



Hewlett Packard
Enterprise

Deploying VMware vSphere on HPE Synergy

Best Practice Guide

February 2017

Contents

Introduction..... 4

HPE Synergy and VMware vSphere Best Practices..... 5

 Cluster Design with HPE Synergy 6

 HPE Synergy Scalable Fabric Design 8

 Server Profile/Server Profile Template Design..... 18

 Ethernet/Storage Connections..... 18

 BIOS Settings 22

 Server Profile Template Best Practice Summary: 23

 HPE Synergy Shared Storage Options 24

 VMware vSphere Hypervisor (ESXi) 6.5 Provisioning..... 27

 HPE Synergy Image Streamer 29

 Server-Side Link Aggregation and vSphere Distributed Switches 31

 Lifecycle Management..... 40

 HPE OneView for VMware vCenter Integration..... 42

Appendix A - HPE Synergy Architecture Overview..... 44

Appendix B - HPE Synergy Configuration Overview 45

Appendix C - Referenced HPE and VMware Best Practice Documents 47

Introduction

IT organizations are looking for an automatic way to configure the hardware necessary to provide only the resources needed to run an application. They need this composable on-demand capability, delivered in a robust infrastructure which is flexible, cost-efficient, and scalable. The solution needs to provide a stable platform for mixed workloads without compromising on performance. What they are looking for should enhance their ability to meet the required service-level agreements to their users.

HPE Synergy is architected from the ground up for Composable Infrastructure. The performance, flexibility, and manageability that HPE Synergy provides, help businesses provision quickly and efficiently on-premise infrastructure resources. HPE Synergy provides the necessary management solution for the entire life cycle of deployed infrastructure. From scaling deployed resources through maintenance and retirement of these resources, HPE Synergy reduces associated IT management costs. Another benefit of HPE Synergy is reduced costs associated with deploying unnecessary stranded resources through rightsizing the deployed infrastructure in support of the workloads deployed on the HPE Synergy solution. It also provides a single unified API to assemble the infrastructure to meet business needs.

Working together, Hewlett Packard Enterprise and VMware have designed a highly efficient IT infrastructure powered by VMware running on HPE Synergy. A composable infrastructure from Hewlett Packard Enterprise combined with VMware® vSphere and HPE Gen9 and Gen10 Servers, enables a best in class virtualization experience. This enables organizations to take advantage of a disaggregated pool of compute, storage, and fabric resources. Allowing the infrastructure administrator to avoid deploying fixed ratios of these resources typically found in traditionally deployed rack server types of infrastructure.

This ***Deploying VMware vSphere on HPE Synergy Best Practices Guide*** provides a set of standard practices for deploying VMware vSphere clusters on HPE Synergy infrastructure.



If the reader of this document is new to the HPE Synergy architecture and HPE OneView concepts, the Appendix section contains background/summary information on HPE Synergy setup and HPE OneView terms used in this document. All the published best practices referenced in this guide are listed in Appendix D.

HPE Synergy and VMware vSphere Best Practices

HPE Synergy and VMware have been technology partners for well over a decade. Technologies from VMware combined with HPE infrastructure are at the core of creating the new software defined datacenter.



HPE Synergy supports key virtualization features including SR-IOV, VMware vSAN, VMware NSX, BIOS Virtualization Workload Settings, and a new stateless provisioning option called HPE Image Streamer.

The best practices section of this document is organized into eight categories. Each section provides a brief narrative of the goals to be achieved and along with a set of recommended best practices.

An important note on this VMware vSphere guidelines/best practices for HPE Synergy, the options for any HPE Synergy solution are numerous. Hardware options and software defined best practices are all derived from customer specific requirements and workloads. This document provides generalized examples of best practice configurations based on the experience of HPE engineers.

The following sections provide VMware/HPE Synergy best practices in the following categories:

1. Cluster Design with HPE Synergy
2. HPE Synergy Fabric Design
3. HPE Synergy Server Templates/Server Profile Templates
4. HPE Synergy Storage
5. HPE Image Streamer
6. VMware vSphere Provisioning
7. Lifecycle Management
8. HPE OneView/VMware vSphere Integration

Cluster Design with HPE Synergy

HPE Synergy Multi-Frame architecture provides several options for designing your VMware vSphere clusters, both within a single rack and spanning multiple racks for high-availability. The first example uses HPE Synergy configured in a 3-Frame logical enclosure with six VMware vSphere clusters configured in the environment. Each 6-node VMware vSphere cluster spans the three HPE Synergy 12000 frames in a design to improve high-availability. All the 36 HPE Synergy compute nodes share a pair of redundant master networking modules providing access to any necessary external Ethernet and FC networks.



The second example illustrates a HPE Synergy multi-frame configuration with the vSphere cluster nodes distributed across racks to improve high-availability. In this example three racks contain nine Synergy 12000 frames. A total of 108 compute nodes are used to create twelve vSphere clusters, each with 6 nodes.

HPE Synergy Multi-Frame technology provides scalability while managing the devices as a single *logical enclosure (LE)*. Network performance is not degraded as frames are added to the logical enclosure. HPE Synergy provides zero oversubscription between all compute nodes in the LE configuration.

Clusters should be homogeneous with all components installed identically across all ESXi hosts. This is to promote a reduction of service related events by making sure all firmware, drivers, and software are at the same revisions.

Other cluster design considerations that are specific to each customer (and therefore not covered in detail in this paper) environment are:

- Number of member compute nodes in the Cluster
- Network design and traffic segmentation
- Distributed vSwitch vs Standard vSwitch
- SAN Volume Size
- VMware vCenter host, processor and memory considerations
- DRS/HA configuration
- Physical infrastructure and virtualization maximums (i.e. virtual machines per ESXi host)




In this third example of an HPE Synergy Multi-Frame configuration, a three frame environment provides shared storage using the D3940 Storage Modules. The solution is self-contained with compute, storage, and networking in a single rack environment. The solution contains high-availability components throughout. This design has the advantage of scalable software defined storage resources being delivered within the same physical infrastructure. Virtualization capacity can be quickly added into your datacenter without reliance on existing SAN arrays, SAN fabrics or management systems.

There are four vSphere clusters, each with six SY480 compute nodes. Each node also functions as a vSAN storage node to evenly distribute storage I/O across the cluster. Each Server Profile has 3 x 800GB SSDs for the caching tier, 7 x 1.6TB SAS drives for the data capacity tier. All logical disks are configured in HBA mode (no RAID). Each vSphere cluster has 60+TB of raw storage available.


Although six-node clusters are used in this example, you have full flexibility to configure larger or smaller cluster sizes based on your specific needs.

Each SY480 Gen10 compute node supports up to 3.0 TB of DDR4 HPE SmartMemory and up to 56 total CPU cores (two sockets/28 cores), capable of hosting hundreds of virtual machines per compute node. The VMware vSAN storage can provide 50+TB of raw storage to each vSphere cluster.



HPE Synergy Scalable Fabric Design

The strength of HPE Synergy comes from a template-driven configuration strategy where you define the best practice once, then instantiate that configuration *N* times across the datacenter. The following diagram illustrates this concept, once the templates are defined, scaling HPE Synergy at rack scale is a simple process of creating/configuring new infrastructure from a single screen. The process starts by defining your best practices through one or more fabric templates (LIG), an enclosure template (EG), and server profile templates (SPT).



For an overview of the HPE Synergy powered by HPE OneView installation/setup process, refer to Appendix B of this document.

Composable Storage, Ethernet and Fibre-Channel fabrics are defined in the Logical Interconnect Groups (LIG) template. Each HPE Synergy 12000 frame has six interconnect bays, each paired for an A-Side/B-Side design to support fabric redundancy. Depending on the interconnects installed, one to three of these LIG templates must be created.

Step 1: Create the template for the HPE Synergy interconnect fabrics (LIG)


Step 2: Create the template for the enclosure group (EG) using Step-1 fabric templates

Step 3: Create a Logical Enclosure using the Step-2 Enclosure Group template

This following section provides general best practices for creating the LIG and EG within HPE OneView. For readers wanting to understand Virtual Connect features and traffic flow in more depth, please refer to the technical whitepaper **Virtual Connect Traffic Flow with HPE Synergy**. The document link is in Appendix D.

General fabric module placement:


1. --Composable Storage with HPE Synergy 12gb SAS Connection Modules in bay set 1 & 4
2. SAN Storage with either the HPE Virtual Connect SE 16Gb FC Modules for Synergy or the Brocade 16Gb Fibre Channel SAN Switch for Synergy in bay set 2 & 5
3. Converged Ethernet/FCoE/FC via HPE Virtual Connect SE 40gb F8 Modules in bay set 3 & 6



For example, a HPE Synergy 12000 frame with six interconnects consisting of two HPE Synergy 12gb SAS Connection Modules, two HPE Virtual Connect SE 40gb F8 Modules, and two HPE Virtual Connect SE 16Gb FC Modules, would have three LIGs configured within the HPE Synergy composer. Within each LIG, uplink sets define the physical uplinks that connect the interconnects to a TOR switch (Ethernet trunks), Fibre-Channel edge switches, and/or iSCSI storage connections, and/or Flat-SAN (to 3PAR StoreServ Systems).

Multi-Module Link Aggregation Group (MC-LAG)


The HPE Virtual Connect SE 40Gb F8 Module is a converged switch capable of Ethernet, iSCSI and Fibre-Channel uplinks. With HPE Synergy, an active/active Multi-Module Link Aggregation Group (MC-LAG) is supported. In redundant mode, a paired set of HPE Synergy interconnect modules are configured as a single logical switch. In an MC-LAG active/active configuration each HPE Virtual Connect SE 40Gb F8 module supports up to six 40Gb uplinks or 24 10Gb uplinks. The following diagram illustrates two valid active/active MC-LAG configurations.



Best practice: Stacked High-Availability switch configuration with HPE Synergy MC-LAG uplink set. Maximum active members in a single LACP group is 8, HPE Synergy supports up to 16 uplinks per module, 32 total. All must terminate to a single or stacked switch with the same port speed. Physical uplinks should be evenly distributed across the stacked switch pair.

Non MC-LAG Active/Active Option

HPE Synergy also supports non MC-LAG configuration to two TOR switches in an active/active configuration using HPE SmartLink. MC-LAG with S-channel downlinks offers better high-availability and performance over Active/Standby, however, it requires a stacked TOR switch configuration.




Best practice: Required if ToR switches are not stacked and uplink HA is required. Recommend 2+ physical uplinks per module. All LACP members in the port group must be same speed. HPE SmartLink must be enabled within HPE OneView Network properties.

The following sections will highlight two configuration examples for the Logical Interconnect Groups.

Example 1: Converged Ethernet/Fibre-Channel

The following screenshot is an example of a converged design from a Multi-Frame environment using Ethernet and Fibre-Channel uplinks from the HPE Virtual Connect SE 40Gb F8 Module pair. In this example we have 4 uplink sets defined within the shared uplink set: one uplink will be used for HPE Synergy Image Streamer capability, one uplink set is defined for the Ethernet MC-LAG trunk, and two uplink sets are for the Fibre-Channel uplinks SAN-A and SAN-B.



Best Practice: Use MC-LAG with active/active uplinks to your TOR switch for optimal load balancing and resiliency. Up to 8 x 40 Gb active uplinks can be configured into a single uplink set. Consider segmenting datacenter management traffic from VM production traffic through separate uplink sets. For large dense virtualization environments, consider segmenting FC traffic from Ethernet using VCFC 16Gb interconnect modules for maximum SAN storage IO performance.

Example 2: Non-Converged

This is a non-converged example from a single frame with redundant interconnect modules. A pair of HPE Virtual Connect SE 16Gb FC Modules are used to segment FC traffic and maximize IOPS to the HPE 3PAR StorServ Storage System. In the example below, the LIG named VCFC is redundant and has two uplinks sets: one for SAN-A, one for SAN-B. New with HPE OneView 4.0 is ISL trunking which provides load balancing across uplinks for this FC configuration.

OneView

Search

Logical Interconnect Groups 3All typesAll resourcesAll labels

+ Create logical interconnect group

Name

SAS

SUS

VCFC

VCFC

Logical Interconnect Group

Redundancy

Redundant

SAN-A

1 network

2 uplink ports

SAN-B

1 network

2 uplink ports

2

1 2 3 41 2 3 41 2 3 41 2 3 4

Q1Q2Q3Q4

Virtual Connect SE 16Gb FC Module for Synergy

1 2 3 4 5 6 7 8

1 2

5

1 2 3 41 2 3 41 2 3 41 2 3 4

Q1Q2Q3Q4

Virtual Connect SE 16Gb FC Module for Synergy

1 2 3 4 5 6 7 8

1 2

February 2018

Mapped VLAN vs Tunnel

HPE Synergy supports both mapped VLAN and tunnel networks. Up to 3966 VLANs are supported in mapped VLAN mode. HPE Synergy supports 200-1000 VLANs in a network set depending on number of frames in the LIG.

To ease VLAN administration and eliminate VLAN restrictions imposed by the HPE Virtual Connect SE 40Gb F8 Module, HPE Synergy also supports tunnel mode where Ethernet traffic is “tunneled” through the uplink. However, when in tunnel mode, the HPE Virtual Connect SE 40Gb F8 Module maintains a single MAC Address table for the tunneled VC network even though it encompasses multiple VLANs (this is normal). The result is that when a host (physical or VM) inside the VC Domain sends a broadcast like an ARP, it is sent out the VC uplink on one VLAN, traverses through the load balancer and is broadcast on the load-balanced VLAN. If that VLAN is also sent to the VC uplink port, the MAC address of the host is learned outside of VC. Like any 802.1d bridge, subsequent traffic sent to that host's MAC address and received on the VC uplink is discarded as VC has learned that the MAC address resides outside the domain. The normal MAC address aging time in HPE Virtual Connect SE 40Gb Modules is 5 minutes, so this condition will exist until the entry is aged out. The problems can occur when two or more VLANs are bridged together (i.e. two VLANs in one broadcast domain) AND those two or more VLANs are included on the same Tunnel Mode network. To prevent this issue either use switches from vendors that do not use VLAN bridging (F5) or when you use upstream switches that are using VLAN bridging (CISCO) either turn tunneling off or don't bridge two or more VLANs and have them in the same tunnel network.


Best Practice: For small/medium environments, consider combining management and VM traffic on an uplink set using mapped VLAN mode. For large dynamic (i.e. DevOps) environments consider using Mapped VLAN mode for management traffic. Then configure the virtual machine (VM) traffic on a tunnel uplink set. Be aware of the bridging issue described above. This increases the number of physical uplinks to the TOR switch but this scalable design physically segments management traffic from VM production traffic and reduces VLAN administration within HPE Synergy.

Number of uplinks in MC-LAG

Each redundant pair of Virtual Connect SE 40Gb F8 modules provide four adjustable downlinks connections (three Ethernet and one Fibre Channel) to each compute module. A Fibre Channel Upgrade license needs to be purchased separately to enable the Fibre channel functionality. Up to six uplinks using QSFP+ interfaces are available for connection to upstream switches. Currently each port supports 40Gb uplinks with using QSFP+ adapters connected to Ethernet switches. The modules paired can support up to 24 connections to upstream Ethernet and Fibre Channel switches by utilizing splitter cables.


With the Virtual Connect SE 40Gb F8 module up to 8 *active* 10Gb OR 40Gb uplinks can be supported in an LACP group. It is possible within HPE OneView to configure up to 16 uplink ports as part of a single uplink set, however, only 8 of those uplinks will be active at any given time. When using LACP, the protocol requires all uplink ports be at the same speed.

Any Ethernet, FC and FCoE networks must be on separate uplink sets. In an HA redundant configuration using the Virtual Connect SE 40Gb F8 module, two QSFP ports (X7/X8) are required for IRF cluster interconnects between modules. As mentioned before, on each module six QSFP ports are available for uplinks providing a total of two modules x six QSFP ports x 40Gb each = 480Gb total uplink bandwidth available.



LACP-S (S-Channel Support)

HPE Synergy supports LACP-S also known as S-channel support, on the downlinks to Converged Networking Adapters (CNA) in the compute nodes. S-channel provides faster fault recovery and load balancing over traditional OS based NIC teaming. MC-LAG and S-Channel configuration are independent of each other. MC-LAG is configured between the interconnect uplinks and TOR switch(s) whereas S-Channel is configured between the interconnect downlinks and the server. MC-LAG combined with S-channel provides greater resiliency for VMware vSphere hosts.



The HPE Virtual Connect SE 40Gb F8 Module for HPE Synergy supports associating connections to a compute module to form a Link Aggregation Group (LAG). LAGs on server ports allow for outage-free interconnect firmware updates, seamless failover, improved traffic load-balancing, and switch-assisted NIC teaming policies.

To create a LAG between connections in a server profile, select the same valid Link Aggregation Group (LAG 1-LAG 16) for each of the connections. A LAG contains exactly two Ethernet connections. With the 3820c CNA up to 4 LAGs can be created, up to a total of 24 LAGs per server profile. Both connections must:

- Have the same requested bandwidth.
- Access the same network or network set.
- Be configured on the same HPE Synergy Converged Network Adapter on both physical ports of the one adapter, and adapter ports must be connected to interconnects in the same logical interconnect.

The teaming or bonding driver in the server OS must be configured to enable the Link Aggregation Control Protocol (LACP) and is required for the LAG to form correctly.

Best Practice: Within Server Profiles enable S-channels on NIC pairs whenever possible, S-channel provides faster fault recovery and load balancing compared to traditional NIC teaming with OS-based load balancing. S-Channel requires VMware distributed virtual switch on the VMware ESXi host.

Internal Networks

Networks are defined within HPE OneView and selected when creating either the Server Profile or a Server Profile Template (SPT). Any networks utilized by either the server profile or a SPT, must exist and be defined in the logical interconnect the compute hosts are connected to. For networks not traversing a connected TOR switch, HPE OneView supports internal networks which are defined within the Logical Interconnect Group (LIG) template definition. VMware vMotion and fault tolerance are excellent examples of internal networks between the hosts and which typically do not transverse the TOR switch.

Best Practice: Use Internal networks to isolate Ethernet traffic which is contained within HPE Synergy Multi-Frame environment. VMware VM migration and Fault Tolerance logging traffic are example of internal networks.

Private Networks

New with HPE OneView 4.0 private VLANs are now supported. Private VLANs (pVLANs) are known as port isolation, a security mechanism in networking where the VLAN communication is restricted to the uplink (peer-to-peer communication is blocked). Private VLAN is much needed Layer 2 security, important in public sector and multi-tenant environments. Note: the private VLANs feature (isolated/community/promiscuous) is not the same as legacy Virtual Connect private vlan type.

Best Practice: Consider using pVLAN when enhanced Layer-2 security is required.

DirectPath I/O and SR-IOV

DirectPath I/O (DPIO) improves performance on a VM by decreasing the number of CPU cycles needed to run the ESX/ESXi hypervisor. It leverages hardware support (Intel VT-d and VMD-V) to allow guest to directly access hardware devices. VMware recommends using DirectPath I/O only with high-priority VMs that require maximum network performance. The technology is a special use case and has several restrictions.

HPE Synergy network adapters support SR-IOV Single root input/output virtualization (SR-IOV) provides isolation of PCI express resources for performance. SR-IOV enables network traffic to bypass the software switch layer typical in a hypervisor stack which results in less network overhead and performance of non-virtualized environments. With SR-IOV technology provided on Flex NICs there is a major advantage for HPE Synergy compute nodes. The SR-IOV architecture on Virtual Connect allows the expansion of up to 512 PCIe functions, depending on hardware capability. For example, if you have an Ethernet adapter supporting 64 virtual functions (VFs) per physical port, on a single FlexNIC you can configure 8 to 64 VFs using Virtual Connect. The VFs can be assigned to one or more VMs each having its own IP and dynamic MAC Address. This provides a significant increase in server consolidation and performance.

Best Practice: For high-performance virtualization environments, consider enabling SR-IOV to improve network performance. Directpath I/O is another VMware performance option which can be enabled for special use cases. However, DPIO has serious HA and other limitations including vMotion and Snapshots.

Server Profile/Server Profile Template Design

Server profiles define the hardware personality of HPE Synergy Compute nodes. Server Profile Templates (SPT) provide a mechanism to store best-practice configurations for multiple server profiles. Server profiles created and deployed from templates provide consistent configurations for servers servicing identical workloads. Server profile templates also provide a single entity for managing change to multiple server profiles that are based off a template. HPE OneView provides compliance alerts/reports to ensure any server profile deviation from its parent template can be acknowledged or updated.

Server Profile (and Templates) include the following settings:

- Basic server identification information
- Template membership
- Hardware type
- Ethernet and Fibre-Channel networking connectivity settings
- Firmware baseline
- Local and SAN storage settings
- BIOS / UEFI configuration
- Boot order
- Physical or virtual UUIDs, MAC addresses, WWN addresses

The following section provides an **example** of a VMware ESXi server profile template including Network Connections, Local Storage, SAN Storage, and BIOS settings.

Ethernet/Storage Connections

A HPE Synergy 480 Gen10 Compute Module with a single HPE Synergy 3820C 10/20Gb Converged Network Adapter with two physical LAN ports, each with 4 flexible NIC functions. The following screen illustrates all 8 NICs configured into 3 NIC teams for management, VM migration/FT, and production traffic.

Server Profile Template Connection Example:


Connections☒ Manage connections

ID	Name	Network	Port	Boot		
1	Mgmt-NS	<u>Mgmt-NS</u> (network set)	Mezzanine 3:1-a	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2 Gb/s				
2	Mgmt2-NS	<u>Mgmt-NS</u> (network set)	Mezzanine 3:2-a	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2 Gb/s				
3	esxFT-NS	<u>esx-FTVmotion-NS</u> (network set)	Mezzanine 3:1-c	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2 Gb/s				
4	esxFT2-NS	<u>esx-FTVmotion-NS</u> (network set)	Mezzanine 3:2-c	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	2 Gb/s				
5	SAN A	<u>SAN-A</u> Fabric attach	Mezzanine 3:1-b	Not bootable		
	Type	Fibre Channel				
	WWNN	Auto				
	WWPN	Auto				
	MAC address	Auto				
	Requested bandwidth	8 Gb/s				
6	SAN B	<u>SAN-B</u> Fabric attach	Mezzanine 3:2-b	Not bootable		
	Type	Fibre Channel				
	WWNN	Auto				
	WWPN	Auto				
	MAC address	Auto				
	Requested bandwidth	8 Gb/s				
7	VMprod-NS	<u>VMprod-NS</u> (network set)	Mezzanine 3:1-d	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				
	Requested bandwidth	8 Gb/s				
8	VMprod2-NS	<u>VMprod-NS</u> (network set)	Mezzanine 3:2-d	Not bootable		
	Type	Ethernet				
	MAC address	Auto				
	Requested virtual functions	None				

When creating connections in a server profile template or server profile, the administrator can set a requested preferred and maximum bandwidth for each connection. When multiple physical functions on the network adapter are assigned to different connections, this allows the administrator to guarantee a minimum level of bandwidth will be available for a specific connection.


Understanding the bandwidth needs of various network functions (management, virtual machine, vMotion) is critical to ensure optimal performance. The following diagram illustrates the physical NIC teams to vmnic to vSwitch/dvSwitch configuration. S-channel is enabled on the dvswitch uplinks.

Traditional Segmentation with Server Profile Max Bandwidth Settings



VMware Network I/O Control

Another best practice example utilizes VMware Network I/O control (NETIOC) which provides granular bandwidth control at the dvportgroup (vlan). Bandwidth reservations can be defined for classes of traffic. VMware distributed resource sharing works with NETIOC to ensure demand does not exceed capacity on any host within the cluster. NETIOC support bandwidth limitation at a VM virtual adapter level.



Local Storage

HPE Synergy supports a variety of local boot options for VMware ESXi servers including local disk (SAS/SATA), SSD, SD cards, and HPE D3940 Storage Modules. Within a SPT, a RAID volume can be defined and is automatically configured when the SPT is used to create the profile that is attached to a physical HPE Synergy compute node. In this example a RAID-1 boot volume will be configured using the two local SAS drives.

Local Storage

Integrated storage controller 

RAID

Initialization will occur on next assignment to server hardware

Name	Type	RAID Level	Number of Drives	Size GB	Drive Technology	Boot	
Boot	Logical drive	RAID 1	2	n/a	SAS HDD	☒	

SAN Storage

HPE Synergy provides fully automated HPE 3PAR StoreServ Storage provisioning and fabric zoning. Ephemeral and private volumes can be created within the SPT creation process using a wizard. Permanent and permanent shared volumes can be attached as illustrated in the screen shot below.

SAN Storage

☒ Manage SAN Storage

Host OS type

VMware (ESXi) 

Volume Attachments

[Expand all](#) [Collapse all](#)

Name	LUN	Pool	Size	Sharing	Enabled Paths	Boot		
▶ Cluster1-DS1	Auto	H4-NL-F5	500.00 GiB	Shared	2	n/a		
▶ Cluster1-DS2	Auto	H4-NL-F5	500.00 GiB	Shared	2	n/a		

Best practice: Create permanent/shared HPE 3PAR StoreServ volumes using the HPE OneView volume template feature. Also make sure multi-pathing is enabled for HA. Consider creating separate CPGs for all HPE Synergy attached storage. Manage the volumes within HPE OneView including snapshots and clones. For environments utilizing the HPE Virtual Connect SE 16Gb FC Module, utilize FC port trunking on the FC uplinks.

BIOS Settings

HPE Synergy 480 Gen10 Compute Modules have Workload Profile BIOS settings. This feature sets all the recommended BIOS settings for a given workload by simply selecting a single choice. There are two workload profile choices for virtualization hosts: Virtualization – Power Efficient or Virtualization – Max Performance. Choose the appropriate workload profile which meets your specific needs.

The screenshot shows the 'Edit BIOS Settings' window. The 'Workload Profile' dropdown menu is open, displaying the following options:

- Default - General Power Efficient Compute
- General Power Efficient Compute
- General Peak Frequency Compute
- General Throughput Compute
- Virtualization - Power Efficient
- Virtualization - Max Performance
- Low Latency
- Mission Critical
- Transactional Database
- High Performance Compute (HPC)
- Decision Support
- Graphic Processing
- I/O Throughput
- Web/E-commerce
- Extreme Efficient Compute
- Custom

The 'Virtualization - Max Performance' option is currently selected. Below the dropdown, the 'BIOS Serial Console and EMS' section is visible, showing three settings:

- BIOS Serial Console Port: Default - Auto
- BIOS Serial Console Emulation Mode: Default - VT100+
- BIOS Serial Console Baud Rate: Default - 115200

At the bottom right, there are 'OK' and 'Cancel' buttons.

Server Profile Template Best Practice Summary:

- Take advantage of HPE Virtual Connect technology and virtual identity settings within the Server Profile. Virtual Connect provides profile mobility across different compute nodes regardless of generation. In this way a simple movement of a profile can be used to scale up, scale down, archive, or repurpose the servers.
- Define your VMware ESXi host best practice through HPE OneView's Server Profile Template (SPT) and include firmware along with standard settings for the servers.
- Utilize S-Channels and NIC Link-Aggregation groups to enhance HA and traffic load balancing
- Set the Virtualization Workload Profile in the SPT so that every server gets those BIOS settings
- Ensure multi-path is configured to any and all SAN volumes
- Set the preferred bandwidth on all storage connections to the maximum speed for the storage system port or fabric speed
- Consider using VMware I/O Control (NetIOC) to intelligently configure bandwidth reservations for classes of traffic (including vSAN), bandwidth reservations per VM, and network load balancing across clusters using DRS
- Configure HPE OneView Network Sets to configure 802.1q trunks on Synergy Converged Network Adapters
- Set firmware baseline in the SPT to ensure consistency and compliance
- Since SPTs are associated with HPE Synergy hardware, provide a unique and descriptive name for the Server Profile Template, i.e. VMware vSphere-sy480
- Define preferred/maximum Synergy Converged Network Adapters bandwidth settings in creating Networks and Network Sets
- Create volume templates (VT) within HPE OneView for your shared datastores
- Configure the appropriate VMware vSphere load balancing algorithms in NIC teams on vSwitch/dvSwitch. Default algorithm is route based on originating virtual port.

Load balancing method	Available for	Resource consumption	Requires physical switch configuration
Route Based on Originating Virtual Port	VSS & VDS	Low	No
Route Based on Source MAC Hash	VSS & VDS	Higher	No
Route Based on IP Hash	VSS & VDS	Highest	Yes
Route Based on Physical NIC Load	VDS	Low	No
Explicit Failover Order	VSS & VDS	Low	No

VSS – vSphere Standard Switch
VDS – vSphere Distributed Switch

- Where vMotion performance is vital, consider provisioning dedicated 10Gbs network with teamed vmknics
- Keep VM hardware version and VMware tools up to date

HPE Synergy Shared Storage Options

HPE Synergy provides unmatched levels of shared storage automation and fabric options. This section highlights the key features of the new fully composable HPE Synergy D3940 Storage Module as well as the fully integrated capabilities with HPE 3PAR StorServ products.

HPE Synergy has several shared storage solutions for VMware vSphere including:

1. Composable Storage Module (D3940) with VMware vSAN
2. Fully automated storage provisioning using 3PAR StorServ Storage Systems
3. Unmanaged external 3rd party storage arrays (i.e. NetApp/EMC)

The following section will summarize best practice information for Options 1 and 2.



Option 1: Shared storage using the HPE Synergy D3940 Storage Module and VMware vSAN.

VMware vSAN and the HPE Synergy D3940 Storage Module is a cost-effective option for providing RAID protected shared datastores for ESXi clusters.

HPE Synergy D3940 Storage Module Summary Specs:

- 12Gb SAS fabric providing 2M IOPS
- Supports SATA, SAS, and, SSDs
- Optimized for HPE StorVirtual VSA
- Up to 40 SFF drives per unit
- Redundant IO paths
- Non-Disruptive Firmware Updates
- RAID support 0,1,1+0,5,6
- Up to 64 logical drives per server
- Use up to four HPE Synergy D3940 Storage Modules per HPE Synergy 12000 Frame hosting the vSAN infrastructure for a total of 480TB of SAS storage



A detailed technical whitepaper is available with the installation/configuration of the HPE Synergy D3940 Storage Modules, please refer to the links in Appendix A for more information.

Best Practice: VMware vSAN is best implemented in a multi-frame environment with one of the HPE Synergy frames providing shared storage services. VMware recommends a minimum three-node VMware vSAN cluster. Utilize all-flash devices in the HPE Synergy D3940 Storage Modules and VMware vSAN network bandwidth should be a minimum of 10Gb for maximum performance.

Option 2: Shared storage using HPE 3PAR StorServ Storage Systems

HPE Synergy with HPE 3PAR StorServ solutions offers unmatched performance and automation in a converged platform today.

HPE OneView utilizes open RESTful API for integration with all HPE 3PAR StorServ systems. These systems are easily added to HPE Synergy in seconds. HPE Synergy supports the multi-tenancy features in HPE 3PAR StorServ such as virtual domains and common provisioning groups (CPGs) to manage/provision all or just a subset of resources in the array.

HPE Synergy supports all leading SAN managers today including Cisco, Brocade, and HPE Networking. Discovered storage fabrics automatically detect fabric ports on the HPE 3PAR StorServ arrays. New with HPE OneView 4.0 is FC ISL load balancing across FC uplinks.

There are several VMware/HPE 3PAR StorServ best practices documents available on the web, the links are provided in Appendix A of this document.

HPE Synergy support both Native FC, FCoE and iSCSI storage fabrics. HPE Synergy supports automated fabric zoning and storage provisioning to HPE 3PAR StorServ arrays including thin provisioning and fabric zoning policies. HPE Synergy supports ephemeral/permanent volumes; private/shared volumes, and thin provisioning. Just like everything else in HPE Synergy shared volumes can be defined through a template for standardization and reuse.



Best Practice: Utilize All-Flash HPE 3PAR StorServ 9000 arrays typically or HPE 3PAR StorServ 20000 arrays for large scale/dense virtualization environments. Use HPE Synergy 16G Fibre Channel Host Bus Network Adapters along with HPE Virtual Connect SE 16Gb FC Modules. Utilize multi-path for storage HA and ISL load balancing (Brocade only). HPE 3PAR StorServ automatically optimizes the storage for maximum performance so storage decisions in server profiles and SPTs are as simple as the size, thin provisioning option, and the storage pool (RAID type) the compute node needs to utilize.

VMware vSphere Hypervisor (ESXi) 6.5 Provisioning

VMware vSphere has multiple deployment options on HPE Synergy. You can choose the options that best meet the needs of your environment, and can use either an interactive, image based, or a scripted installation when installing VMware vSphere software on HPE Synergy compute nodes.

VMware vSphere ISO image options

VMware vSphere offers different image (ISO file) options. You can use the VMware base image, an HPE Custom Image, or your own custom image.

- **VMware base image:** The VMware base image includes the base VMware vSphere Hypervisor (ESXi) 6.5 as the hypervisor and an initial set of drivers. This image does not include server-specific management components, or new or updated drivers. If you choose to use the VMware base image, we recommend that you add HPE management components to the image after installing VMware vSphere. You can use the VMware base image to deploy many ProLiant servers. However, there is a subset of ProLiant servers that requires you to use the HPE Custom Image or develop a custom ISO image that includes the required device drivers.
- **HPE Custom Image:** The HPE Custom Image offers the easiest method to deploy VMware vSphere Hypervisor (ESXi) 6.5 on ProLiant servers. This image includes HPE management tools, utilities, and drivers that allow easy deployment of a fully managed ProLiant server. The HPE Custom Image is available at hpe.com/info/esxideownload.
- **Custom ISO image:** You can create your own ISO image for deploying VMware vSphere Hypervisor (ESXi) 6.5. To create a custom ISO image, use the VMware Image Builder included with VMware vSphere PowerCLI. You can start with either the VMware base depot or the HPE custom depot, and then customize it to include the VMware vSphere Installation Bundles (VIBs) that you want to load at boot time. VIBs can include additional drivers, HPE components, and other third-party components. The output of Image Builder will be a custom ISO image that includes the VMware base image or HPE Custom Image and all the added VIBs.

HPE Custom Images

The HPE Custom Image provides an integrated, well-tested image. HPE tests customized images against all applicable servers. Tests include scenarios for updates, and upgrades, and for configuration changes.

HPE Custom Images simplify configuring and deploying VMware vSphere on ProLiant servers. You can use an HPE Custom Image with any applicable ProLiant server; however, certain servers require an HPE Custom Image or customer-created custom image with the appropriate device drivers. For a list of servers that require a custom image, see the HPE ProLiant Server VMware Support Matrix at hpe.com/servers/vmwarecert.

To ensure you have all the components for a VMware vSphere installation, we recommend you do the following: 1. Load the HPE Custom Image, which includes all drivers and HPE additions. 2. Update the system firmware to align with the driver versions included in the HPE Custom Image.

We recommend using the HPE Custom Image to ensure you have all the drivers and HPE additions required to install VMware ESXi to ProLiant servers using the latest hardware. In addition, you should update the system firmware to align with the driver versions included in the HPE Custom Image. See the SPP documentation at hpe.com/info/spp and the “HPE ProLiant server and option firmware and driver support recipe” document on vibsdepot.hpe.com for information on SPP releases supported with HPE Custom Images.

For information on updating system components, see the “VMware: Best practices for maintaining HPE firmware, driver, and system software using the HPE Service Pack for ProLiant (SPP) and the HPE Software Delivery Repository (SDR)” white paper.

SAN configuration

One common implementation for VMware ESXi is a boot-from-SAN configuration. Installing the operating system in this configuration means that the enabling device drivers must be available in the VMware vSphere installation images. To simplify installing a boot-from-SAN configuration on ProLiant servers, the HPE Custom Image includes the required drivers (FC, FCoE, and iSCSI) for both embedded and supported optional devices that provide SAN connectivity. The result is a seamless process for setting up a boot-from-SAN VMware vSphere environment on ProLiant servers.

HPE Custom Image updates

HPE releases new custom VMware ESXi images when there is a need to update any critical components of the image. For example, when HPE introduces a new server that needs a new or updated networking driver, HPE then releases a custom image that includes the necessary drivers for that new server. VMware typically releases major, minor, or updated releases of VMware vSphere several times a year. We do not deliver a new custom image based on VMware patches or hot fixes. HPE supports application of VMware patches and hot fixes to the HPE Custom Image.

You can get HPE custom VMware vSphere images at hpe.com/info/esxideownload, where you can also download your required version of ESXi

- VMware vSphere Host Provisioning Options:
- HPE Synergy Image Streamer (iSCSI image boot)
- VMware vSphere AutoDeploy (PXE boot)
- HPE Insight Control Server Provisioning (ICSP)
- ILO Advanced virtual media

VMware vSphere Host Boot Options

- Local boot: SD Boot, M.2, local SATA/SAS with optional RAID protection
- SAN Volume: Boot from SAN (BFS) HBA
- Composable Storage: HPE D3940 Storage Module boot volume

Best Practice: HPE provides VMware ESXi custom images (VMware vSphere ISOs) with the latest drivers. HPE Synergy supports a variety of HPE and VMware provisioning tools and host boot options.

HPE Synergy Image Streamer

HPE Synergy Image Streamer is an optional appliance which can be configured into a HPE Synergy solution or added into an existing HPE Synergy environment. It provisions VMware vSphere ESXi nodes in minutes, all consistently based on a golden image. Image Streamer is fully integrated into HPE OneView delivered via the HPE Synergy Composer appliance and includes grow cluster capability within the VMware vCenter console using the OV4VC 9.0 plugin.

HPE Synergy Image Streamer provides an alternative method for deploying stateless VMware ESXi hosts that are not dependent on PXE or DHCP infrastructure.



Please refer to the HPE Synergy Image Streamer Quickspecs at the following link [here](#).

Manage Physical Servers like Virtual Machines

Deploy at Cloud-like speed for fast Time-to-Service by using HPE Synergy Image Streamer

Enhanced server profiles manage stateless Compute Modules like VMs

Virtual Machine Operation

- Create VM template with OS
- Deploy template to VM
- Update VM
- Hibernate VM
- Move VM
- Delete VM

Physical Server Equivalent

- Create logical server profile with OS
- Activate server profile on compute module
- Update server profile
- Deactivate server profile
- Move server profile
- Delete server profile




Image Streamer Features:

- Designed for scale-out server environments
- A full-feature framework for creating/maintaining golden OS images
- Supports VMware and Linux platforms
- Supports custom attributes for server personalization
- High-availability architecture for management and network boot
- Automated image streamer repository replication across the management domain
- No local boot drives, CAPEX/OPEX savings compared to legacy environments
- Deploy operating systems in minutes not hours => activation replaces OS scripted installations
- Built to support rolling upgrades in cluster environments

Best Practice: Image Streamer is an innovative provisioning option for HPE Synergy. An existing HPE Synergy environment can be upgraded with Image Streamer with no disruption to running systems. An Image Streamer HA pair can support up to 60 ESXi compute nodes. Upgrading ESXi 6.0 to 6.5 is as simple as changing the gold disk image pointer and rebooting the compute node.

Server-Side Link Aggregation and vSphere Distributed Switches

HPE Synergy environments using the Virtual Connect SE 40Gb F8 Module are capable of providing highly-available networking through Link Aggregation of the server network interfaces. Combined with LACP support on a vSphere Distributed Switch, ESXi hosts can contain multiple link aggregation groups (LAGs) on a distributed switch to aggregate the bandwidth of physical NICs on ESXi hosts that are connected to LACP port channels. Benefits include:

- Seamless failover between adapter ports.
- Enhanced server traffic load-balancing.
- Use of switch assisted NIC teaming policies.

Requirements:

- HPE Synergy Composers running OneView 4.0 or higher
- HPE Synergy Virtual Connect modules running firmware 1.2.0 or later
- VMware vCenter 6.x or higher

Detailed instructions for configuring vSphere Distributed Switches is available in the references in Appendix C the following examples represent the general flow.

Creating Network Connections in HPE OneView

When creating a Server Profile in HPE OneView, it is best practice to create two NICs for every Network, or Network Set assigned to a server. This provides path redundancy through the Virtual Connect modules.

Connections

[Expand all](#) [Collapse all](#)

	ID	Name	Network	Port	Boot
▶ ●	1		Management-Master (network set)	Mezzanine 3:1-c	Not bootable
▶ ●	2		Management-Master (network set)	Mezzanine 3:2-c	Not bootable
▶ ●	3	Deployment Network A	Deploy VLAN99	Mezzanine 3:1-a	iSCSI primary
▶ ●	4	Deployment Network B	Deploy VLAN99	Mezzanine 3:2-a	iSCSI secondary
▶ ●	5		SAN-A FCoE VLAN2101	Mezzanine 3:1-b	Not bootable
▶ ●	6		SAN-B FCoE VLAN2102	Mezzanine 3:2-b	Not bootable
▶ ●	7		Production-Master (network set)	Mezzanine 3:1-d	Not bootable
▶ ●	8		Production-Master (network set)	Mezzanine 3:2-d	Not bootable


Pairs of server NICs can be assigned to a Link aggregation group as part of the Server Profile or Server Profile Template. When these NICs are assigned to a VMware Distributed Virtual Switch enabled with LACP, they will form a highly available NIC team.

▼ ● 7	<u>Production-Master</u> (network set)	Mezzanine 3:1-d	Not bootable
Interconnect	<u>BOT-CN75150107, interconnect 3</u>		
Type	Ethernet		
MAC address	8E:8D:99:10:03:62 (v)		
Requested virtual functions	None		
Requested bandwidth	2.5 Gb/s		
Allocated bandwidth	2.5 Gb/s		
Max bandwidth	20 Gb/s		
Link aggregation group	LAG2		
▼ ● 8	<u>Production-Master</u> (network set)	Mezzanine 3:2-d	Not bootable
Interconnect	<u>MID-CN751500Z9, interconnect 6</u>		
Type	Ethernet		
MAC address	8E:8D:99:10:03:63 (v)		
Requested virtual functions	None		
Requested bandwidth	2.5 Gb/s		
Allocated bandwidth	2.5 Gb/s		
Max bandwidth	20 Gb/s		
Link aggregation group	LAG2		

Note: Do not assign a LAG to networks used to provision an Operating System via PXE boot. Without OS level drivers forming the LAG at time of the PXE request, the server will fail to communicate on the network. VMware Auto Deploy is an example of this scenario. Place the bootable NICs on a discrete deployment network and do not assign them a LAG

Create or Modify an Existing Distributed Virtual Switch

In order to ensure proper communication, the vDS and OneView Network Sets must contain the same VLANs.



OneView Search

Network Sets 3 All resources All labels

+ Create network set

Name

- Management-Master
- Production-Master**
- Test

Production-Master Overview

General

Preferred bandwidth 2.5 Gb/s

Maximum bandwidth 20 Gb/s

Used by none

Networks

Prod7	1117	Prod10	1120	Prod3	1113	Prod5	1115	Prod9	1119
Prod8	1118	Prod2	1112	Prod4	1114	Prod6	1116	Production	1111

On the vDS, click Configure and navigate to the LACP setting.

Synergy-LAG-dvSwitch Actions

Getting Started Summary Monitor **Configure** Permissions Ports Hosts VMs Networks

Settings

- Properties
- Topology
- LACP**
- Private VLAN
- NetFlow
- Port mirroring
- Health check
- Resource Allocation
- System traffic
- Network resource pools
- More
- Network Protocol Profiles

LACP

The enhanced LACP support on a vSphere distributed switch lets you connect ESXi hosts to LAGs.

Migrating network traffic to LAGs

+


LAG Name	Ports	Mode	VLAN
This list is empty.			

No items selected

Create the new LAG that will be associated with physical NICs of the Compute Module. Select whether the LAG negotiation mode with Virtual Connect mode is Active or Passive and the preferred load balancing mode. The default Load Balancing is *Source and destination IP address*.


For the LACP port channel, VMware recommends that the distributed switch and the physical switch must use the same hashing algorithm.

The hashing algorithm on the Virtual Connect modules for the LACP on downlinks is Source and Destination MAC Address. The same load balancing mode can be used on the distributed switch to follow the VMware's recommendations.




Click OK to save the new Link Aggregation Group.


Using the vCenter Action menu, select Distributed Port Group → Manage Distributed Port Groups.




Select the Teaming and failover policy to edit.



Select the port groups that will be associated with the teaming and failover policy.




Change the failover order to include the LAG in the Active uplinks, and move the default dvSwitch uplinks to Unused.




Note: VMware does not support mixing LAGs and standalone uplinks in the same Distributed Port Group.


Add vSphere hosts to the Distributed Switch by clicking Actions → Add and Manage Hosts. Select the hosts to add to the dvSwitch.




Choose which network adapter tasks to perform. In this scenario both the physical adapters and existing VMkernel will be migrated to the Distributed Switch.




Ensure you select the proper vmnic and use the Assign uplink button to associate it to one of the LAG ports.



Repeat this step with other vmnics being added to the LAG.




Migrate any VMkernel network adapters from the standard vSwitch to the corresponding dvSwitch port group.



Finish the migration wizard by allowing vCenter to check for potential conflicts and clicking OK to migrate the host.

After the migration is complete, you can examine the host to see how the vmnics are connected to the Link Aggregation group.



Note: Depending on the number of adapters created in OneView, and the number of network sets it may be necessary to create multiple dvSwitches, one for each NIC pair in OneView.

When the Distributed vSwitch with LACP is applied to a host, the errors in the OneView profile will automatically clear.

▼ ♦ **An error has occurred on connection 1. Interconnect BOT-CN75150107, interconnect 3 port 19 subport c is not an active member of the link aggregation group LAG1.** 2/23/18 2:17:46 pm Locked unassigned ▼
33 minutes ago

Resolution Verify the switch-assisted load-balancing using 802.1AD/LACP based NIC teaming is configured on the server MID-CN751500Z9_bay 7. If the problem persists, please contact your authorized support representative and provide them with a support dump.

Notes

Write a note

Health category Connection

► Event details

To verify the LAG status in VMware the following ESXCLI command can be run.

```
# esxcli network vswitch dvs vmware lacp status get
```

When the LACP LAG is properly formed, the ESXCLI will show a state of Bundled.

```
[root@ESX-1:~] esxcli network vswitch dvs vmware lacp status get
Management-Nexus__VDS
DVSswitch: Management-Nexus__VDS
Flags: S - Device is sending Slow LACPDUs, F - Device is sending Fast LACPDUs
- Device is in passive mode
LAGID: 308676574
Mode: Passive
Nic List:
  Local Information:
    Admin Key: 1
    Flags: SP
    Oper Key: 1
    Port Number: 32774
    Port Priority: 255
    Port State: AGG, SYN, COL, DIST,
    Nic: vmnic6
  Partner Information:
    Age: 00:00:21
    Device ID: 50:65:f3:df:4f:de
    Flags: SA
    Oper Key: 1252
    Port Number: 879
    Port Priority: 128
    Port State: ACT, AGG, SYN, COL, DIST,
    State: Bundled

  Local Information:
    Admin Key: 1
    Flags: SP
    Oper Key: 1
    Port Number: 32775
    Port Priority: 255
    Port State: AGG, SYN, COL, DIST,
    Nic: vmnic7
  Partner Information:
    Age: 00:00:21
    Device ID: 50:65:f3:df:4f:de
    Flags: SA
    Oper Key: 1252
    Port Number: 2671
    Port Priority: 128
    Port State: ACT, AGG, SYN, COL, DIST,
    State: Bundled
```

Lifecycle Management

Maintaining infrastructure, its firmware, and hypervisor operating systems can be a challenge. VM density and shrinking maintenance windows requires vigilant planning and coordination. HPE has incorporated key features into HPE Synergy and to its infrastructure components to facilitate the update process.




Firmware management

HPE Synergy Composer powered by HPE OneView provides fast, reliable, and simplified firmware and driver management across many HPE Synergy components. HPE OneView manages firmware to reduce manual interactions and errors, in addition to minimizing downtime. Firmware updates of management appliances and shared infrastructure are non-disruptive to the production workload.

An HPE Synergy Custom SPP (Support Pack for ProLiant) bundle is a comprehensive collection of firmware and system software components. Using the firmware management features built in to the appliance, you can define firmware baselines and perform firmware updates across many resources. HPE Synergy Custom SPP Bundles enable you to update firmware on HPE ProLiant servers, controllers, storage, and enclosures.


HPE Synergy provides policy-based firmware management including non-disruptive firmware updates to shared infrastructure components. Future enhancements will include a new template specifically defined for hypervisor cluster environments and features to allow for rolling upgrades within the cluster.

Best Practice: Establish firmware baselines in the SPT. Keep the HPE Synergy firmware up to date with minimal disruption with features like frictionless firmware updates (exclusive to HPE Synergy).



Health Monitoring

HPE Synergy is managed by HPE OneView, an intuitive infrastructure management tool that manages compute, storage, and networking through software defined intelligence. HPE OneView provides proactive health monitoring across the datacenter using its innovative dashboard view. The HPE OneView Dashboard provides an at-a-glance view of your datacenter's health and capacity. This view can be customized for each user when they login. SNMP alerts can be sent up to an enterprise monitoring system as well as SMTP alerts can be filtered and then sent to individual users or teams. HPE provides a free to use product called Global Dashboard which provides a consolidated view of your HPE OneView managed resources.



HPE OneView Dashboard

VMware vCenter Integrations:


- Comprehensive system health and alerting in VMware vCenter
- Deep levels of HPE server, blade, firmware, and infrastructure inventory and management
- End-to-end view of the HPE networking environment
- Monitor and manage health, configuration, and capacity of HPE Storage

VMware vRealize Operations Manager Integrations:

- Provides integrated and highly automated performance, capacity, configuration compliance, and cost management tools to the VMware vRealize Operations Manager custom graphical user interface (GUI)
- Uses the VMware analytics engine to analyze what is normal, and then applies that baseline to a dynamic server environment

HPE OneView for VMware vCenter Integration

HPE OneView for VMware vCenter is software for the VMware vCenter management console and enables the VMware vSphere administrator to quickly obtain context-aware information about HPE servers and HPE storage in their VMware vSphere environment directly from within VMware vCenter. This enables the VMware vSphere administrator to easily manage physical servers and storage, data stores and virtual machines, all from a single familiar console. By providing the ability to clearly view and directly manage relationships between virtual machines, and HPE Infrastructure, the VMware administrator's productivity increases, as does the ability to ensure quality of service.



Features and Benefits

- Simplify administration through integration of the physical and virtual infrastructure.
- Fast problem identification by getting hardware events into the VMware vSphere management console.
- Easier control by launching your trusted HPE management tools with a single click from the VMware vSphere dashboard
- Proactively manage or view changes with detailed relationship dashboards of server, networking and storage
- Maintain stability and reliability with firmware inventory and deployment
- Simple on demand server and storage capacity provisioning
- Automated deployment of HPE StoreOnce VSA, enables an industry leading secure backup and recovery solution in your VMware environment
- Visualize complex configuration relationships

HPE OneView for VMware vCenter 9.0 includes the following updates:

General

- Migrated to HTML5 vSphere Client/Web Client for User Interface
- Support for HPE OneView 4.0.0
- Edit option for HPE OneView Controllers for the user role
- Edit deployment plan supported for Image Streamer.
- HPE OneView for VMware vCenter for Server


General

Added vCenter Pro-active High Availability provider that will allow you to evacuate VMs from HPE OneView, managed host that have reported hardware alerts.

- HPE Insight Control server provisioning credentials enhancements.

New hardware support

- Supports HPE Synergy D3940 Storage Module
- Supports HPE Synergy 12Gb SAS Connection Module



Licensing


Storage management with HPE OneView for VMware vCenter can be used at no charge. Licensing is enabled as part of the download and installation sequence.

Server management with HPE OneView for VMware vCenter is licensed with HPE OneView and HPE Insight Control. New advanced provisioning capabilities are available for HPE OneView licensed hosts only.

Best Practice: Use this plug-in in VMware vCenter installations.

Appendix A - HPE Synergy Architecture Overview

The HPE Synergy Architecture is the next generation converged infrastructure, specifically engineered for hybrid IT environments. HPE Synergy provides the resiliency, scale, and features needed for enterprise-level virtualization clusters.



Feature Summary:

- New interconnects and converged network adapters provides twice the downlink bandwidth @ 50Mbps. Uplink TOR ports speeds have more than doubled to 100Mbps.
- HPE Synergy's unique scale-out fabric allows exponential infrastructure scale in a matter of minutes, not days or weeks
- Out-of-band, fully-redundant management network providing hardware auto-discovery, health, and fabric topology validation
- HPE Synergy multi-frame fabric expansion provides high availability with wire-speed scale
- A variety of shared storage options including VMware vSAN utilizing the HPE Synergy D3940 Storage Modules and/or fully integrated HPE 3PAR StorServ storage systems
- Storage I/O options including FCoE, Native Fibre-Channel, iSCSI, and Flat-SAN
- Non-disruptive firmware updates to shared infrastructure components like FC switches and HPE Virtual Connect
- High availability network features including active/active MC-LAG and S-channels
- New policy based firmware management, embedded into HPE OneView to efficiently manage firmware in a cluster environment.
- Server Profile Templates ensure consistency across the enterprise through standardization and compliance
- Fully automated fabric zoning and SAN volume provisioning (HPE 3PAR StorServ)
- Multi-tenancy HPE 3PAR StorServ storage capabilities including domains (segmentation) and storage pools (quotas)
- Powerful unified API to a growing partner ecosystem including Arista, Capgemini, Chef, Docker, Microsoft, NVIDIA, and VMware

Appendix B - HPE Synergy Configuration Overview

Implementing VMware vSphere best practices must first start with understanding the process of powering up HPE Synergy for the first time, hardware discovery, and applying your best practice (through templates) to configure HPE Synergy for use.


The following sequence is an example of bringing up a new HPE Synergy environment and the major tasks involved in the process. Detailed installation and configuration steps are in the HPE Synergy Product Documentation. This summarized list is included for reference in the following best practices section.

Installation sequence example with HPE Synergy and Fibre-Channel/3PAR StoreServ.




1. HPE Synergy rack installation and power
2. Cabling HPE Synergy out-of-band management network to TOR switch
3. Cabling HPE Synergy interconnect modules
 - a. Master module IRF module interconnects
 - b. Network/Fabric TOR uplinks
 - c. Optional Image Streamer uplinks
4. HPE Synergy discovery and First-Time Setup (via console)
5. Create address pool for HPE Synergy Management (and Image Streamer if applicable)
6. Create Networks and Network Sets in the HPE Synergy Composer
 - a. Includes Ethernet and FC fabrics
7. Create a template (LIG) for each fabric type
8. Create a template (EG) for the enclosure group
9. Create the logical enclosure (LE)
10. Optionally perform a big-bang firmware update
11. Add SAN manager and map discovered fabrics to HPE OneView SAN networks
12. Add HPE 3PAR StorServ Storage Systems and import storage pools (CPGs)
13. Create volume templates for the shared permanent volumes
14. Create HPE 3PAR StorServ shared volumes (datastores) from volume templates
15. Create Server Profile Templates (SPT) and Server Profiles (SP)

The following diagram maps the relationships of HPE OneView templates to logical instantiations which in turn is mapped back to physical infrastructure. The process starts by creating best practice configuration using HPE OneView templates. Once the fabric (LIG) and enclosure templates (EG) are defined, the final step is the creation of the Logical Enclosure (LE) which configures HPE Synergy into a managed state.



The left side of the diagram illustrates the templates defined in the HPE Synergy Composer. The HPE OneView templates and logical instances on this relationship diagram are accessible from the HPE OneView main pulldown.



HPE OneView main menu pulldown

Appendix C - Referenced HPE and VMware Best Practice Documents

The following technical whitepapers were referenced in the *Deploying VMware vSphere on HPE Synergy Best Practices Guide*. The best practices provided in this document are considered *generalized* best practices for hypervisor environments. Many customer specific requirements must be considered for properly designing virtualization clusters including capacity, scalability, performance, workload, and availability.

HPE Synergy Related Best Practices

- [HPE Synergy HPE OneView 4.0 Support Matrix](#)
- [HPE Synergy with VMware vSAN and D3940](#)
- [HPE Synergy UEFI Users Guide](#)
- [HPE Synergy Image Streamer POC Cookbook](#)
- [HPE Synergy 40Gb F8 Switch Module VLAN Configuration](#)
- [Best Practices for HPE Synergy Firmware and Driver Updates](#)
- [Implementing Multi-Chassis LAG and HPEN FlexFabric 59xx Series](#)
- [Implementing Multi-Chassis LAG and Arista 7000 Series Switches](#)
- [Implementing Multi-Chassis LAG and Cisco Nexus 5000 Series Switches](#)

VMware Best Practices

- [VMware vSphere Availability](#)
- [VMware vSphere 6.5 Performance Best Practices](#)
- [VMware vCenter Server HA and Performance Best Practices](#)
- [VMware vSphere 6.5 Configuration Maximums](#)
- [VMware vSphere Networking 6.5](#)

HPE 3PAR StorServ Best Practices

- [HPE 3PAR StorServ ESXi 6.5 Best Practices](#)

© Copyright 2017 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.

This document contains confidential and/or legally privileged information. It is intended for Hewlett Packard Enterprise and Channel Partner Internal Use only. If you are not an intended recipient as identified on the front cover of this document, you are strictly prohibited from reviewing, redistributing, disseminating, or in any other way using or relying on the contents of this document.

4AA4-xxxxENW, December 2017