Hewlett Packard Enterprise

NonStop Technical Boot Camp 2023 TBC23-TB58 Best Practices for Configuring and Managing HPE Virtualized NonStop Systems for Mission Critical Workloads

Marcelo de Azevedo, Bryce Kosinski, Lars Plum, Mark Thompson, and John Zimsky

September 2023

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 overview
- HPE Virtualized NonStop adoption and hardware environments chosen by customers
- Best practices for planning and configuring vNS systems
- VMware patching considerations
- Considerations for upgrading vNS systems
- ESXi host upgrade methods

HPE Virtualized NonStop overview

VMware Cloud for HPE Virtualized NonStop

VMware vCenter Server®

VMware ESXi™ hosts (compute node(s))

VMware vRealize® Orchestrator

HPE NonStop VMs (vNS CPUs, vCLIMs, vNSC) deployed in ESXi hosts

VMware vSphere® Web Client

Datastores

Storage for vNS (network storage or drives in compute nodes)

25, 40, 50, or 100 GbE fabric switches

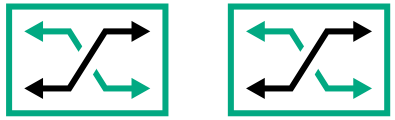

10 or 25 GbE storage fabric switches (optional)

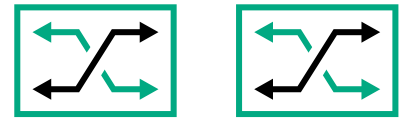

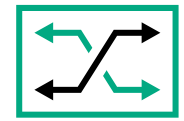

1 GbE maintenance LAN switch

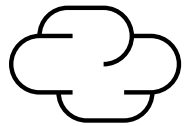

Notes:

- VMware vCenter Server and VMware vRealize Orchestrator can be deployed in a dedicated ESXi host (recommended), or they can run in VMs deployed in the same ESXi hosts where the HPE NonStop VMs run
- vNS does not require a dedicated vCenter. It can be integrated into an existing vSphere environment managed under vCenter

Resources required to deploy a vNS system

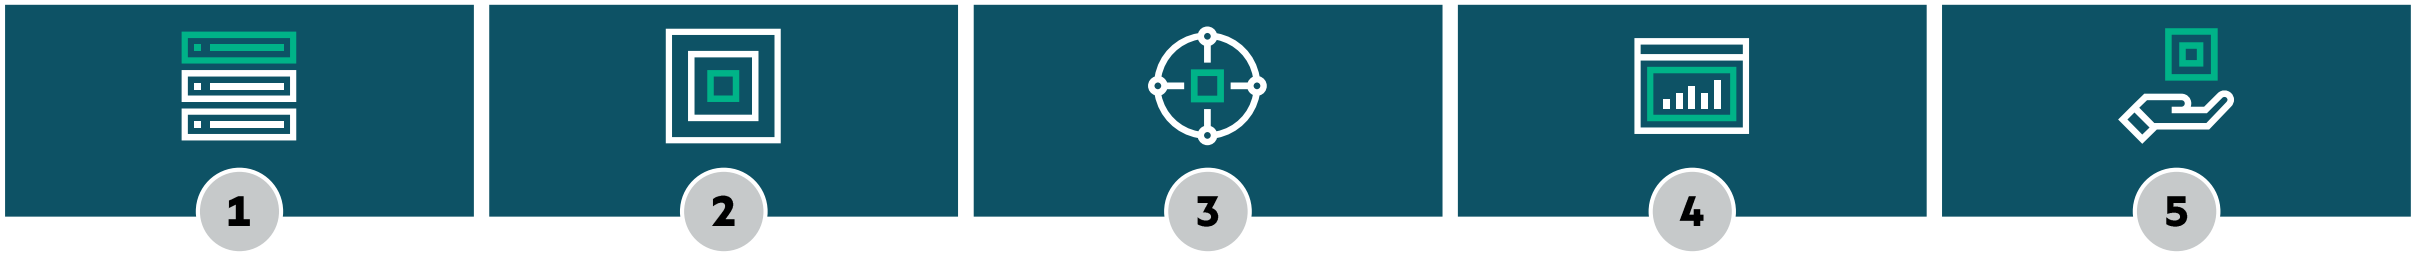

1	2	3	4	5
Servers	**Storage**	**Networking**	**Software**	**Services**

• Based on these Intel® Xeon® x86 processors:
  • Minimum: E5-2600 v3 family
  • Latest supported: 3rd Gen Intel Xeon Scalable processors
	• Block level storage
	• External network storage or internal server storage
	• Use HPE NonStop's storage mirroring design or vSAN RAID1 mirroring
	• No sharing of storage hardware between primary and mirrored volumes
		• System interconnect: redundant Ethernet switches with minimum 25GbE ports
		• System interconnect NICs: Mellanox ConnectX®-6 VPI or EN (100 GbE), or ConnectX®-4 Lx VPI or EN (25 GbE)
		• IP vCLIM: 10 GbE NICs based on specific Intel® or Broadcom® processors
			• vSphere Enterprise Plus Edition (6.5 or above)
			• vCenter and vRO (not dedicated)
			• Windows Servers 2022 Standard Edition
			• HPE NonStop L-series SUT
			• vCLIM
			• vNSC
			• Virtual BackBox (for BR)
				• Hardware installation & startup
				• VMware installation & startup
				• vNS installation & startup

Refer to the **Hardware architecture guide for HPE Virtualized NonStop on VMware** for more information

HPE Virtualized NonStop adoption and hardware environments chosen by customers

HPE Virtualized NonStop adoption

- Several in production, more are getting ready
- 75% on HPE hardware
- Several customer and partner testimonials in the public domain
- Published case studies and testimonials: Dell, EuroClear, Sumitomo Mitsui Card, Red Link S.A., and 4Tech Software

HPE-IT business applications running on vNS

- A total of 13 vNS systems
- Twelve ESXi hosts for the production VMware environment +
- Twelve ESXi hosts for the DR VMware environment
  - One rack at each data center
  - Major reduction in footprints
  - Reduced electrical feeds
  - Multiple vNS systems share the ESXi hosts
  - For example, ESXI hosts 1 – 6 are shared by two vNS systems
  - For those six ESXi hosts, HPE-IT assigned NonStop systems with lower workloads
- Eight ESXi hosts for development and QA VMware environment
  - One rack at development data center
- Add a full DR environment for all production systems
  - Shadowbase
  - AutoSYNC
- NonStop systems reduced from ten to seven
  - Added one more development vNS system
- Maximum vNS processor memory
- Two IP vCLIMs per vNS system
- Six Storage vCLIMs per production vNS system
  - Lots of archival NonStop DP2 volumes
  - Continued use of SMF
- One pair of vNSC for all vNS systems
- Retain Expand-over-IP configuration

Hardware environments

Type	What is deployed
System configurations	1c, 2c, 4c, 6c; 2 to 8 CPUs
Hardware vendor	HPE (75%), non-HPE (25%)
IT design	Private cloud
Server categories	Rackmount Blades
Storage	Internal Hyperconverged External: iSCSI and FC
Ethernet fabric switches	HPE StoreFabric M-series HPE FlexFabric Cisco Dell
Maintenance LAN switch	HPE Aruba Cisco Dell

Best practices for planning and configuring vNS systems

Best practices for planning and configuring vNS systems

- The following technical manuals (available at **www.hpe.com/info/nonstop-ldocs**) cover planning and configuration aspects of vNS systems in detail
  - *HPE Virtualized NonStop deployment and configuration guide for VMware*
  - *Hardware architecture guide for HPE Virtualized NonStop on VMware*
- The above manuals have been greatly enhanced in recent years
  - HPE NonStop ISV partners have reported being able to plan, configure and deploy vNS systems following instructions on these manuals
  - However, HPE Services also offers vNS planning, configuration, and deployment services
- This presentation is not a substitute for the above referenced manuals
- However, this presentation summarizes some of the key considerations to plan, configure, and maintain vNS systems

Best practices for planning vNS systems: Key stake holders

- Identify the key stake holders early on:
  - Business groups owning the workload
  - Cloud IT groups
    - VMware, network, and storage architects
    - Security experts
    - Administrators and support personnel
  - NonStop team
    - Developers
    - System administrators
- Discuss vNS architecture and application requirements with all teams
- Discuss VMware monitoring, alerts, and patching policies with VMware and security teams
- Consider enrolling key stake holders in the HPE Virtualized NonStop on VMware – Architecture, Installation and Configuration class offered by NonStop Education (for more details, visit https://www.nonstop-academy.com/)

Best practices for planning vNS systems: Storage options

Storage option for DP2 volumes	Considerations
Internal drives in ESXi hosts without vSAN	• DP2 disk revives required after loss of an SCLIM or ESXi host due to planned maintenance or failure • Rolling upgrades possible by bringing down one ESXi host at time • Configuration supported: 2c-2d with 2 paths to storage
External storage arrays 	• DP2 disk revives not required after loss of an SCLIM or ESXi host • Rolling upgrades possible by bringing down one ESXi host at time or migrating vNS VMs to a spare • Configurations supported: 2c-2d and 4c-2d with 4 paths to storage • Storage arrays have built-in fault tolerance, but two arrays provide another level of fault tolerance by allowing us to put primary and mirror drives on separate arrays
VMware vSAN	• DP2 disk revives not required after loss of an SCLIM or ESXi host • However, if an ESXi host in a vSAN cluster is lost and replaced the vSAN RAID array must be rebuilt • Rolling upgrades possible by bringing down one ESXi host at time or migrating vNS VMs to a spare • Configurations tested so far: 2c-1d and 2c-2d with a single RAID-1 VMware vSAN datastore

- For more details, see session TBC23-TB57 HPE Virtualized NonStop Storage

Best practices for planning vNS systems: VM image file placement

Storage option for DP2 volumes	VM image files placed on ESXi host boot drives	VM image files placed on a datastore other than ESXi host boot drives
Internal drives in ESXi hosts without vSAN	• Either option offers similar fault isolation • Loss of access to HSS.iso or CLIM OS VMDK image impacts only VMs in the same ESXi host • Rolling upgrades not a consideration (not supported in this configuration)	
External storage arrays	• Optimal option for fault isolation • Loss of access to HSS.iso or CLIM OS VMDK image impacts only VMs in the same ESXi host • Copying vCLIM OS VMDKs can take up to 30 min on hosts with HDDs. This can be reduced by using SSD/NVMe as OS disks • This is not a concern for vCPUs, because HSS.iso files are small	• Non-optimal for fault isolation – loss of an external storage array impacts all VMs with image files on that array • Store VMDKs of vCLIM failover pairs in separate fault zones for availability • Faster rolling upgrades (VM image files need not be copied elsewhere)
VMware vSAN	• Same considerations listed above	• Similar to above, but the fault isolation concern does not apply (since vSAN is a cluster of ESXi hosts)

Best practices for planning vNS systems

Use `faultZone` option 2 (default) for system deployment

`faultZone` option	2 (default)	1	0
Fault-tolerant rules are enforced on CPUs	✓	✓	No
Fault-tolerant rules are enforced on CLIMs	✓	No	No
Fault-tolerant rules are enforced on NSK volumes	✓	No	No
Each CPU VM is required to run in a separate host	✓	✓	No
The number of CPU VMs must be an even number	✓	✓	No
CLIM VMs of the same type must run in different hosts	✓	No	No
NSK volume is required to have both `primaryClim` and `mirrorClim` to provide both paths to the volume	✓	No	No

HPE Virtualized NonStop security

Security of HPE NonStop software

- Nothing changes!!
  - All security features are available
  - All security practices are applicable
- OS security features are intact
- PCI related practices continue to be applicable
- Secure the stack referring to the *HPE NonStop Security Hardening Guide*

Security of the platform

- Responsibility of the private cloud
- Follow VMware and hardware vendor's security guidelines
- Best practices:
  - Deploy HPE NonStop in the most secure area of the cloud with both physical and logical security
  - Physically isolate the switches used for the fabric
  - Secure and protect HPE NonStop volumes
  - See section **Securing Virtualized NonStop (vNS) systems** in the *HPE NonStop Security Hardening Guide* for more details

Best practices for planning vNS systems: How many ESXi hosts?

- The *Hardware architecture guide for HPE Virtualized NonStop on VMware* provides detailed information on how to plan for the ESXi hosts in which the vNS VMs will be deployed
- Deploying a combination of a vCPU, Storage vCLIM, and/or an IP vCLIM per ESXi hosts is a commonly used approach to achieve a fault-tolerant configuration
  - However, see `faultZone` option 2 discussed earlier for more details
- It is possible that at least one of the ESXi hosts has enough cores and memory to deploy vNS VMs (vCPU, vCLIMs) and support appliances (vNSC, vBB, vCenter, and vRO)
- However, it is generally preferable to deploy the support appliances (vNSC, vBB, vCenter, and vRO) in a separate ESXi host
- This allows the ESXi hypervisor in that host, as well as the VMware vCenter and vRO VMs, to be updated separately before a planned ESXi upgrade of the underlying VMware environment
- Consider also planning for a spare ESXi host for rolling upgrades (discussed in more detail later)
- Create diagrams, spreadsheets, and/or documents depicting the planned system architecture:
  - Mapping of the various VMs to the underlying ESXi hosts
  - Network architecture of the system
  - Storage mappings (DP2 volumes to LUNs, datastores, and vCLIMs)

Virtualized NonStop on VMware

Sample architectural diagram

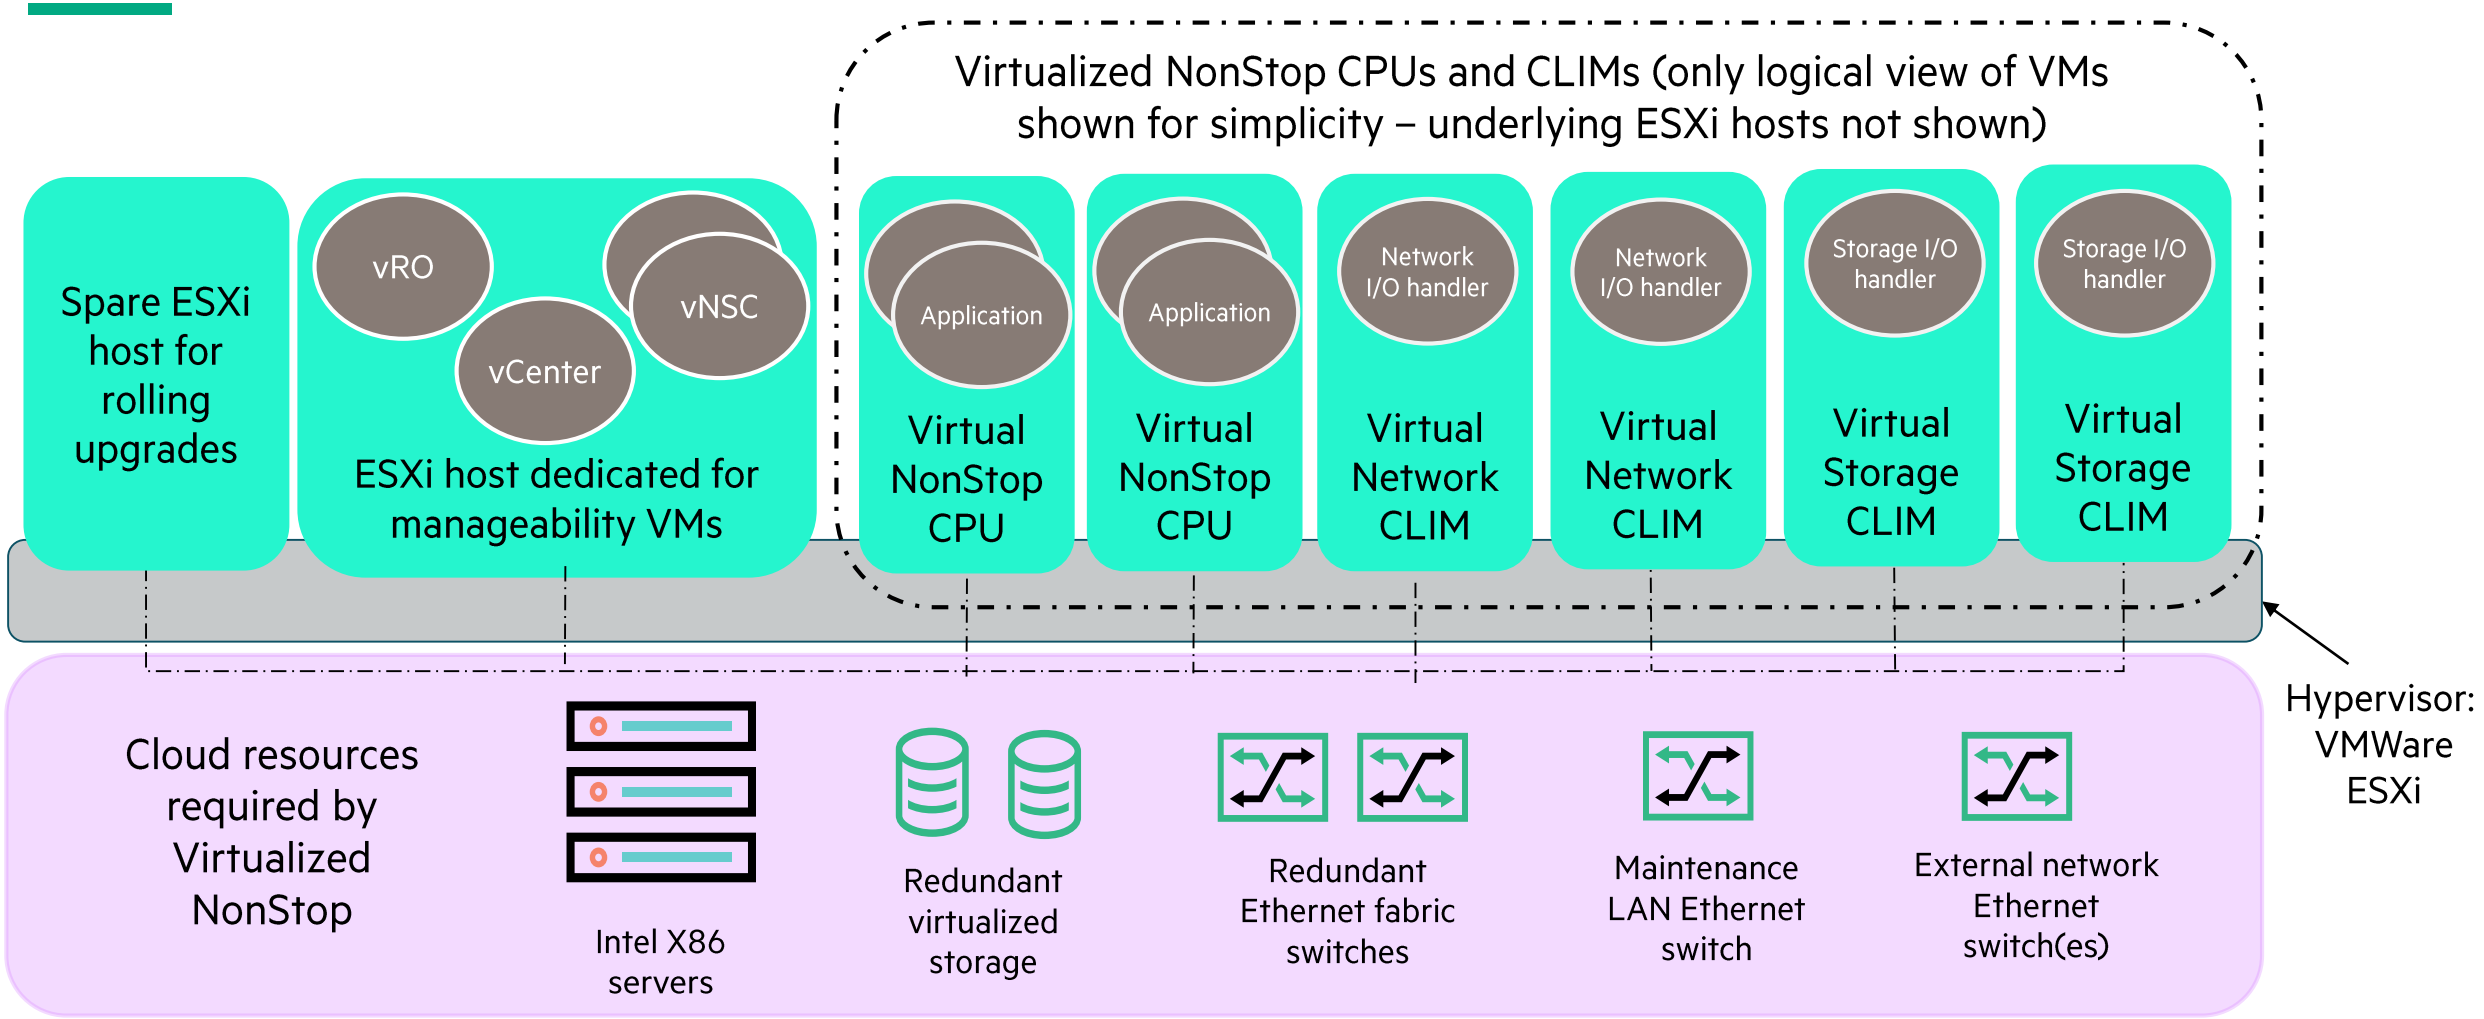

New NonStop $SYSTEM disk snapshot and restore feature

- New feature released with L23.08. Relevant SPRs:
  - T0853L03^DCR CLIM DVD Software
  - T0682L02^BBN OSM Service Connection Suite
  - T0634H06^ACY OSM Console Tools
  - T0954V04^AAV ETI-NET BackBox H4.11 (required only if backup destination is tape)
  - T0964V01^AAJ VTR R1.05 (required only if backup destination is tape)
- Minimum required RVU: L21.06 plus the SPRs above
- Supported HPE NonStop systems:
  - HPE Virtualized NonStop
  - Converged HPE NonStop NS4 X4, NS8 X4, NS7 X2 & X3, and NS3 X2 & X3
  - Virtualized Converged NonStop NS2 X2 & X3
- Requirements:
  - Either disk or tape devices can be configured as the destination for $SYSTEM disk snapshots
  - The destination disks or tapes used must be connected to the same storage CLIM pair as $SYSTEM is connected
  - The size of the destination disks or tapes must be equal to or greater than those of the $SYSTEM volume

Automatic TCP/IP Connection Failover on vNS

- The L23.08 RVU introduces Automatic TCP/IP connection failover for IPv4 connections on HPE Virtualized NonStop vCLIMs
- This feature was originally introduced in L19.03 RVU with Gen9 and Gen10 IP and Telco CLIMs on HPE NonStop X platforms
- As of L23.08 the same feature is supported on HPE Virtualized NonStop systems
- For more information, see the *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- Automatic TCP/IP Connection Failover can be configured on PCI pass-through, SR-IOV, or VMXNET3 network interfaces. Considerations:
  - VMXNET3 interfaces do not have visibility to physical link pulse loss
  - Thus, automatic failover of a VMXNET3 network interface upon physical link pulse loss is not supported
  - However, automatic failover of a VMXNET3 network interface is supported in case of a vCLIM failure or interface failure
  - vCLIM failover pairs (primary and backup vCLIMs) in the same vNS system must be deployed in different ESXi hosts for fault tolerance

VMware patching considerations

VMware patching considerations

- At first glance, VMware releases ESXi patches fairly frequently. What does this mean for vNS?

ESXi version	Initial release date	Number of patches released	Number of patches recalled	Average interval between patches with bug fixes, security fixes, or both (days)	Average interval between patches with security fixes (days)
ESXi 7.0 U3	10/5/2021	14	3	49	108
ESXi 7.0 U2	3/9/2021	5	0	86	172
ESXi 7.0 U1	10/6/2020	6	0	99	124
ESXi 7.0	4/2/2020	2	0	82	82

- The data above is based on ESXi release notes posted on https://docs.vmware.com/en/VMware-vSphere/index.html as of September 2, 2023
- All ESXi patches released up to the above date require host reboots and VM migration or shutdown
- See also https://kb.vmware.com/s/article/2143832 for a summary of ESXi versions and build numbers

When to patch VMware?

- Just as any other software, VMware releases patches periodically to address functional and security defects in the product
- In the last 12 months, VMware has released 8 patches for ESXi 7.0 (Source: https://kb.vmware.com/s/article/2143832)
- NonStop also releases patches for its software stack as SPRs
- Customers install patches after reviewing its content and determining the benefit the patch offers to their environment
- According to VMware (Source: https://www.vmware.com/security/advisories.html#severity=Critical&severity=Important), on ESXi, there was one “Critical” and two “Important” patches released over the past 12 months
  - The “Critical” patch (https://www.vmware.com/security/advisories/VMSA-2022-0033.html) was a vulnerability in the USB 2.0 controller which vNS VMs do not use
  - Hence a dedicated set of hosts running vNS VMs and running as an isolated cluster in a VMware environment, may not have been impacted by the vulnerabilities being addressed by this patch
  - The “Important” patches had fixes for CVEs of severity 7.2 and below

When to patch VMware? (continued)

- While the past trend is no guarantee of the future, there is no data to suggest that vNS requires frequent VMware ESXi patching
- A properly isolated VMware environment running vNS VMs will require ESXi upgrades at the time of the RVU upgrades and upon availability of critical patches addressing CVEs impacting that environment
- The available data does not suggest that the latter is very frequent
  - A cursory sampling of customers using vNS for production applications show VMware being patched between once a year to once every 2 years
- In addition to VMware ESXi patching, RVU upgrades also provide an opportunity for SPP (Service Pack for Proliant) upgrades on vNS systems deployed on HPE Proliant servers

Considerations for upgrading vNS systems

vNS availability for various upgrade scenarios (Note: see caveats in next slide)

Upgrade scenario	Requires system to be brought down to perform the upgrade	Upgrade can be done in a non-disruptive manner due to redundant components	Can be performed by upgrading one host at a time, without bringing the vNS system down	Can be performed without disrupting vNS (VMware vCenter, vRO and VMware Tools deployed on a separate host)
L-Series RVU upgrade	•			
Applying any SPR(s) that require a system load	•			
Host firmware upgrades that require a host reboot			•	
Host hardware upgrades that require host downtime			•	
Storage array firmware upgrade		•		
Storage array hardware upgrade		•		
Fabric switch firmware upgrade		•		
Fabric switch hardware upgrade		•		
VMware ESXi hypervisor upgrade			•	
VMware vCenter upgrade				•
VMware vRealize Orchestrator (vRO) upgrade				•
Upgrade of other VMware components besides the ESXi hypervisor, such as VMware Tools, etc.				•

vNS availability for various upgrade scenarios: Caveats

- A storage array hardware upgrade that takes the array offline entails a full outage in a configuration with a single storage array holding all the NSK volumes
  - However, a storage array firmware upgrade can be tolerated with a single array with redundant storage controllers by upgrading the storage controller firmware one controller at a time
- In a configuration with separate arrays for the primary and mirror NSK volumes, a storage array hardware upgrade that takes an array offline will require revives of the drives
  - FCHECK scans may also be required against the array drives prior to taking the array offline
- If CLIM VMDK images were placed on a storage array, an upgrade that takes the array offline will impact all the CLIMs defined on the array
  - This would lead to an outage if primary IP connections don't use automatic failover
- Likewise, upgrade scenarios that take an ESXi host offline require primary IP connections configured with automatic failover

General considerations when upgrading vNS systems

- Consider adding a spare ESXi host to each vNS system to perform rolling upgrades
- Consult the information below when planning a vNS system upgrade:

Information	Where to find:
VMware versions supported for vNS	**RVU, vSphere, and vCenter support matrix** section in the *HPE Virtualized NonStop Deployment and Configuration Guide for VMware*
ESXi custom images for HPE Proliant servers for various VMware versions	https://www.hpe.com/us/en/servers/hpe-esxi.html
Mapping between ESXi custom images and Service Pack for Proliant (SPP) versions	http://vibsdepot.hpe.com/mapping/SPP-HPE_Custom-Image-vibsdepot-mapping-Gen9-later.pdf

ESXi host upgrade methods

Down system upgrade

Scenarios without a spare

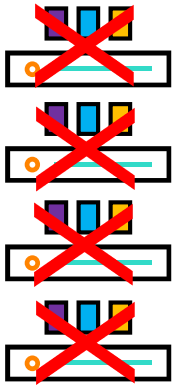

Bring system down & upgrade all hosts

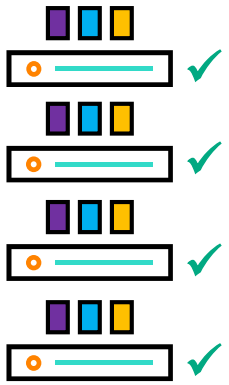

Bring system back up

- Down system upgrade (upgrade all hosts in parallel)
- Fastest method, but no pre-check that host upgrade is successful (thus **not recommended**)

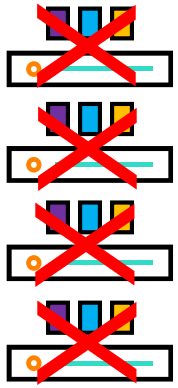

Bring system down & upgrade one host

Bring upgraded host back up

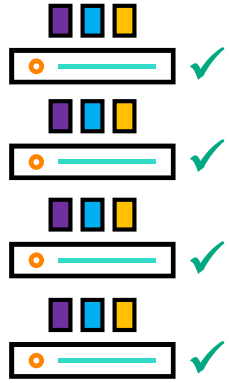

Upgrade remaining hosts and bring system back up

- Down system upgrade, do one host first to confirm upgrade is successful, then proceed to other hosts
- Second fastest method, but requires using one of the production hosts to pre-check that host upgrade is successful
- This method is actually used frequently to upgrade QA and development lab systems owned by HPE NonStop Engineering

Down system upgrade

Scenario with a spare

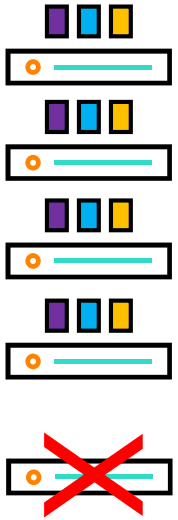

Bring spare down & upgrade spare

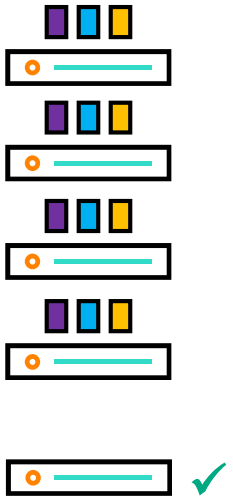

Bring spare back up

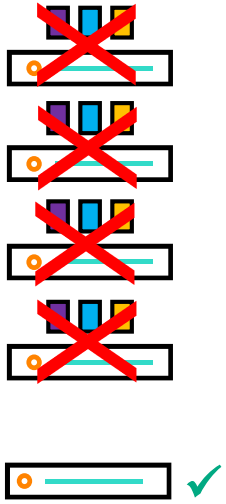

Bring system down & upgrade hosts

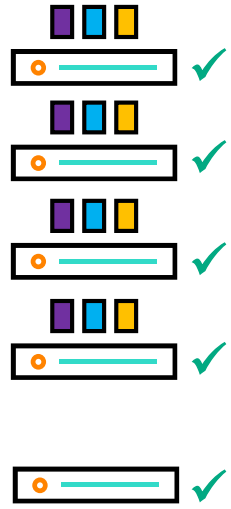

Bring system back up

- Down system upgrade, do spare host first to confirm upgrade is successful, then proceed to other hosts
- Second fastest method, uses a spare host to pre-check that host upgrade is successful

Rolling upgrade

Scenario without a spare

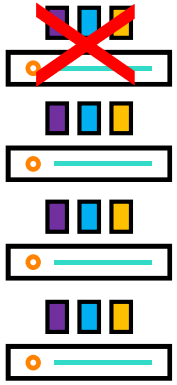

Bring host 1 down
& upgrade host 1

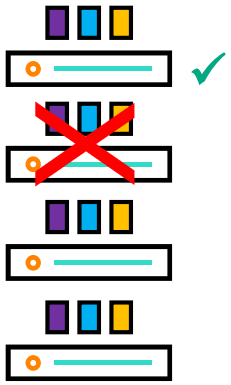

Bring host 1 back up

Bring host 2 down
& upgrade host 2

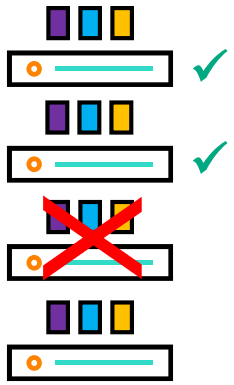

Bring host 2 back up

Bring host 3 down
& upgrade host 3

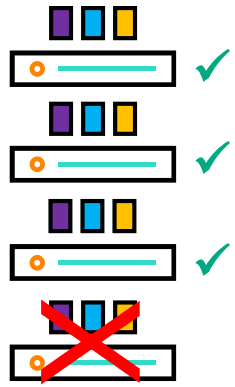

Bring host 3 back up

Bring host 4 down
& upgrade host 4

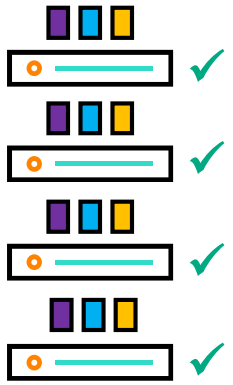

Bring host 4 back up

- Takes longer than a parallel upgrade, but vNS system stays up
- No pre-check that host upgrade is successful
- This method is used for production vNS systems in the field

Rolling upgrade

Scenario with one spare + VMs moved to the “previous host” after each ESXi host is upgraded

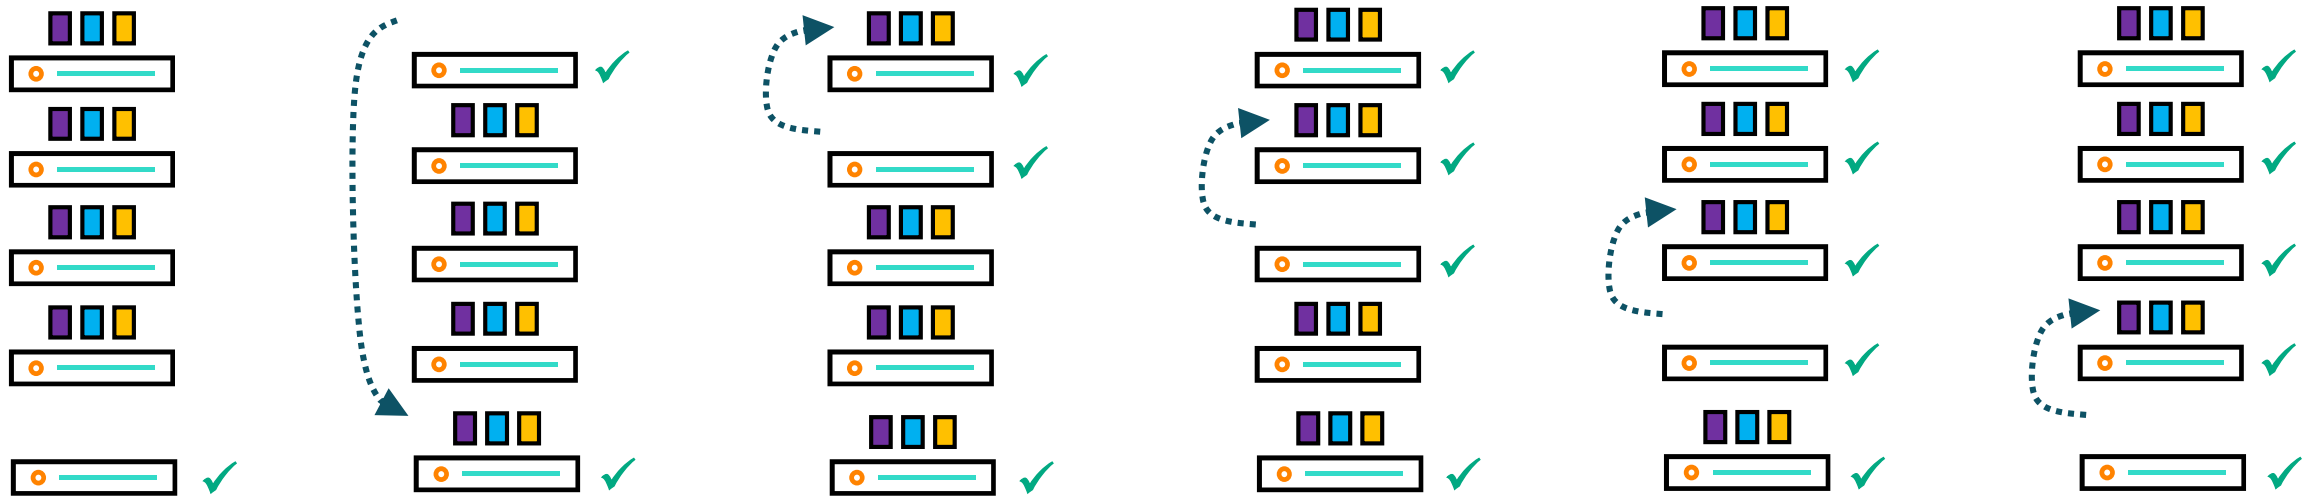

Upgrade spare host

Move host 1 VMs to spare & upgrade host 1

Move host 2 VMs to host 1 & upgrade host 2

Move host 3 VMs to host 2 & upgrade host 3

Move host 4 VMs to host 3 & upgrade host 4

Move VMs from spare to host 4

- Takes longer than a parallel upgrade, but vNS system stays up
- Uses a spare host to pre-check that host upgrade is successful
- Cabling changes may be required for certain vNS system configurations, making this method disadvantageous in such cases

Rolling upgrade

Scenario with one spare + VMs return to original location after each ESXi host is upgraded

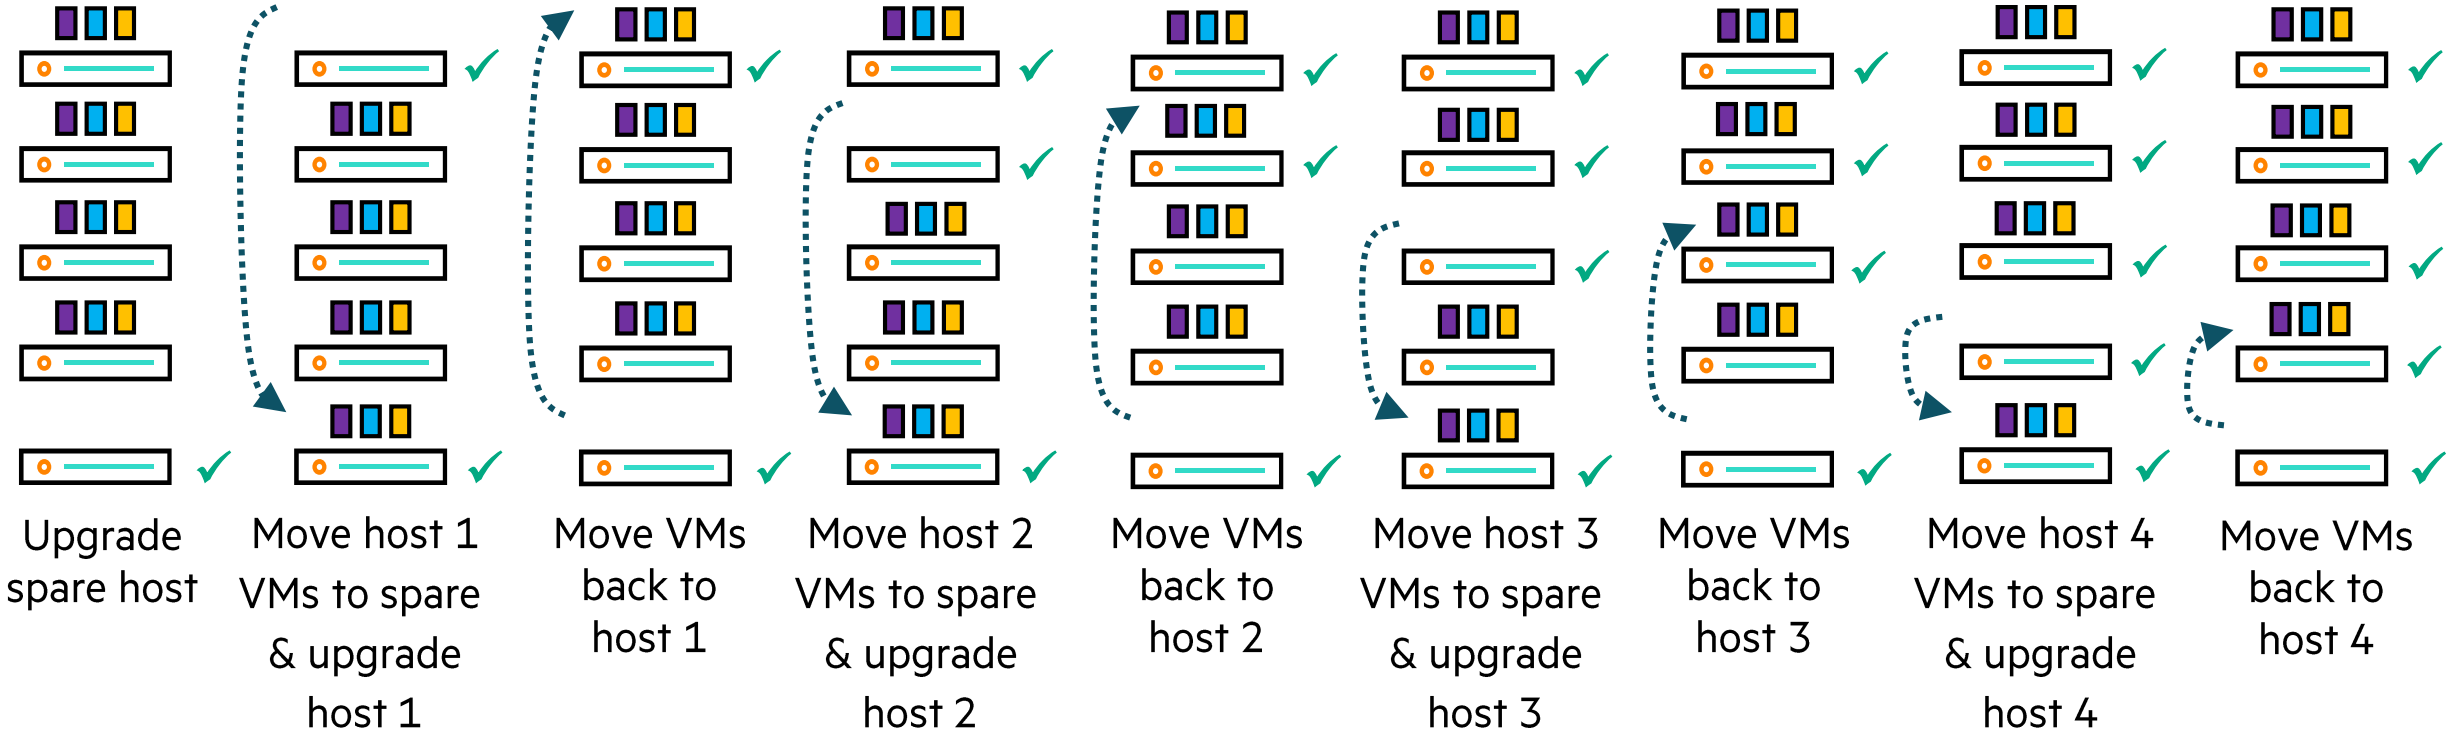

- Slowest method, but vNS system stays up
- Uses a spare host to pre-check that host upgrade is successful
- This method is used for production vNS systems in the field
- Cabling changes not required

Summary of ESXi host upgrade methods

Upgrade method	Down system upgrade?	Rolling upgrade?	Spare host required?	Considerations
Down system upgrade, upgrade all hosts in parallel	Yes	No	No	Fastest method, but no pre-check that host upgrade is successful (**not recommended**)
Down system upgrade, do one host first to confirm upgrade is successful, then proceed to other hosts				Second fastest method, but requires using one of the production hosts to pre-check that host upgrade is successful
Down system upgrade, do spare host first to confirm upgrade is successful, then proceed to other hosts			Yes	Second fastest method, uses a spare host to pre-check that host upgrade is successful
Rolling upgrade without spare host	No	Yes	No	Takes longer than a parallel upgrade, but vNS system stays up. No pre-check that host upgrade is successful
Rolling upgrade with spare host + VMs moved to the “previous host” after each ESXi host is upgraded			Yes	Takes longer than a parallel upgrade, but vNS system stays up. Uses a spare host to pre-check that host upgrade is successful. Cabling changes may be required for certain vNS system configurations, making this method disadvantageous in such cases.
Rolling upgrade with spare host + VMs return to original location after each ESXi host is upgraded				Slowest method, but vNS system stays up. Uses a spare host to pre-check that host upgrade is successful. Cabling changes not required.

Related or recommended talks on the subject

Talk ID	Title	Presenter	Date & Time
TBC23-TB55	HPE Virtualized NonStop Continues the Journey to the Cloud (repeats on Wednesday)	Marcelo de Azevedo Ken James Bryce Kosinski Spencer Kropp Lars Plum	Tuesday, Sep 12, 2023 11:00 AM - 12:00 PM Wednesday, Sep 13, 2023 10:30 AM - 11:30 AM
TBC23-TB63	Evolving Your vNS Environment to Keep Your NonStop Business Running	Mark Thompson	Wednesday, Sep 13, 2023 11:30 AM - 12:30 PM
TBC23-TB56	Provide the mission-critical advantage to workloads in your private cloud with HPE Virtualized NonStop	Marcelo de Azevedo Bryce Kosinski	Wednesday, Sep 13, 2023 2:45 PM – 3:45 PM
TBC23-DEV3	HPE Virtualized NonStop developer chat room	Marcelo de Azevedo Lars Plum Bryce Kosinski	Thursday, Sep 14, 2023 9:00 AM – 10:00 AM
TBC23-TB57	HPE Virtualized NonStop Storage	Marcelo de Azevedo Lars Plum	Thursday, Sep 14, 2023 11:15 AM - 12:15 PM

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*

NonStop Partnership– It's a Beautiful Thing!

Thank you for attending this talk TBC23-TB58 Best Practices for Configuring and Managing HPE Virtualized NonStop Systems for Mission Critical Workloads

For more information send email to:

P.Kamath@hpe.com
Marcelo.DeAzevedo@hpe.com
Bryce.Kosinski@hpe.com
Lars.Plum@hpe.com
Mark.T@hpe.com
John.Zimsky@hpe.com

HPE Slides and Materials Usage

This content is protected

This presentation is the property of Hewlett Packard Enterprise and protected by copyright laws of the United States. The material in this presentation is provided to attendees of the NonStop Technical Boot Camp 2023 as part of their registration and attendance at the event. Attendees are free to use this material and share it with others within their own company.

This material may not be quoted, copied, communicated or shared with third parties or mutual customers without permission from HPE. To request permission to share material in this presentation outside of your company, send an email to mark.pollans@hpe.com explaining the usage you are intending and your request will be considered.