Hewlett Packard Enterprise

eGTUG 2024

HPE Virtualized NonStop Continues the Journey to the Cloud

Arun Fernandes and Lars Plum

April 2024

Forward-looking statements

This is a rolling (up to three year) Roadmap and is subject to change without notice

This document contains forward looking statements regarding future operations, product development, product capabilities and availability dates. This information is subject to substantial uncertainties and is subject to change at any time without prior notification. Statements contained in this document concerning these matters only reflect Hewlett Packard Enterprise's predictions and / or expectations as of the date of this document and actual results and future plans of Hewlett Packard Enterprise may differ significantly as a result of, among other things, changes in product strategy resulting from technological, internal corporate, market and other changes. This is not a commitment to deliver any material, code or functionality and should not be relied upon in making purchasing decisions.

Agenda

HPE Virtualized NonStop: Navigating to public clouds

Current areas of focus

- **HPE Virtualized NonStop on public clouds: initial solution**
- **HPE Virtualized NonStop on public clouds: long-term solution**

Additional areas of investigation

Q & A

HPE Virtualized NonStop: Navigating to Public Clouds

Lessons learned from the field: vNS deployments in public clouds

- Although vNS adoption on private VMware clouds continues to increase, several NonStop customers have also shared their plans to migrate workloads to public clouds
- Driving forces for this trend include:
  - “Cloud first” strategy driven by IT
  - Migrate IT to an “aaS” procurement model
  - Meet ESG (environmental, social, and governance) goals

Our response to customer demand for vNS in the cloud

- Both converged and Virtualized NonStop can already be deployed in Equinix data centers through co-location
- Two hyperscalers (Microsoft Azure and GCP) offer solutions that meet these requirements in some geographies
- The next few slides summarize customer requirements, current hyperscaler offerings, and gaps to be addressed for a vNS solution in the cloud

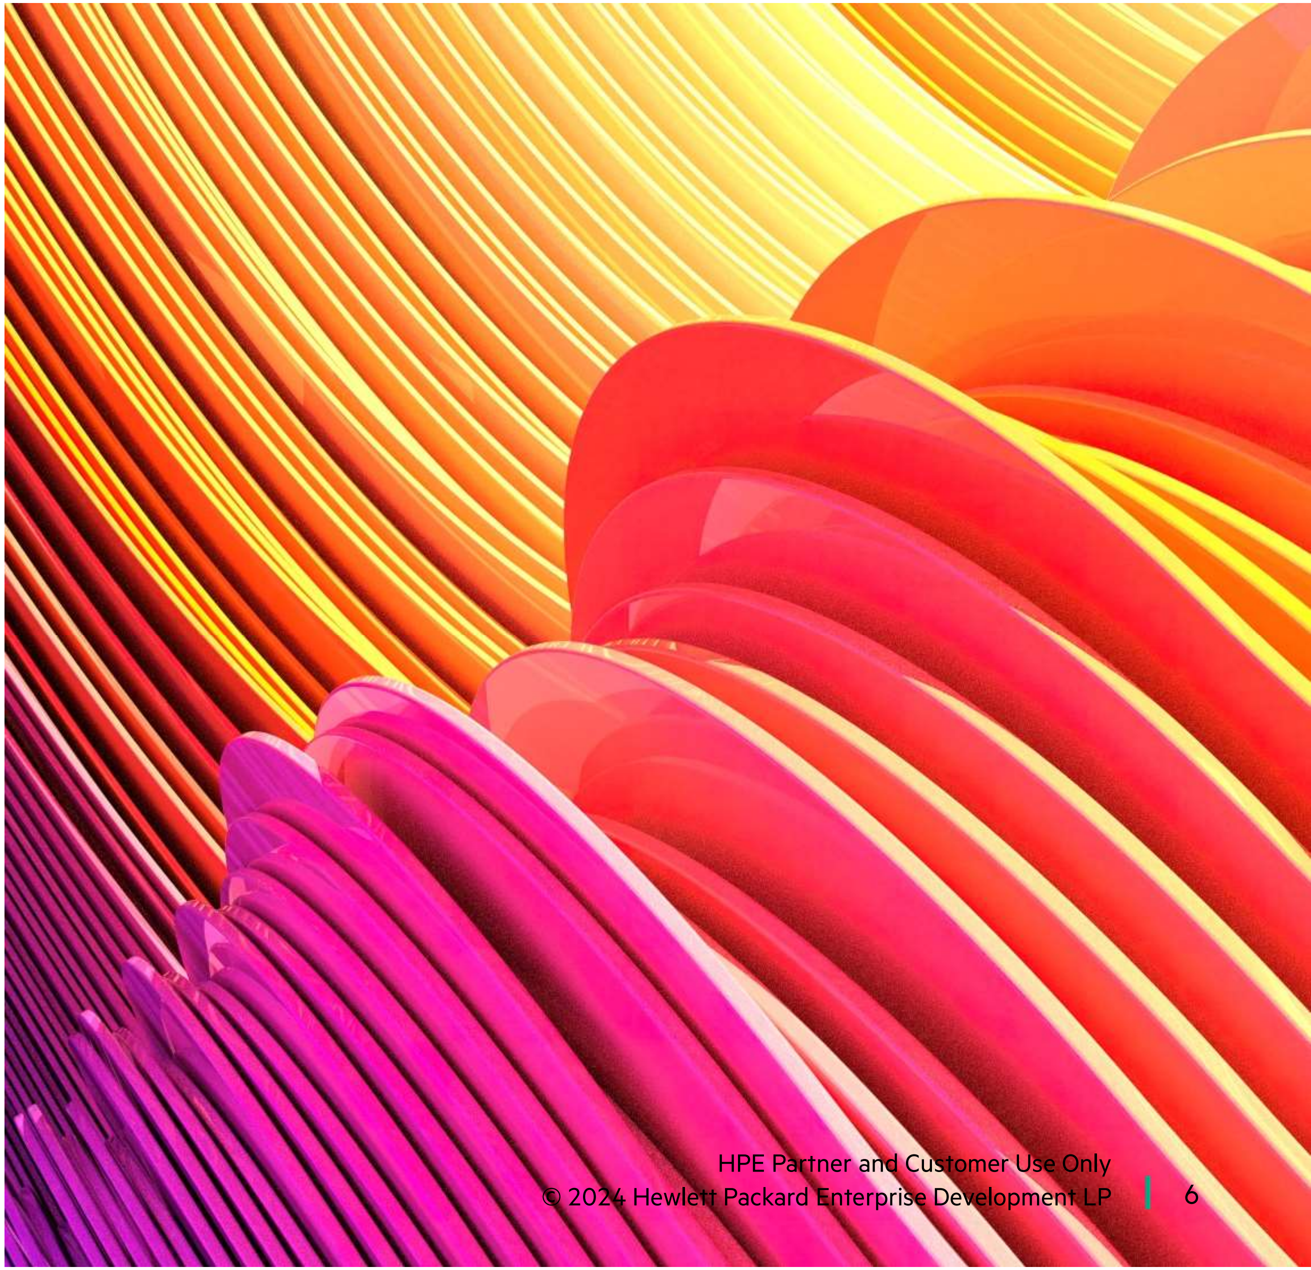

HPE Virtualized NonStop in public clouds

High-level requirements identified through customer surveys

Function	Requirement
Hyperscalers	Inquiries received for Microsoft Azure, Google Cloud Platform, and Amazon AWS
Production system geographies	Inquiries received for in-country data centers in multiple geographies
Development and test system geographies	Some inquiries suggest flexibility for out-of-country data centers in multiple geographies
Hardware	General fleet servers preferred to broaden deployment options and reduce costs
Software compatibility	Must run existing L-Series applications without changes
Deployment and configuration	Ease of deployment and configuration
Security	Secure system and data
Availability	Achieve the same availability as the current vNS solution
Manageability and operations support	Inquiries received for services such as HPE GreenLake Managed Services for NonStop, but cloud providers will also play a role in the upkeep of the infrastructure
Billing	Support for HPE GreenLake and consumption based billing
Licensing	Bring Your Own License (BYOL) model is a possibility

HPE Virtualized NonStop in public clouds

Offerings identified during our discussions with hyperscalers

Function	Hyperscaler offerings
RoCE NICs and VMware	Supported by two out of three hyperscalers with specific solutions in select geographies
System fabrics	Private redundant Ethernet switches
Storage	Redundant storage options identified so far include VMware vSAN and internal drives in hosts
Security	Dedicated servers with secure access through the cloud provider network security architecture
General fleet servers	Widely available as dedicated servers in multiple geographies, but have general-purpose Ethernet NICs and the hyperscaler's native hypervisor (e.g. Hyper-V for Microsoft Azure) in lieu of RoCE NICs and VMware
Orchestration	Multi-cloud orchestration frameworks such as Terraform available VMware vRO is a possible offering with specific solutions in select geographies
Fault-tolerant deployment	Supported through anti-affinity rules, but approaches to specify VM placement vary between offerings
Dedicated cores and memory	Supported, but approaches to specify these dedicated resources vary between hypervisors
Infrastructure upkeep	Hyperscalers generally prefer to be in charge of hardware and virtualization infrastructure maintenance and upgrades, but will generate alerts for planned and unplanned maintenance. Therefore, the scheduling of maintenance windows must be mutually agreed between the customer and the hyperscaler.

HPE Virtualized NonStop in public clouds

Main gaps identified during our discussions with hyperscalers

Function	Current areas of focus
RoCE NICs and VMware	• Two hyperscalers (Microsoft Azure and GCP) offer solutions that meet these requirements in some geographies • NonStop R&D has worked with these hyperscalers to deploy several Proof of Concept vNS systems in their data centers • However, availability of this solution will be limited to geographies that offer servers with RoCE adapters and VMware

General fleet servers	• Hyperscalers populate their broadly available general fleet servers with general-purpose Ethernet adapters
	• Moreover, general fleet servers generally run the hyperscaler's native hypervisor (e.g., Hyper-V in Microsoft Azure)
	• Because of this, NonStop R&D has achieved significant progress towards:
	    • A long-term Ethernet-based vNS solution that can leverage public cloud general fleet servers
	    • Support for additional hypervisors beyond VMware ESXi and Linux KVM (supported by vNS and NS2, respectively)

Orchestration	• NonStop R&D has achieved significant progress in developing generalized vNS Deployment Tools based on the popular Terraform multi-cloud orchestration framework
Infrastructure upkeep	• NonStop R&D discussed operational scenarios such as host upgrades with hyperscalers • Although live vMotion of virtual machines is commonly used by hyperscalers for host upgrades in their VMware-based solutions, vNS supports rolling upgrades of ESXi hosts through vMotion of powered off VMs

Current areas of focus

vNS deployment on public clouds

- NonStop R&D is pursuing two programs in parallel for vNS on public cloud solutions:
  - An initial solution that leverages RoCE and VMware offerings from Microsoft Azure and GCP
  - A long-term solution that can be deployed on general fleet servers from all major hyperscalers

vNS on public cloud – Initial solution

- Our discussions with hyperscalers identified two cloud offerings that can be leveraged for an initial vNS on the cloud solution:
  - Microsoft Azure Skytap
  - Google Cloud VMware Engine
- These cloud offerings support RoCE NICs and VMware
- They are available in several, but not all, geographies supported by Azure and GCP

Microsoft Azure Proof of Concept (PoC) vNS systems

- A 1st vNS PoC system was deployed in a Microsoft Azure data center in the UK in October 2022 and announced during the 2022 TBC
  - This system has been upgraded since, but it is still running
  - This PoC system passed stress and fault tolerance tests performed by the NonStop QA team without issues
- More recent activities include 3 PoC systems in 3 different geographies for 3 different NonStop customers
  - One of these PoCs has already completed successfully
  - Two additional PoCs are expected to start soon
- A vNS on Azure PoC system was used for the following demos during the 2023 E-BITUG in May:
  - HPE NonStop SQL/MX DBS (Database Services) on public cloud
  - HPE Shadowbase replication between a vNS on Azure PoC system and a vNS system owned by TCM (an HPE NonStop ISV partner)

Azure

Microsoft Azure Proof of Concept (PoC) vNS systems (continued)

- Common aspects of all vNS on Azure POCs carried out so far:
  - The PoCs leverage a Microsoft Azure solution referred to as Skytap, which supports RoCE NICs and VMware
  - The PoCs have redundant fabric switches and use internal drives in the servers as VMware datastores
  - The VM system fabric interfaces were deployed as SR-IOV interfaces, as currently required by our vNS solution
  - The VMware environment includes vRO (vRealize Orchestrator), allowing the system to be deployed with our vNS Deployment Tools.

Azure

Google Cloud Platform Proof of Concept (PoC) vNS systems

- A 1st vNS PoC system was deployed in a Google Cloud Platform (GCP) data center in the US in October 2022 and announced during the 2022 TBC
  - This system has been upgraded since, but it is still running
  - This PoC system passed stress and fault tolerance tests performed by the NonStop QA team without issues
- The 1st PoC system was deployed through direct interactions with vCenter
- A 2nd PoC system was deployed by cloning VMs from the 1st PoC system and changing IP addresses as needed
  - This 2nd PoC system was deployed in the same ESXi hosts used for the 1st PoC system
- The 2nd POC on GCP system was used for a more comprehensive internal POC as follows:
  - Safe Balloting POC between a vNS system on GCP and a vNS system in an HPE lab in Alpharetta, GA
  - For the Safe Balloting POC, the Gravic Shadowbase Validation Architecture Module (VAM) was utilized to verify the transactional integrity between the GCP and Alpharetta vNS systems

Google Cloud

Google Cloud Platform Proof of Concept (PoC) vNS systems

- A 2nd and separate GCP cloud has since been provisioned for a customer PoC
  - This 2nd cloud includes vRO, allowing vNS system deployment through the vNS Deployment Tools
  - This customer PoC recently completed successfully
- Common aspects of all vNS on GCP POCs carried out so far :
  - The PoCs leverage the Google Cloud VMware Engine (GCVE) solution, which supports RoCE NICs and VMware
  - The PoCs have redundant fabric switches and use VMware vSAN datastores
  - The VM system fabric interfaces were deployed as SR-IOV interfaces, as currently required by our vNS solution

Google Cloud

vNS on public cloud – Long-term solution

- The long-term solution will support general fleet servers from all major hyperscalers
- Main goals:
  - Leverage general-purpose Ethernet NICs for system fabric traffic
  - Support for hyperscaler native hypervisors
  - Multi-cloud orchestration

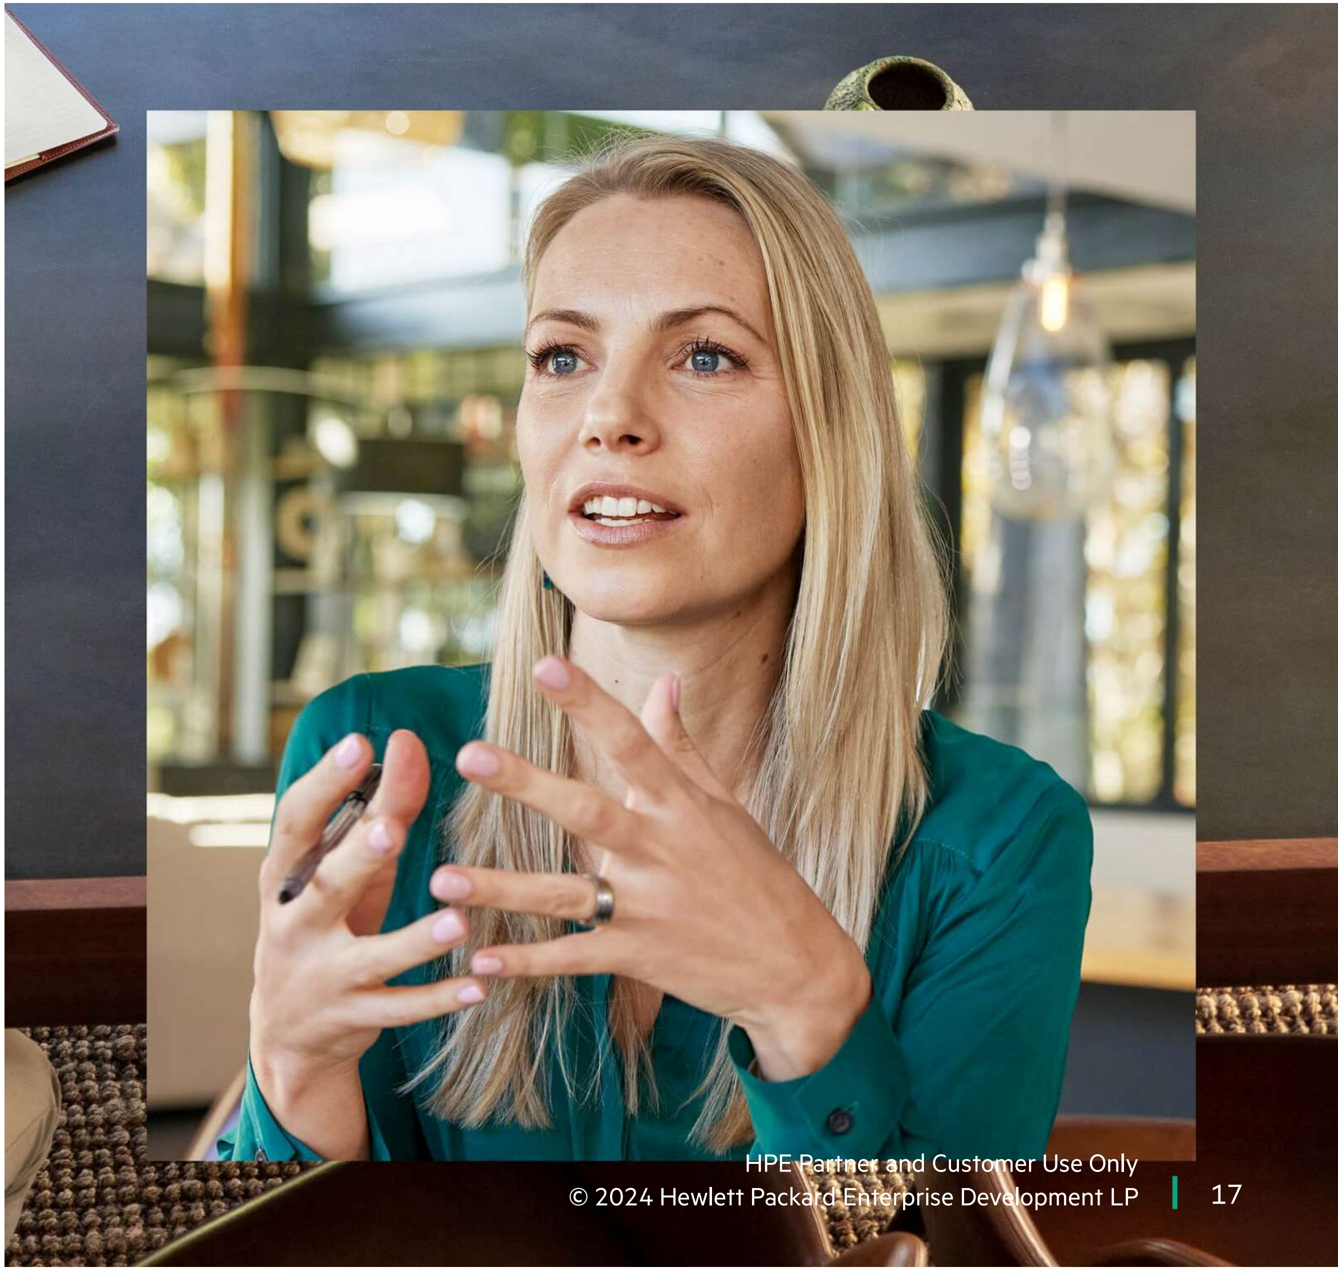

Current vNS stack with RoCE NICs

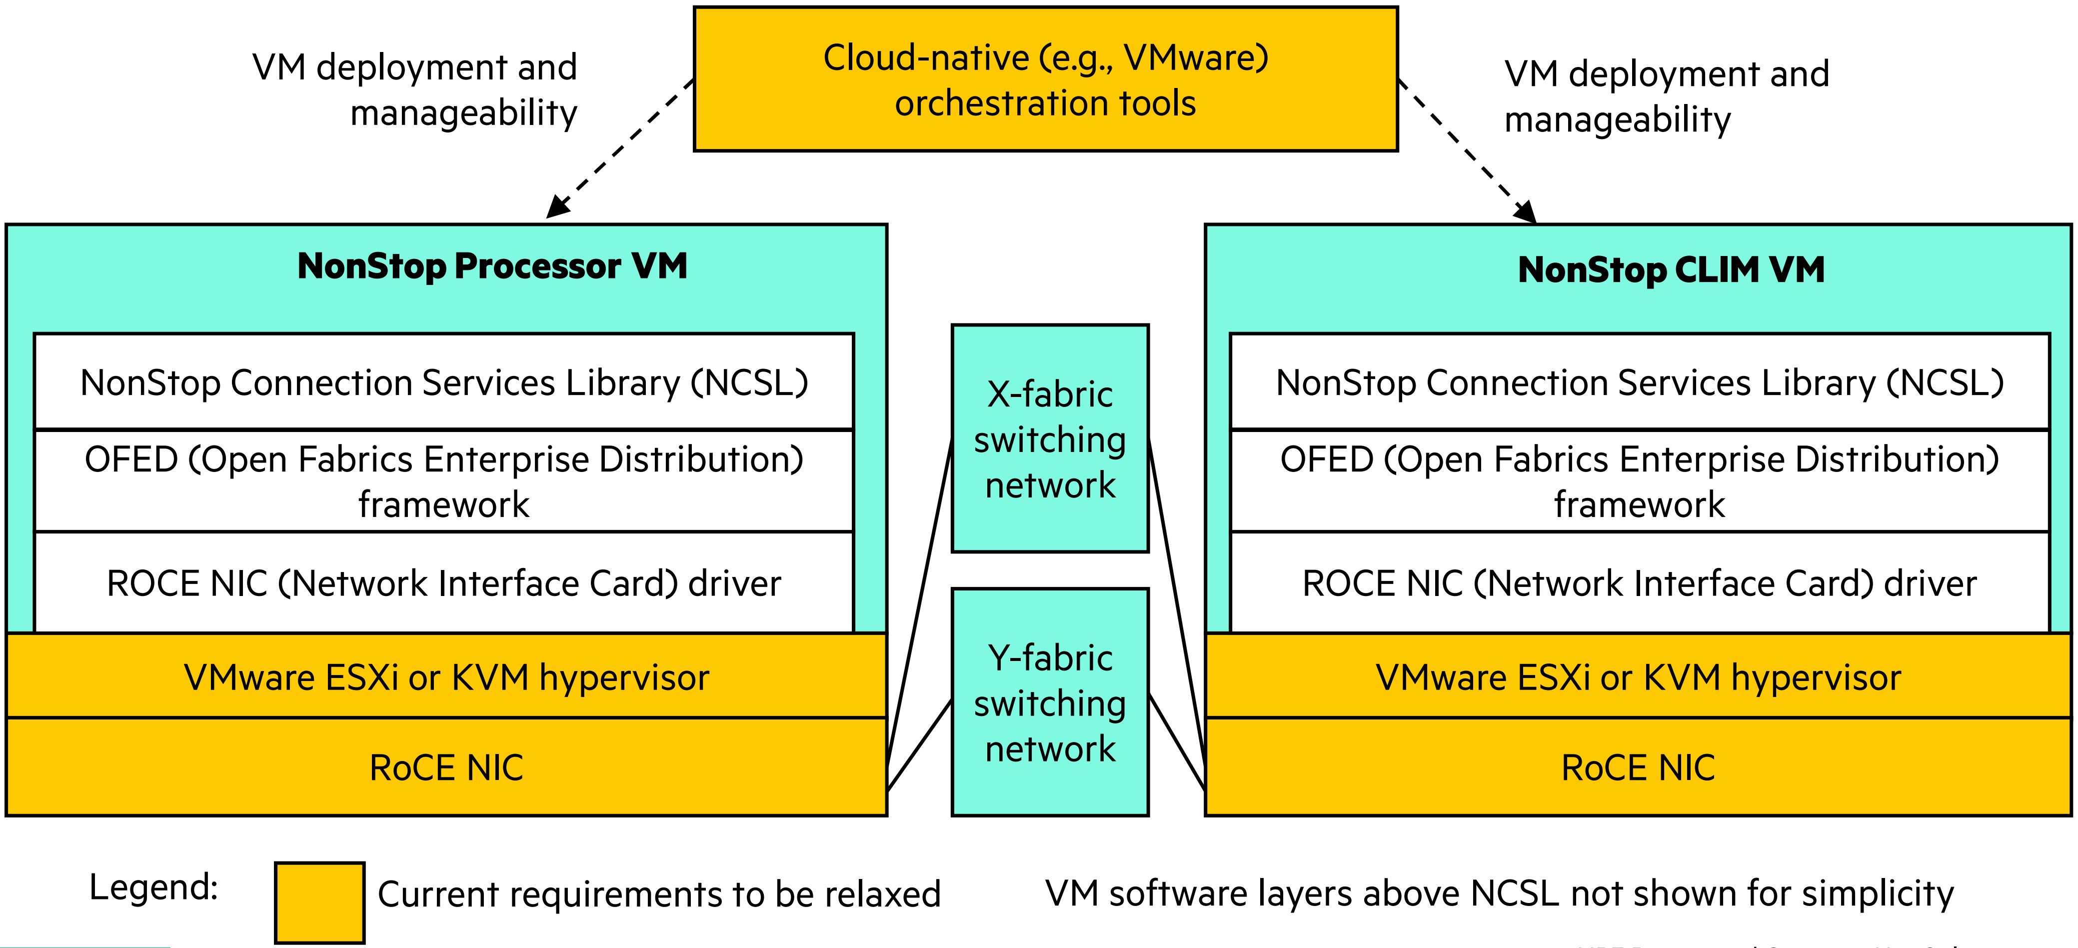

VM software layers above NCSL not shown for simplicity

vNS stack with Ethernet NICs

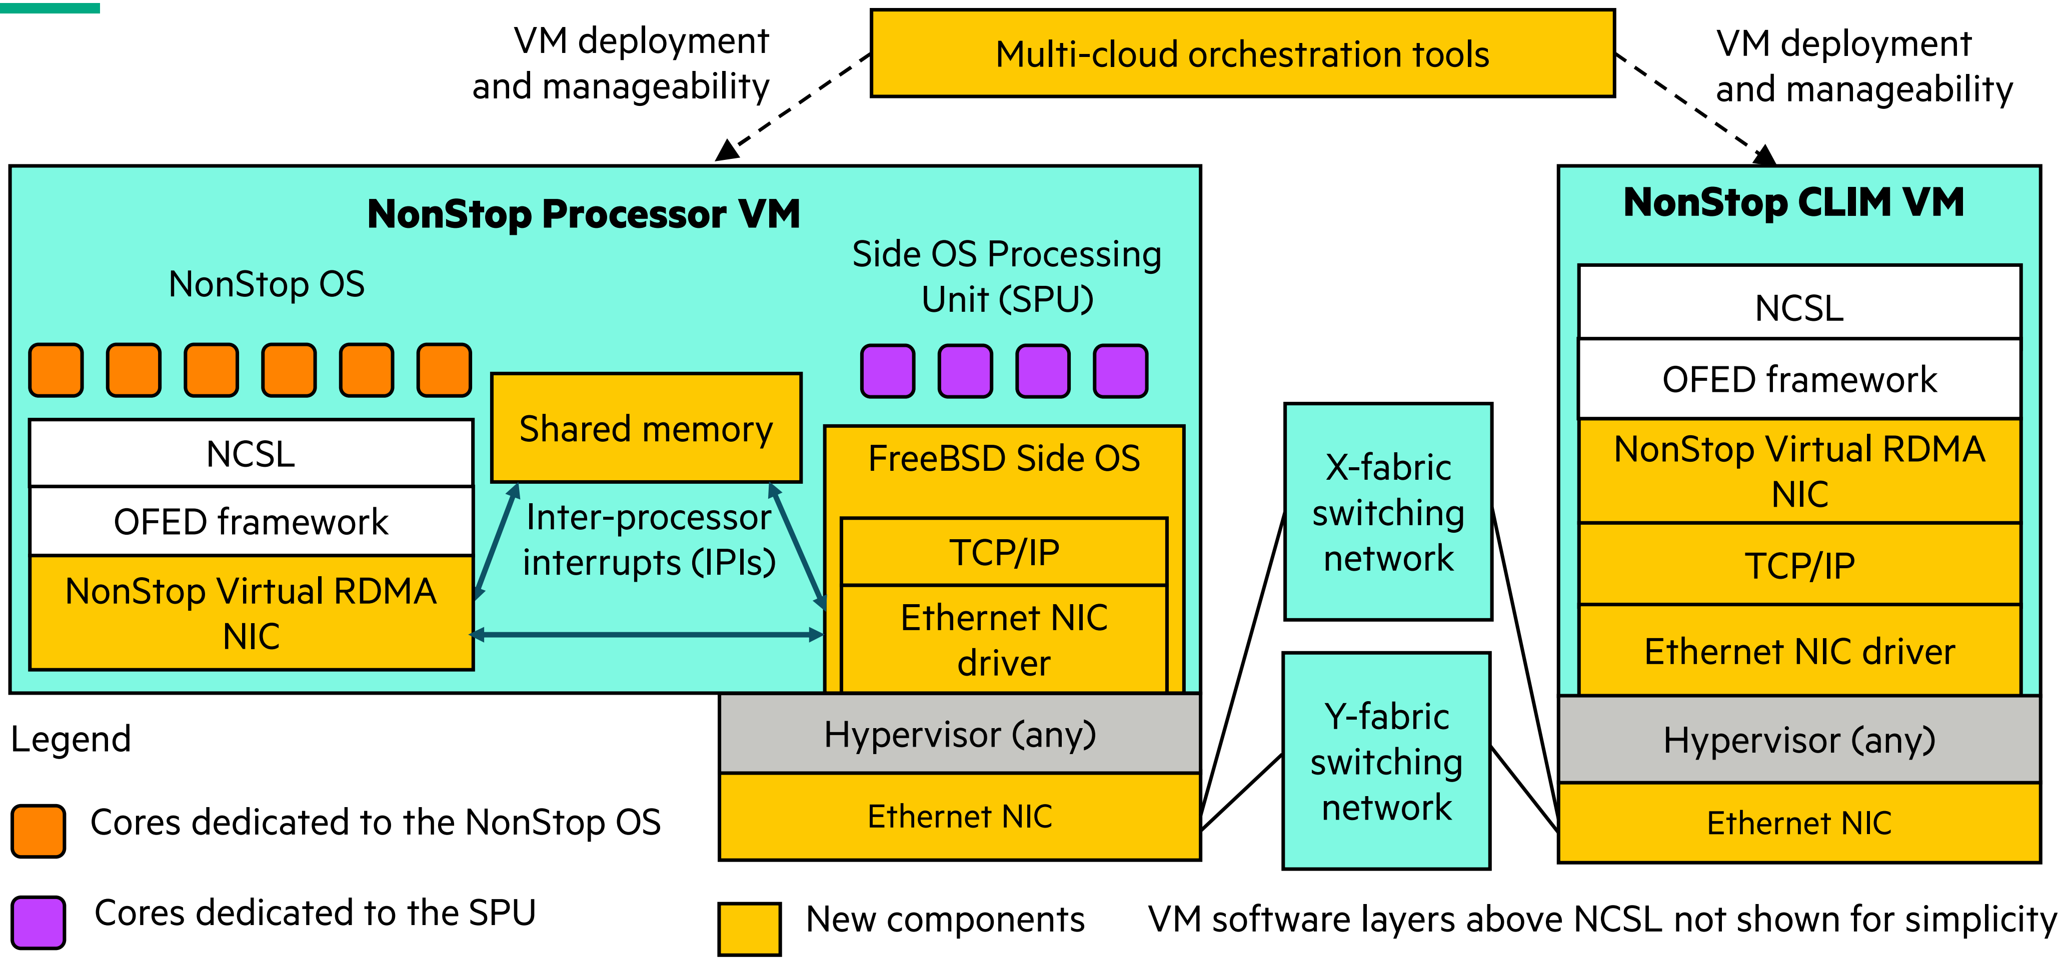

NonStop processor VM stack comparison

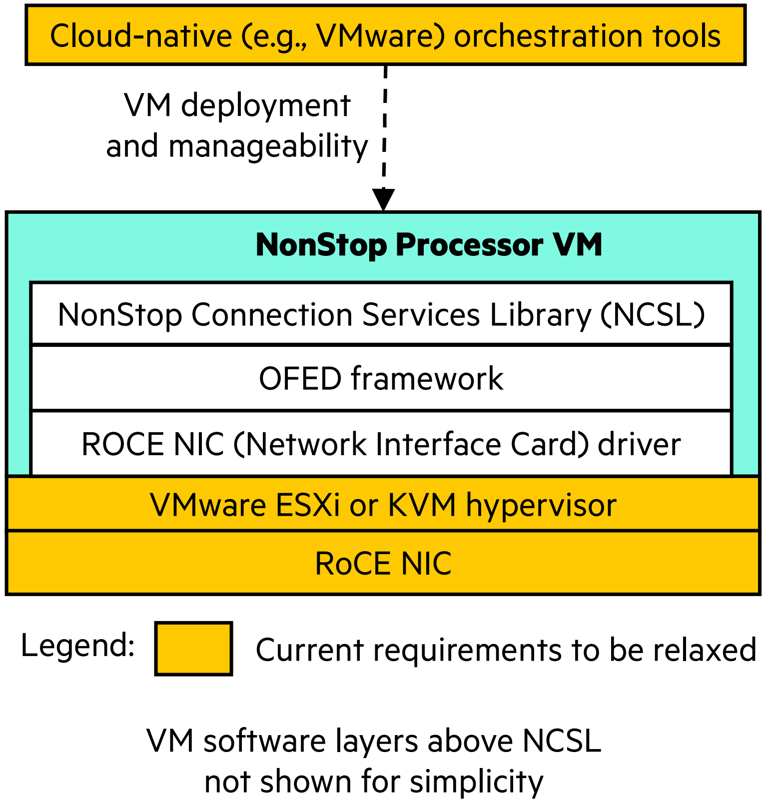

Current stack with RoCE NIC

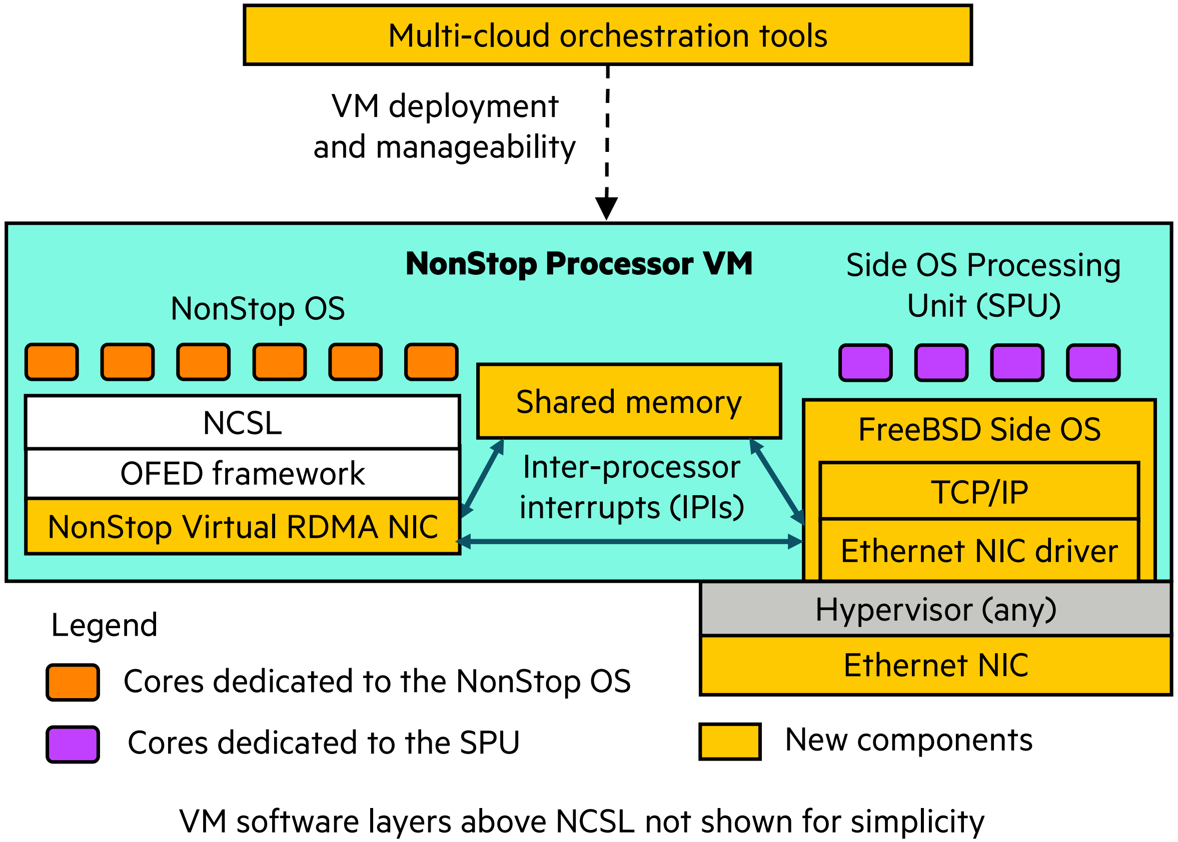

New stack with NVRDMA and Ethernet NIC

NonStop CLIM VM stack comparison

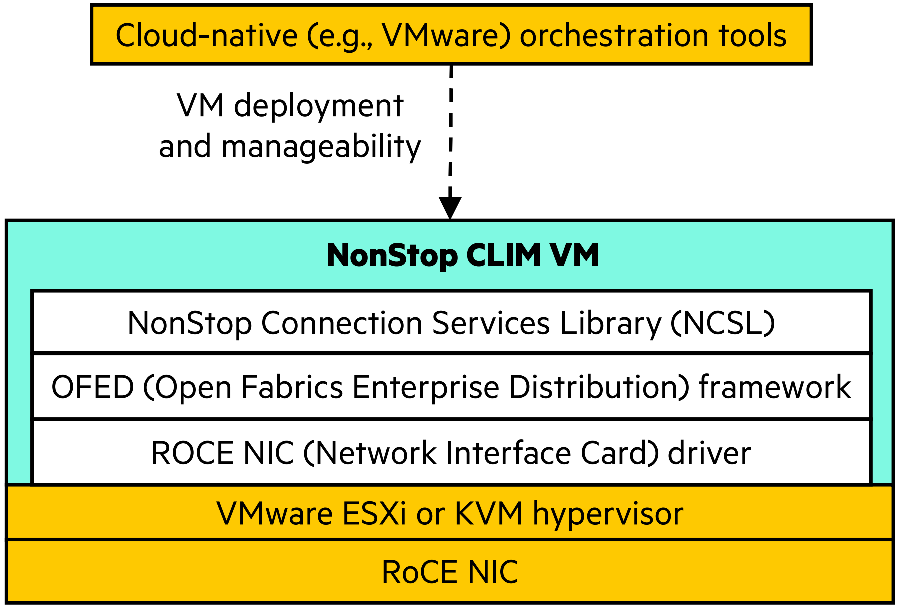

Legend: Current requirements to be relaxed

VM software layers above NCSL not shown for simplicity

Current stack with RoCE NIC

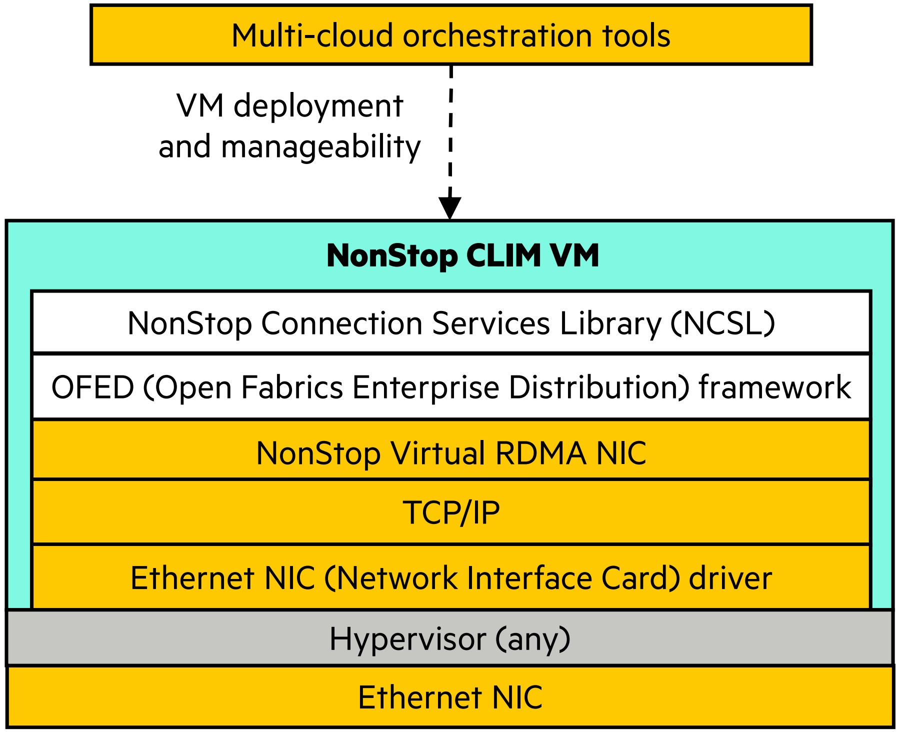

Legend: New components

VM software layers above NCSL not shown for simplicity

New stack with NVRDMA and EtherNet NIC

Key innovations in vNS stack with Ethernet NICs

- Goal: evolve vNS to a software stack that can be deployed on general fleet servers from all hyperscalers but which also preserves the existing interface for all vNS stack components above driver layers
- Note that general fleet servers can be dedicated to the vNS system and not be shared for other VMs
- Main components of our solution:

NonStop Virtual RDMA (NVRDMA)	• Removes RoCE NIC requirement – Runs over TCP/IP, standard Ethernet NICs, and all major hypervisors • Emulates RDMA (Remote Direct Memory Access) protocols in software • Superior performance compared to other competing approaches, namely Linux Soft-RoCE and VMware Paravirtual RDMA (PVRDMA)
Side OS Processing Units (SPUs)	• Ancillary FreeBSD OS runs side-by-side with the NSOS in NonStop processor VMs • TCP/IP stack and hypervisor interactions offloaded to the SPU for performance and hypervisor independence
Multi-cloud orchestration	• Deployment tools interface with the Terraform orchestration framework

TACL prompt achieved

- TACL prompt achieved on new vNS stack on Ethernet adapters
- This snapshot was taken while running a stress test on the system
- Additional milestones reached:
  - CPU reload
  - Latency performance tuning (first pass)

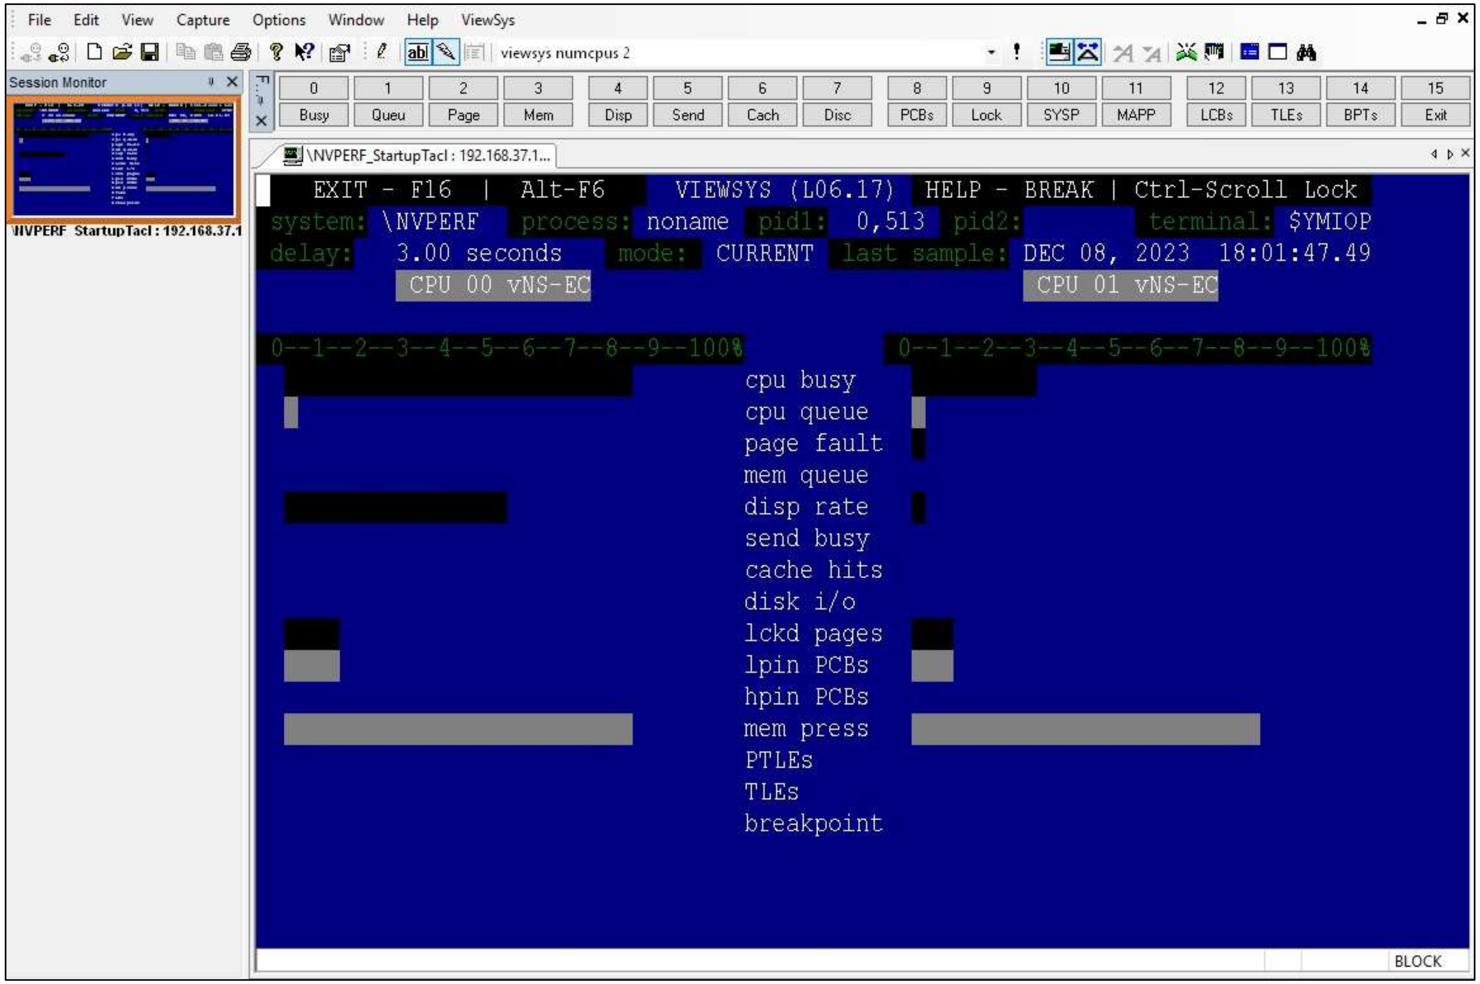

vNS solution with RoCE vs Ethernet fabric NICs

Fabric NICs	RoCE	Ethernet
Fabric switches	Ethernet	
Fabric link speeds	25, 40, 50, and 100 Gbps	
Fabric redundancy	X and Y fabrics	
RDMA implementation	Hardware	Software (NVRDMA)
Provides abstraction of InfiniBand protocols to NonStop Connection Services Layer (NCSL) and software layers above NCSL?	Yes, on both NonStop CPU and CLIM VMs	
Virtualization options for fabric interfaces	SR-IOV or PCI-PT	SR-IOV, PCI-PT, or virtualized interfaces such as VMXNET3
Virtualization options for CLIM network interfaces	SR-IOV, PCI-PT, or virtual interfaces (e.g., VMXNET3 and Virtio)	
Fabric NICs available in public cloud data centers?	Yes, with RoCE-based offerings from select hyperscalers in specific geographies	Yes – broad availability in general fleet servers
Supports any applications that run on L-Series without changes?	Yes	

NonStop Virtual RDMA (NVRDMA)

- NVRDMA performance tests indicate we have a viable solution
- Currently working towards running system-level benchmark tests with NVRDMA
- QA stress tests carried out on NonStop servers running did not find any adverse impacts on critical system operations, particularly time-sensitive messages and I/O operations

Current performance: NVRDMA on Linux @ 100 Gbps

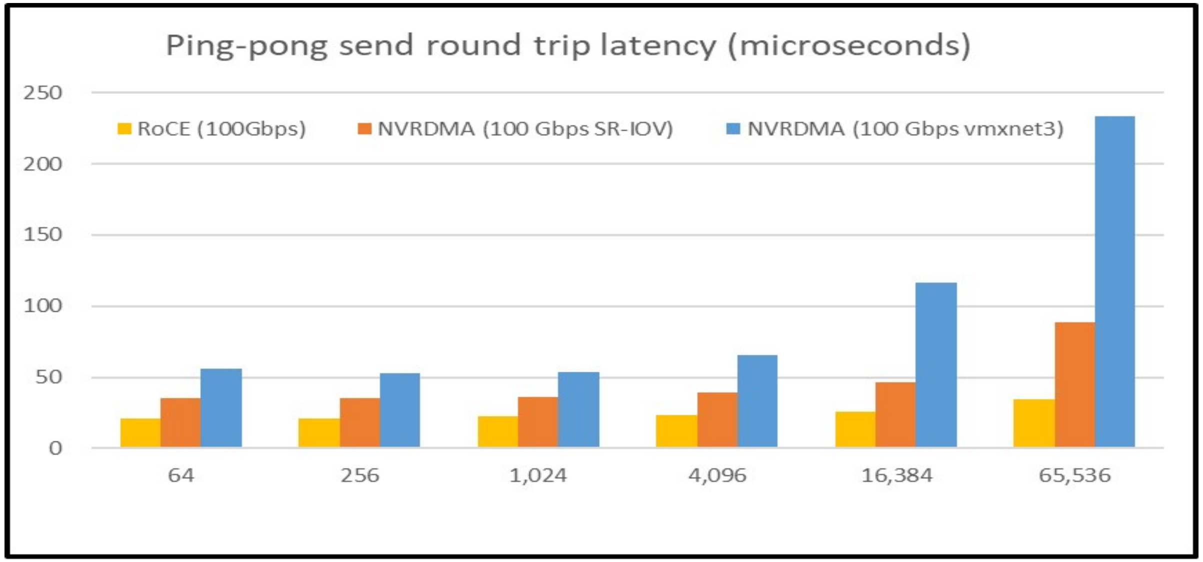

Data Size (bytes)	RoCE (100 Gbps)	NVRDMA using SR-IOV (100 Gbps)	NVRDMA using vmxnet3 (100 Gbps)
64	21	35	56
256	21	36	53
1,024	22	36	53
4,096	23	39	66
16,384	26	47	117
65,536	34	88	234

Notes:

- Graph shows Nvidia NICs using RoCE, and same Nvidia NICs using NVRDMA in both SR-IOV and vmxnet3 virtualization modes

Support for both RoCE and general-purpose Ethernet NICs

- Our goal in enhancing the NonStop fabric software stack is to support general-purpose Ethernet NICs in addition to (as opposed to in lieu of) RoCE NICs
- Customers aiming for the lowest possible latencies might want to consider a vNS solution with RoCE NICs

Support for hyperscaler native hypervisors

- A Virtualized NonStop system has virtual machine (VM) instances that run 3 different guest OSes:
  - NonStop CPU VMs run the NonStop OS (NSOS)
  - NonStop CLIM VMs run the HPDE (HPE Debian Enablement) Linux OS
  - The virtual NSC runs the Windows Server OS
- Similar to other Linux OS distributions, Debian can run as a guest OS on many hypervisors
- Likewise, the Windows Server OS is also widely supported on many hypervisors
- However, the NSOS can currently be deployed as a guest OS only on two hypervisors
  - VMware ESXi (hypervisor used in the current vNS solution)
  - Linux KVM (hypervisor used in the HPE Virtualized Converged NS2 system)
- NVRDMA and the SPU architecture will allow running the NSOS on hyperscaler native hypervisors including:
  - Hyper-V (Microsoft Azure hypervisor)
  - Google KVM (this is based on KVM but has security hardening enhancements implemented by Google)
  - AWS Nitro (also based on KVM)

Support for multi-cloud vNS deployment tools

- NonStop engineering has implemented environment-specific deployment tools for virtualized NonStop solutions in recent years. These include:
  - vNS Deployment Tools for VMware
  - vNS Deployment Tools for OpenStack
  - vmconfig (a script-based tool used in the HPE Virtualized Converged NS2 system)
- However, popular open-source frameworks have emerged to support multi-cloud orchestration software
- In particular, Terraform is supported by all major hyperscalers and is also available for on-premises private clouds such as VMware
- NonStop engineering has achieved significant progress towards multi-cloud vNS deployment tools that interface with Terraform
- These future multi-cloud vNS deployment tools will support vNS deployments in all major hyperscalers

Additional areas of investigation

Other investigation areas: Security model

- Security is a known concern for public cloud workloads
- The security model for vNS in the public clouds will need to consider:
  - Security hardening best practices for the vNS system
    - See section **Securing Virtualized NonStop (vNS) systems** in the *HPE NonStop Security Hardening Guide* for more details
  - Isolation of the cloud resources used by the vNS system including dedicated servers, storage, and networking
  - Security hardening best practices for external network traffic and system manageability access
  - Protecting the data
  - Compliance to industry security standards such as PCI-DSS, GDPR, HIPAA, etc.
- Shared resources such as ToR (Top of Rack) switches are a possibility
  - System interconnect traffic will use dedicated VLAN tags
  - Hardware encryption of system interconnect traffic is also a possibility with “smart” Ethernet NICs, but this will require further investigation

Other investigation areas: Support and billing models

Support models

- Possible options to be explored
  - Infrastructure support from the cloud vendor and NonStop support from HPE
  - HPE supporting the entire stack (infrastructure and NonStop)
  - Cloud vendor supporting the entire stack
- Other challenges to be investigated
  - Compliance with mission-critical SLAs
  - Synchronizing maintenance windows between cloud environment and NonStop

Billing models

- Possible options to be explored
  - Separate invoices – one from the cloud vendor and other from HPE
  - One consolidated invoice
- Other challenges
  - Real-time view of consumption data

Resources

- Product Website: **https://www.hpe.com/us/en/servers/nonstop.html**
- Technical Manuals (available at **www.hpe.com/info/nonstop-ldocs**)
  - *HPE Virtualized NonStop deployment and configuration guide for VMware*
  - *Hardware architecture guide for HPE Virtualized NonStop on VMware*

Thank you for attending this talk
HPE Virtualized NonStop Continues the Journey to the Cloud

For more information send email to:

P.Kamath@hpe.com
Lars.Plum@hpe.com
Arun.Fernandes@hpe.com

HPE Slides and Materials Usage

This content is protected