Small hardware footprint to host two HPE vNS systems!!


Here’s an interesting proposition. How small a hardware can be if it were to host two fully functional HPE vNS systems. Well, below is a configuration that can accomplish it. The core layout and the hardware BoM has been provided here. The cabling is left to the reader as an exercise!

Two 2 CPU HPE vNS systems can be deployed on two 2-socket servers. On each of the two sockets of these servers a CPU, IP vCLIM, and storage vCLIM of one of the HPE vNS systems can be deployed. Each server can thus host six VMs (CPU, IP vCLIM, storage vCLIM belonging to one system x 2) and either a HPE vNSC or a vBB. The processors must have sufficient number of cores to host all these VMs.

The configuration in this section provides an example for two Entry Class 2-core systems suitable for development. Each system uses 8-cores for CLIMs. Internal storage with HDDs is used to provide storage resources. Each CPU is provisioned with 64 GB of memory and each HPE vNS system is provisioned with a little more than 2 TB storage for NSK and user volumes. As described in the section [Multipath access between CLIM and storage volumes](#), such a configuration does not support 4-path storage.

Layout of the cores for VMs

Below is the layout of the processor core assignment to the VMs across the two processor sockets:



Note: The actual core assignment is done by ESXi and may vary from the above

FIGURE 17. Layout of VMs over processor cores for the two HPE vNS systems



The Server layout




FIGURE 18. Layout of servers and switches for the two HPE vNS systems

Hardware bill of materials to host two HPE vNS systems on a small footprint

Qty	Product #	Product Description	Purpose
1	P9K10A	HPE 42U 600x1200mm Adv G2 Kit Shock Rack	Rack
1	P9K10A#001	HPE Factory Express Base Racking Service	Rack
2	P9Q53A	HPE G2 Basic 3Ph 8.6kVA/C19 NA/JP PDU	Rack
2	P9Q53A#0D1	Factory Integrated	Rack
1	AF630A	HPE LCD8500 1U US Rackmount Console Kit	Rack
1	AF630A#0D1	Factory Integrated	Rack
2	868703-B21	HPE ProLiant DL380 Gen10 8SFF CTO Server	Server
2	868703-B21#0D1	Factory Integrated	Server
2	P23553-L21	Intel Xeon-G 5220R FIO Kit for DL380 G10	Server
2	P23553-B21	Intel Xeon-G 5220R Kit for DL380 G10	Server
2	P24472-B21#0D1	Factory Integrated	Server
16	P00924-B21	HPE 32GB (1x32GB) Dual Rank x4 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit	Server
16	P00924-B21#0D1	Factory Integrated	Server
2	870548-B21	HPE DL Gen10 x8 x16 x8 Rsr Kit	Server
2	870548-B21#0D1	Factory Integrated	Server
4	656596-B21	HPE Ethernet 10Gb 2P 530T Adptr	Server
4	656596-B21#0D1	Factory Integrated	Server
2	700751-B21	HPE FlexFabric 10Gb 2P 534FLR-SFP+ Adptr	Server
2	700751-B21#0D1	Factory Integrated	Server
2	P06251-B21	HPE IB HDR100/EN 100G 2p 940QSFP56 Adptr	Server
2	P06251-B21#0D1	Factory Integrated	Server
4	872479-B21	HPE 1.2TB SAS 10K SFF SC DS HDD	Server
4	872479-B21#0D1	Factory Integrated	Server
2	804338-B21	HPE Smart Array P816i-a SR Gen10 Ctrlr	Server
2	804338-B21#0D1	Factory Integrated	Server
4	865414-B21	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit	Server
4	865414-B21#0D1	Factory Integrated	Server
2	733660-B21	HPE 2U Small Form Factor Easy Install Rail Kit	Server
2	733660-B21#0D1	Factory Integrated	Server



Qty	Product #	Product Description	Purpose
2	P01366-B21	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit	Server
2	P01366-B21#0D1	Factory Integrated	Server
2	ROP77A	HPE SN2010M 18SFP28 4QSFP28 P2C TAA Swch	Switch
2	ROP77A#0D1	Factory Integrated	Switch
10	870759-B21	HPE 900GB SAS 15K SFF SC DS HDD	Storage
10	870759-B21#0D1	Factory Integrated	Storage
7	C7535A	HPE Ethernet 7ft CAT5e RJ45 M/M Cable	Maint. LAN
7	C7535A#0D1	Factory Integrated	Maint. LAN
4	845406-B21	HPE 100Gb QSFP28 to QSFP28 3m DAC	Fabric
4	845406-B21#0D1	Factory Integrated	Fabric
1	JL253A	Aruba 2930F 24G Switch	Maint. LAN
1	JL253A#0D1	Factory Integrated	Maint. LAN
1	J9583B	Aruba X414 1U Universal 4-post Rack Mount Kit	Maint. LAN
4	BD714AAE	VMw vSphere EntPlus 1P 1yr E-LTU	Server
2	E6U59ABE	HPE iLO Adv Elec Lic 1yr Support	Server
2	E6U59ABE#0D1	Factory Integrated	Server
2	P11060-B21	MS WS19 (16-Core) Std FIO Npi en SW	Server
4	P11064-DN1	MS WS19 (16-Core) Std Add Lic AMS SW	Server
1	P11073-DN1	MS WS19 RDS 5USR CAL AMS LTU	Server
1	P11077-DN1	MS WS19 5USR CAL en/fr/es/xc LTU	Server

REFERENCES

HPE Virtualized NonStop Deployment and Configuration Guide for VMware available at HPESC (hpe.com/info/nonstop-ldocs).

QuickSpecs: HPE ProLiant DL380 Gen10 server (hpe.com/psnow/doc/a00008180enw).

QuickSpecs: HPE Synergy 480 Gen10 compute module (hpe.com/psnow/doc/a00008520enw).

QuickSpecs: HPE Synergy 660 Gen10 compute module (hpe.com/psnow/doc/a00008522enw).

Technical white paper: HPE NonStop OS—Provide the availability and scalability advantage to your business at a low TCO
hpe.com/psnow/doc/4aa0-6149enw



LEARN MORE AT

hpe.com/info/nonstop

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

© Copyright 2019, 2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel, Intel Xeon, Intel Xeon Silver, and Intel Xeon Gold are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Windows and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware ESXi, VMware vCenter, VMware vCloud, VMware vRealize, VMware vRealize Orchestrator, VMware vSphere, VMware vSphere Enterprise Plus Edition, VMware, VMware vCenter Server Appliance, VMware vSAN, and VMware vRealize Orchestrator Appliance are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.