HPE Storage

Technical white paper

HPE GreenLake for Block Storage and VMware vSphere 7 deployment considerations

For HPE Alletra Storage MP, HPE Alletra 9000 Storage, and HPE Primera Storage

HPE GreenLake

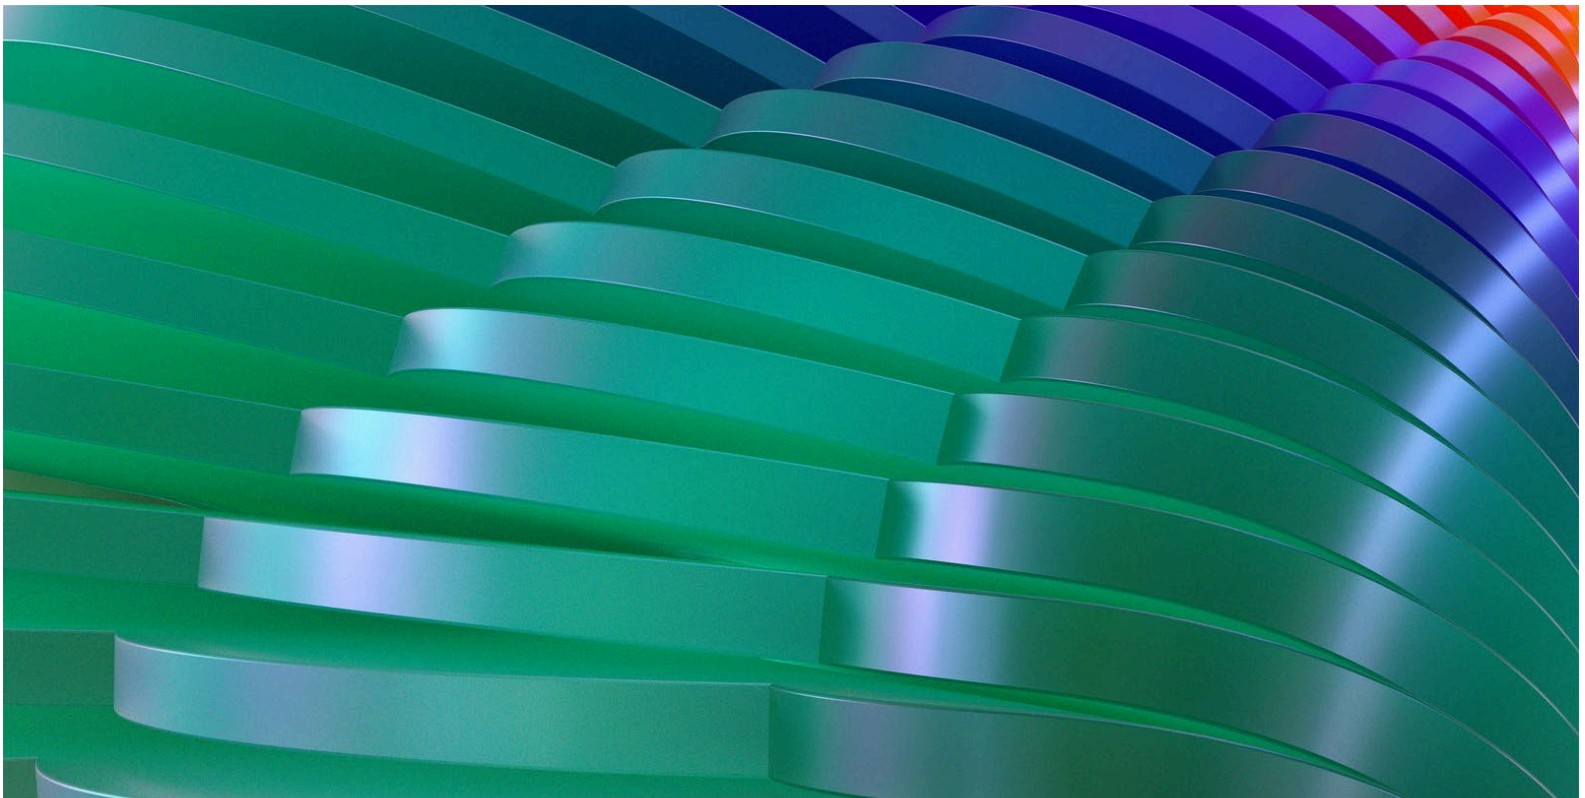

Contents

Executive summary

VMware vSphere® and HPE Storage deliver an elite solution that combines server virtualization with array-based features and optimization for an integrated solution. Complementary features inherent in both platforms increase overall utilization, provisioning agility, and administrative efficiency.

HPE GreenLake for Block Storage built on HPE Alletra Storage MP, HPE Alletra 9000 Storage, and HPE Primera Storage arrays are the next generation of Tier 0 storage platforms that serve mission-critical workloads, virtualization, and cloud-computing environments by providing the performance, scalability, availability, and simplified management that clients demand. They achieve this excellence through an innovative system architecture that offers active-active controller nodes, high-speed full-mesh interconnects, system-wide striping, a services-centric OS, and support for NVMe — among other features. For more information, see Technical deep dive: HPE GreenLake for Block Storage built on HPE Alletra Storage MP.

This paper provides guidance and information for deploying VMware vSphere 7 update 3 and later with HPE storage arrays mentioned above, hereafter referred to as "HPE Alletra MP/9000 or HPE Primera." The intended audience includes solution architects, vSphere administrators, and storage administrators.

HPE Storage array management

Multiple management tools are available for HPE Storage arrays. For HPE Alletra MP/9000 or HPE Primera arrays, the Data Services Cloud Console (DSCC) is a secure cloud-based infrastructure and data management platform. It offers numerous microservices, such as: initial setup, volume and host provisioning, dashboards for capacity and performance, event, and log views. Other management options include the on-board UI and the CLI.

Note
HPE Alletra arrays primarily use the DSCC for management. They are not supported by the HPE SSMC, so they cannot be managed through it. HPE Primera arrays can be managed through the DSCC or the HPE SSMC.

VMware host definition

To manage servers running VMware vSphere on HPE Storage arrays, you must first create a host definition entry that specifies a valid Host Persona for each VMware ESXi™ host connected to the HPE Storage system, and the storage protocol used (FC, iSCSI, NVMe, or direct connection). **Host Persona 11** is the required setting for any VMware ESXi host that is connected to HPE Alletra MP/9000 or HPE Primera arrays. In the DSCC, you can do this by selecting **VMware (ESXi)** in the **Create Host** OS drop-down menu, as shown in Figure 1.

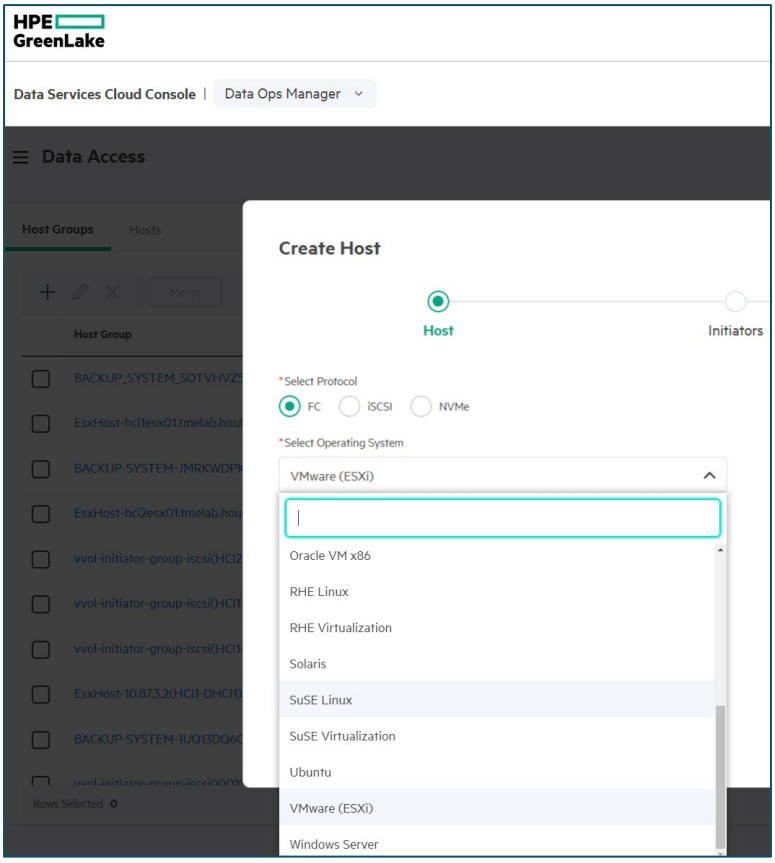

Figure 1. Host OS VMware (ESXi) persona in DSCC

In the CLI, the command has the following format:

```
cli% createhost -persona 11 <ESXi host name> <WWPN1> <WWPN2> … <WWPNn>
```

For example, the following command creates a host named "DL360G10-VMW" with the two host HBAs WWPNs:

```
cli% createhost -persona 11 DL360G10-VMW 50014380231F9AA5 50014380231F9AA7
```

Adaptive Queue Depth Throttling

Each HPE Alletra MP/9000 or HPE Primera port has a finite queue depth, and the I/O queues are shared among the connected host HBA ports on a first-come, first-served basis. If an ESXi host sends an I/O request to an HPE Alletra MP/9000 or HPE Primera port whose queue is full, the ESXi host receives a **queue full** SCSI response from the HPE Alletra MP/9000 or HPE Primera array. The ESXi host's default reaction to this response is to recognize it as a valid command and continue sending I/O requests to that port. When other operating systems receive a **queue full** response, they back off or stop sending I/O, but the ESXi host does not interpret the **queue full** response as an indication to curb its I/O. I/O commands sent to an HPE Alletra MP/9000 or HPE Primera port that has reached its maximum queue depth are not processed beyond the initial **queue full** SCSI response. The overflow of requests has the potential to cause erratic performance and result in VMs becoming unresponsive, which might lead to a crash of the ESXi host. The strategy for avoiding this condition is to enable an Adaptive Queue Depth Throttling algorithm introduced by VMware® (see VMware KB 1008113).

The algorithm adjusts the LUN queue depth in the VMkernel I/O stack. It is activated when the storage array indicates I/O congestion by returning a **queue full** SCSI status. When congestion is detected, the VMkernel throttles the LUN queue depth and attempts to gradually restore the queue depth after congestion conditions subside.

Standard VMFS volumes have a default queue depth of 32, whereas VMware vSphere® Virtual Volumes™ (hereafter referred to as "vVols") default to 128 for the protocol endpoint. Without Adaptive Queue Depth Throttling, administrators are forced to manually limit the number of VMs per ESXi host to reduce the risk associated with any particular VM overrunning I/O queues. Administrators are also forced to manually tune the number of VMs if they detect congestion — a reactive, slow, and error-prone process. By automating congestion control, administrators can confidently create a higher number of VMs per ESXi host without the need for manual congestion control.

To enable queue depth throttling through the vSphere ESXi Shell, set the **QFullSampleSize** and **QFullThreshold** values:

```
# esxcli system settings advanced set -o /Disk/QFullSampleSize -i 32
```

```
# esxcli system settings advanced set -o /Disk/QFullThreshold -i 8
```

These parameters can be set from the vSphere Client through the **Configure** tab for each ESXi host (see Figure 2). In the **Configure** tab, select **Advanced System Settings**, and set **Disk.QFullSampleSize** to 32 and **Disk.QFullThreshold** to 8.

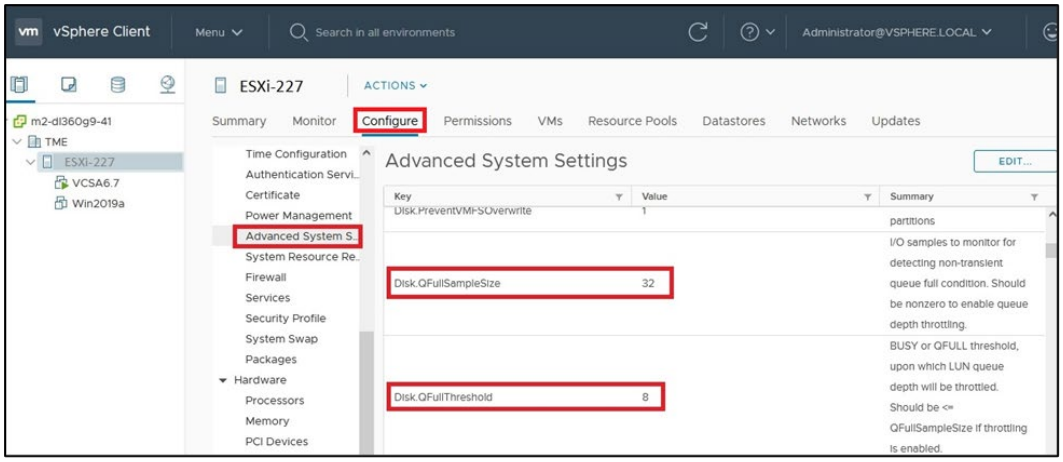

Figure 2. Queue depth throttling settings in vSphere Client

A setting of **QFullSampleSize = 0** turns off the throttling algorithm. With the algorithm turned off, the ESXi host does not adjust I/O load to **queue full** responses from the storage array.

Adaptive Queue Depth Throttling can be used in conjunction with the Storage I/O Control (SIOC) feature. If Adaptive Queue Depth Throttling is enabled, it is important to enable it for all hosts that access the HPE Alletra MP/9000 or HPE Primera array.

Multipath settings

To maintain a constant connection between an ESXi host and its storage, vSphere supports multipathing. vSphere native multipathing has two important plug-ins:

- **Storage array type plug-in (SATP)**, which manages path failover and monitors path health.
- **Path selection plug-in (PSP)**, which defines the path selection policy and selects a path for each scheduled I/O according to the algorithm defined by the PSP policy.

A PSP policy of round robin is the preferred multipath implementation for ESXi, as shown in Figure 3.

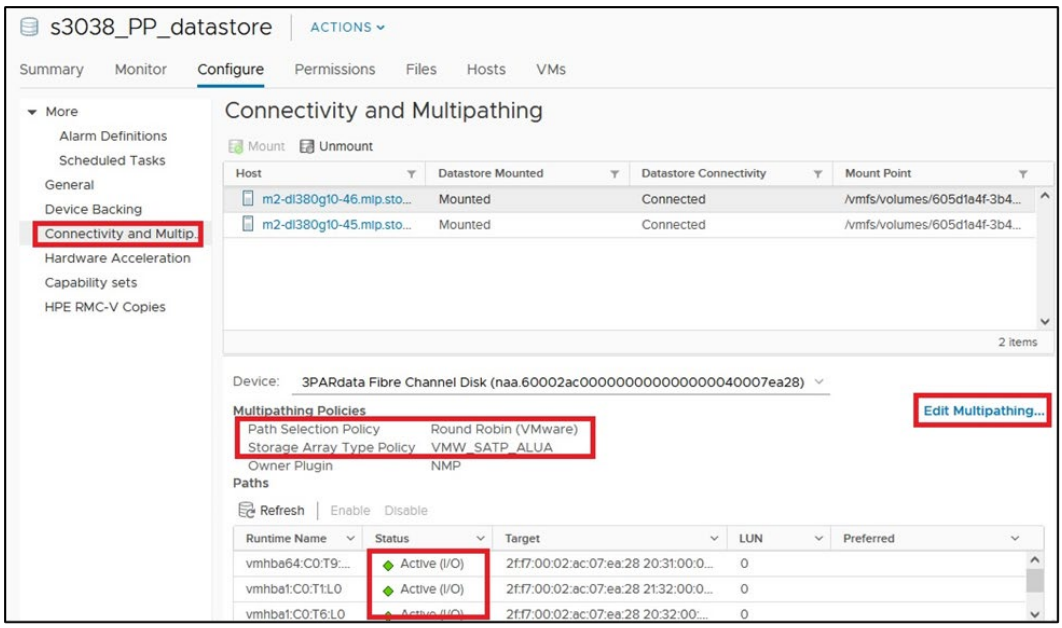

Figure 3. Multipathing settings in vSphere Client

For manual configuration of datastores, you can edit the multipathing settings for each datastore and change the PSP to **Round Robin (VMware)**. For automatic configuration of LUNs, ESXi multipathing for round robin must be configured by using a custom SATP PSP rule through the `esxcli` command.

For storage systems with HPE Alletra MP/9000 or HPE Primera, SATP claim rules are defined using the **3PARdata** identifier.

For ESXi 6.5 through ESXi 6.7 U1, use the following command to set the ALUA custom SATP rule for round-robin multipath:

```
# esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -P "VMW_PSP_RR" -O "iops=1" -c "tpgs_on" -V
"3PARdata" -M "VV" -e "HPE Custom iSCSI/FC/FCoE ALUA Rule"
```

For ESXi 6.7 U2 through ESXi 7.0 U1, set the ALUA custom SATP rule that incorporates IOPS and latency-based sub policy for round-robin multipath:

```
# esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -P "VMW_PSP_RR" -o "throttle_sll" -O
"policy=latency" -c "tpgs_on" -V "3PARdata" -M "VV" -e "HPE Custom iSCSI/FC/FCoE ALUA Rule"
```

For ESXi 7.0 U2 and later, HPE recommends the default **SATP/PSP VMW_SATP_ALUA** rule for HPE 3PAR data virtual volume (VV) devices incorporated in the release distribution as of ESXi 7.0U2 as sufficient for multipath configuration. This rule applies to HPE Alletra MP/9000 or HPE Primera.

The ESXi host must be rebooted for SATP rule creation, changes, additions, or removals to take effect.

Storage I/O Control

Storage I/O Control (SIOC) enables you to throttle VMs that are accessing a datastore that has exceeded a specified latency. SIOC controls I/O on a datastore or VM level. This means that when a specified threshold is reached, SIOC determines on a datastore-wide level which hosts, and thus which VMs, receive a specific amount of the resources. SIOC prevents a single VM from claiming all I/O resources for a datastore in a cluster.

SIOC provides finer-grained control than Adaptive Queue Depth Throttling provides; however, the two features may be used together. SIOC reacts to latency increases for a datastore, whereas Adaptive Queue Depth Throttling responds to **queue full** and **device busy** messages from the HPE Alletra MP/9000 or HPE Primera array.

SIOC identifies high storage latency at the VMFS level and takes corrective action to protect high-priority VMs from noisy neighbors.

Priority Optimization and Intelligent Quality of Service

HPE Alletra MP/9000 and HPE Primera arrays have a manual quality-of-service (QoS) feature called **Priority Optimization** that enables you to set latency, IOPS, and throughput goals for a set of volumes. For instance, you could give production volumes a higher priority than development volumes. Priority Optimization places limits on I/O requests with lower-priority policies to help workloads with higher priority achieve their performance targets.

Although QoS policies have always been available for VMFS datastores, QoS policies can now also be set on vVols[1], either through the Storage Policy-Based Management (SPBM) capabilities in VMware vCenter®, as shown in Figure 4, or by using the HPE Storage Integration Pack for VMware vCenter.

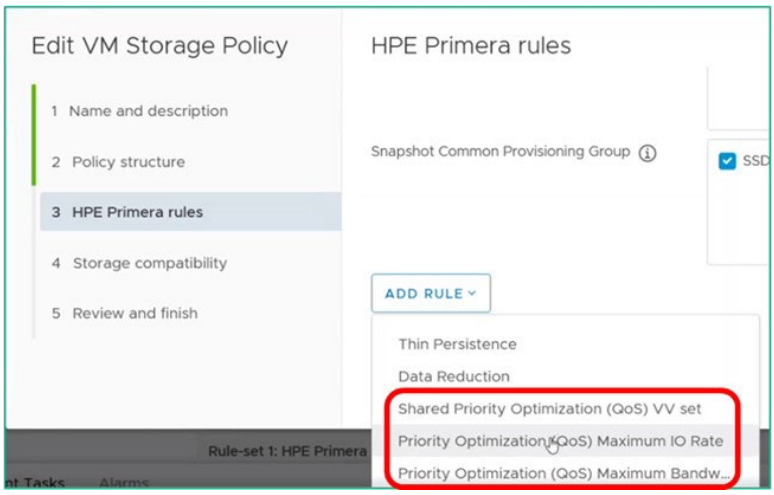

Figure 4. QoS for vVols in vSphere Client

Priority Optimization is a very sophisticated feature that has helped many service providers achieve their SLAs. However, some customers do not want or need such granularity. They simply want to make sure that low-priority workloads do not become noisy neighbors and affect more important workloads. Intelligent QoS enables that goal.

Intelligent QoS automatically applies QoS policies on the bully volumes that the user has tagged as medium-priority or lower-priority workloads. This automation uses machine-learned AI models in the HPE Alletra MP/9000 or HPE Primera OS to eliminate performance problems that higher-priority workloads would otherwise have experienced.

Note
Priority Optimization and Intelligent QoS are mutually exclusive. If Priority Optimization is configured on the array, Intelligent QoS cannot be enabled.

You can combine Priority Optimization or Intelligent QoS with SIOC and Adaptive Queue Depth Throttling to provide I/O congestion control in three stages:

1. First, I/O requests observe the QoS limits that are set on the HPE Alletra MP/9000 or HPE Primera array.
2. Next, I/O requests are affected when SIOC reacts to latency increases.
3. Finally, I/O requests react when Adaptive Queue Depth Throttling responds to **queue full** messages from QoS.

Coordination and some experimentation might be necessary to achieve the best results.

[1] VMware vSphere Virtual Volumes (vVols) support for HPE GreenLake for Block Storage built on HPE Alletra Storage MP is pending.

Storage DRS

Storage DRS (SDRS) is a datastore cluster object that gives users the capability to automate the provisioning and maintenance of VMs. SDRS provides both initial and ongoing placement recommendations for VMDKs as well as automated or manual load balancing of VMDKs, based on I/O latency and capacity utilization metrics.

NOTE
Because SDRS operates on datastore objects, it does not manage vVols, which are managed at the per-VM level using storage policies offloaded to the storage array. Thus, it is unsupported to manage a vVol datastore with an SDRS cluster.

When SDRS migrates a VMDK, I/O decreases on the original datastore, but increases on the target. For that reason, combining HPE Priority Optimization with SDRS requires careful QoS rule planning.

The QoS rule for any potential SDRS destination datastore should include enough headroom to accommodate the additional I/O capacity of the migrated VMDK. This necessity can lead to a managerial challenge in larger ESXi clusters, but the challenge can be addressed by defining a set of affinity rules for each VM to reduce its number of target datastores during an SDRS migration.

If the workload I/O characteristics of a VM are not well known, HPE recommends that you use manual SDRS migrations instead of automated ones to help detect any I/O congestion that might occur after the migration process. If you use SDRS to move datastores, HPE recommends not setting a QoS latency goal on the VVset that holds the datastore. Table 1 lists characteristics of the various control mechanisms.

Table 1. Characteristics of various resource control mechanisms

	HPE priority optimization	vSphere SIOC	Adaptive Queue Depth Throttling	vSphere SDRS
I/O control technique	• Set minimum goal and maximum limit for IOPS and bandwidth • Set latency goal • Set priority level	Enforce predefined I/O shares for each VM	Control queue depth of SAN LUN in VMkernel	Migrate VM to other datastore
Reacting on	• System-Busy level growing • Maximum limit reached	I/O latency growing	Queue Full or Device Busy at SAN LUN or port level	I/O latency and space utilization growing
Granularity	• VVset • Virtual domain • System	All VMs in a datastore	All hosts using the SAN LUN for the datastore or a particular port on the HPE array	Single VM
Managed from	• HPE DSCC • HPE CLI • HPE API • HPE SSMC (HPE Primera only)	• vSphere • CLI	• vSphere • CLI	• vSphere • CLI

Virtual storage controllers

vSphere provides several virtual storage controllers, with the default depending upon the guest OS and virtual hardware version. There are four types of storage controllers for VMs:

- SCSI (BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, VMware Paravirtual SCSI)
- AHCI SATA
- IDE
- NVMe

For more information, see SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility.

For virtual SCSI controller types, shown in Figure 5, HPE recommends using the VMware Paravirtual SCSI adapter (PVSCSI) for optimal performance in storage that is connected by a SAS bus. Environments with I/O-intensive applications benefit the most from this adapter because it significantly reduces CPU utilization and provides potential for increased throughput as compared to the default storage controllers. For more information about configuring the PVSCSI adapter and a list of guest operating systems that support it, see VMware KB 1010398: Configuring disks to use VMware Paravirtual SCSI (PVSCSI) controllers.

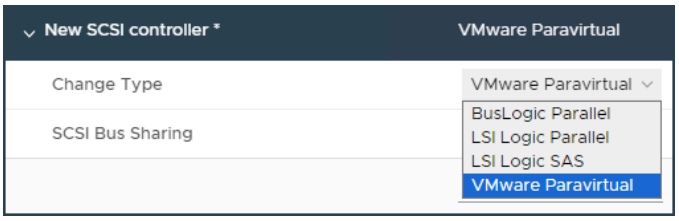

Figure 5. SCSI controller types in vSphere Client

NVMe virtual controllers, shown in Figure 6, are designed to work with all-flash storage that is connected through a PCIe bus instead of a SAS bus. NVMe virtual controllers significantly reduce the overhead for processing guest OS I/O, compared to AHCI SATA or SCSI controllers. For more information about using NVMe controllers, see VMware KB 2147714: Using Virtual NVMe from ESXi 6.5 and virtual machine Hardware Version 13.

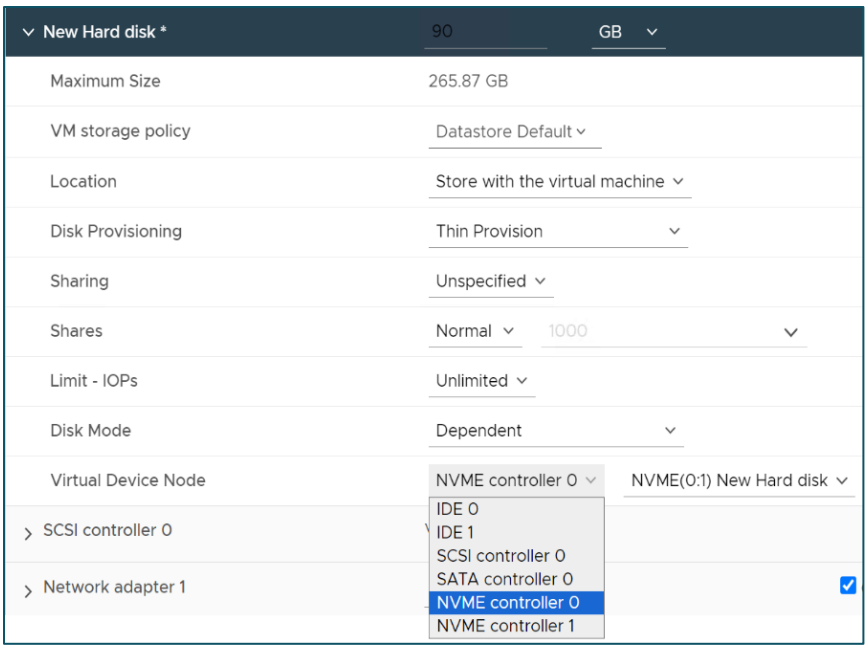

Figure 6. Virtual controller types in vSphere Client

HPE Alletra MP/9000 or HPE Primera volume types

The HPE Alletra MP/9000 or HPE Primera arrays offer two choices for volume types:

- Thin-provisioned virtual volumes (TPVV)
- Thin provisioned with data reduction virtual volumes (TDVV)

On the HPE Alletra MP/9000 or HPE Primera array, start with TPVV for greatest performance. If the workload of the VMs on the datastore is conducive to data reduction (deduplication and compression), you can always tune the volume online.

vSphere Storage APIs for Array Integration

The vSphere Storage APIs for Array Integration (VAAI) enable vSphere to leverage storage resources to improve performance, resource utilization, and scalability by leveraging more efficient array-based operations to perform VMware I/O. Common operations that are offloaded to the array include the following:

- **Full Copy/XCOPY** to copy or migrate data within the same array
- **Block Zero/Write Same** to zero-out disk space
- **Block Delete** to reclaim space by using the SCSI UNMAP feature
- **Atomic Test & Set (ATS)** for locking of files on the VMFS datastore

VAAI primitives are controlled by these advanced settings:

- HardwareAcceleratedLocking
- HardwareAcceleratedMove
- HardwareAcceleratedInit

For more information about VAAI, including how to check/enable/disable it, see VMware KB 1021976: Frequently Asked Questions for vStorage APIs for Array Integration.

Of all these VAAI primitives, ATS requires particular attention. ATS usage can be disabled at different levels:

- Disable ATS usage per ESXi host (which disables ATS usage for all VMFS datastores on the given ESXi host)
- Disable ATS usage per VMFS datastore (which disables ATS usage for all ESXi hosts that are using a given VMFS datastore)
- Disable ATS usage only for VMFS heartbeats

Regardless of your ESXi version, HPE does not recommend disabling ATS heartbeat unless you see the following message in the `vmkernel.log` file:

"`ATS miscompare detected between test and set HB images at offset XXX on vol YYY`"

If you see this ATS miscompare message, then disable ATS usage only for VMFS heartbeats. Leave the other ATS settings at the host and datastore level unchanged. To disable ATS for VMFS heartbeats, use the following command:

```
# esxcli system settings advanced set -i 0 -o /VMFS3/UseATSForHBOnVMFS5
```

Disabling ATS usage only for VMFS heartbeats disables the heartbeat that was implemented with VAAI. The result is that the ESXi host uses plain SCSI reads and writes to update its heartbeat on VMFS datastores.

VMware vSphere Virtual Volumes

Note
VMware vSphere Virtual Volumes (vVols) support for HPE GreenLake for Block Storage built on HPE Alletra Storage MP is pending.

Hewlett Packard Enterprise has partnered with VMware to define, develop, and test vVols, a new storage architecture that was introduced in vSphere 6. vVols can be used for VM storage, and they provide an alternative to VMFS datastores. Unlike VMs that share a VMFS datastore that maps to a LUN on the HPE Alletra MP/9000 or HPE Primera array, vVols VMs are made up of several volumes that map to separate LUNs. With vVols, each VMDK of the VM, the VM's configuration file, and its swap space, snapshots for disks and memory are all individual volumes that are stored natively on the HPE Alletra MP/9000 or HPE Primera array. The volumes that map to each vVols VM are created (and destroyed) on demand, with data operations offloaded to the array. As Figure 7 shows, vVols VMs have the advantage that policies can be applied to individual VMs instead of VMs inheriting whatever policy is set on the VMFS datastore. For example, vVol VM1 can have a QoS policy while vVol VM2 has a replication policy.

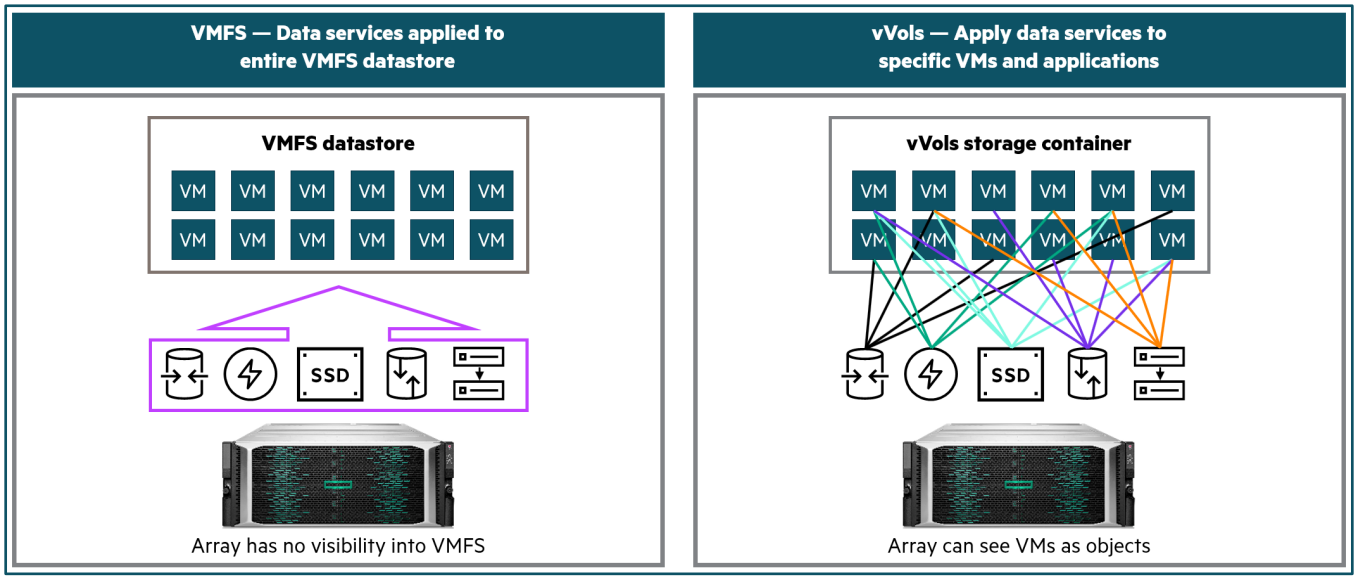

Figure 7. VMFS versus vVols

Using vVols with HPE Alletra MP/9000 or HPE Primera arrays changes the storage architecture of vSphere from a LUN-centric model to a VM-centric model and enables storage arrays to become VM-aware.

The following components make up vVols, as shown in Figure 8:

- **Storage container:** An organizational unit that does not pre-allocate storage space but provides a logical grouping of vVols and is seen as a virtual datastore by ESXi hosts.
- **Protocol endpoint:** A logical I/O proxy that serves as the data path between ESXi hosts to VMs and their respective vVols.
- **VASA Provider:** A software component that mediates out-of-band communication for vVol traffic between VMware vCenter Server®, the ESXi hosts, and the storage array.
- **Storage Policy-Based Manager (SPBM):** A set of rules that define storage requirements for VMs based on capabilities provided by the storage array.

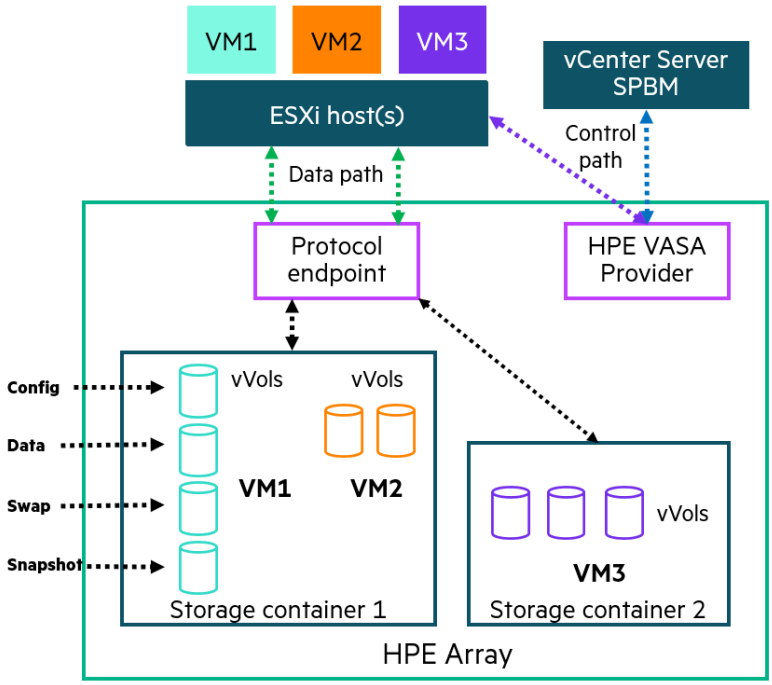

Figure 8. vVols architecture

For more information about vVols, see VMware vVols with HPE Alletra 9000 Storage or VMware vVols with HPE Primera Storage.

vSphere Metro Storage Cluster and HPE Peer Persistence

vSphere Metro Storage Cluster is a stretched storage cluster architecture that enables a vSphere cluster to span data centers, as shown in Figure 9. HPE Peer Persistence software enables a pair of HPE Storage arrays, located at metropolitan distances (using synchronous replication), to act as peers to each other and present a unified system to the ESXi hosts. Using HPE Peer Persistence enables you to configure a high-availability VMware storage solution between two sites in which failover and failback of access to the volumes across arrays remains transparent to the ESXi hosts, the VMs, and the applications running on those VMs.

The solution is designed to provide automatic and transparent failure recovery with a recovery point objective (RPO) of zero (no data loss) and a recovery time objective (RTO) of zero (no application outage) in the case of many failures.

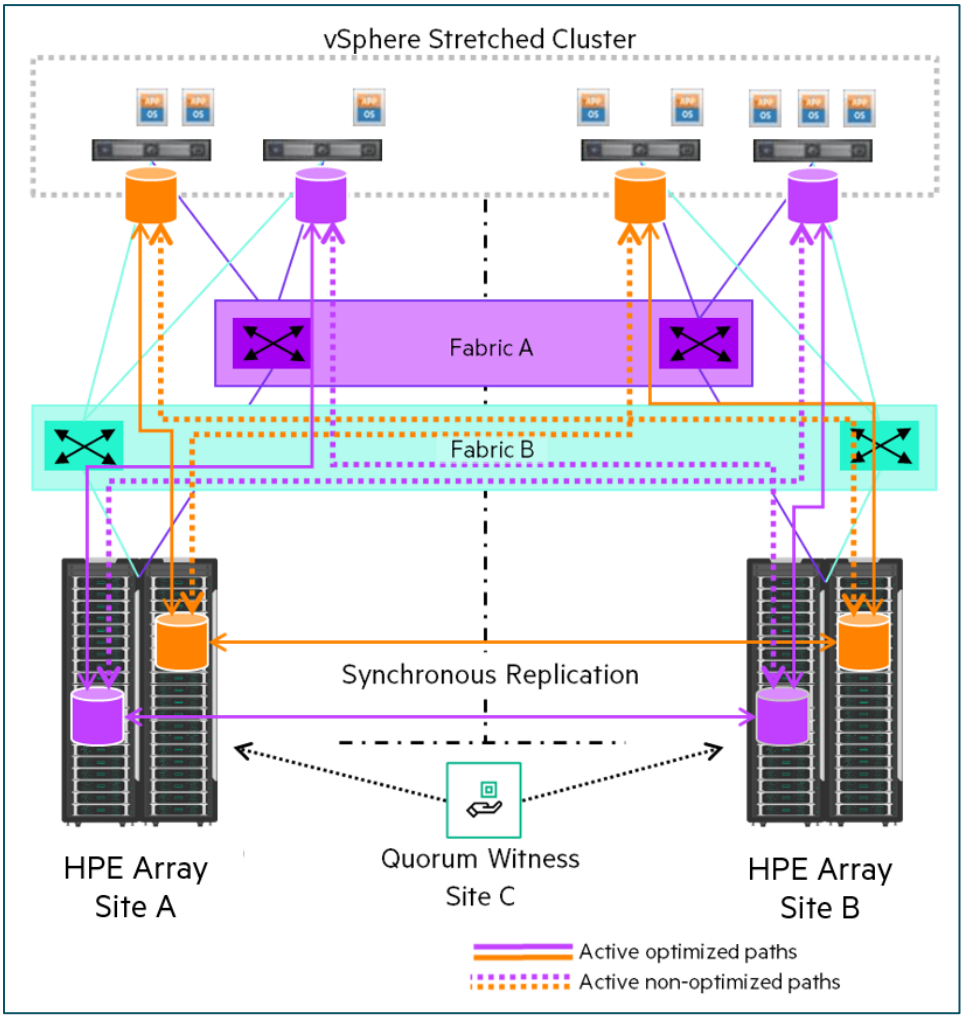

Figure 9. HPE Active Peer Persistence

For more information, see HPE Active Peer Persistence and VMware KB article 77061: Implementing vSphere Metro Storage Cluster (vMSC) using HPE Primera or Alletra 9000 Peer Persistence.

Note
vVols are not currently supported with HPE Peer Persistence.

Site Recovery Manager

VMware Site Recovery Manager™ is a business continuity and disaster recovery solution that integrates vSphere infrastructure with Remote Copy and Virtual Copy through the Site Recovery Manager adapter (SRA) from HPE. Unlike HPE Peer Persistence, which is supported only with synchronous replication and campus/metro distances, Site Recovery Manager is not limited to synchronous replication and can span wider areas.

Site Recovery Manager helps with planning, testing, and running the recovery of VMs between a protected vCenter Server site and a recovery vCenter Server site, as shown in Figure 10. It simplifies disaster recovery and increases reliability, and it is supported with traditional VMDKs or vVols.

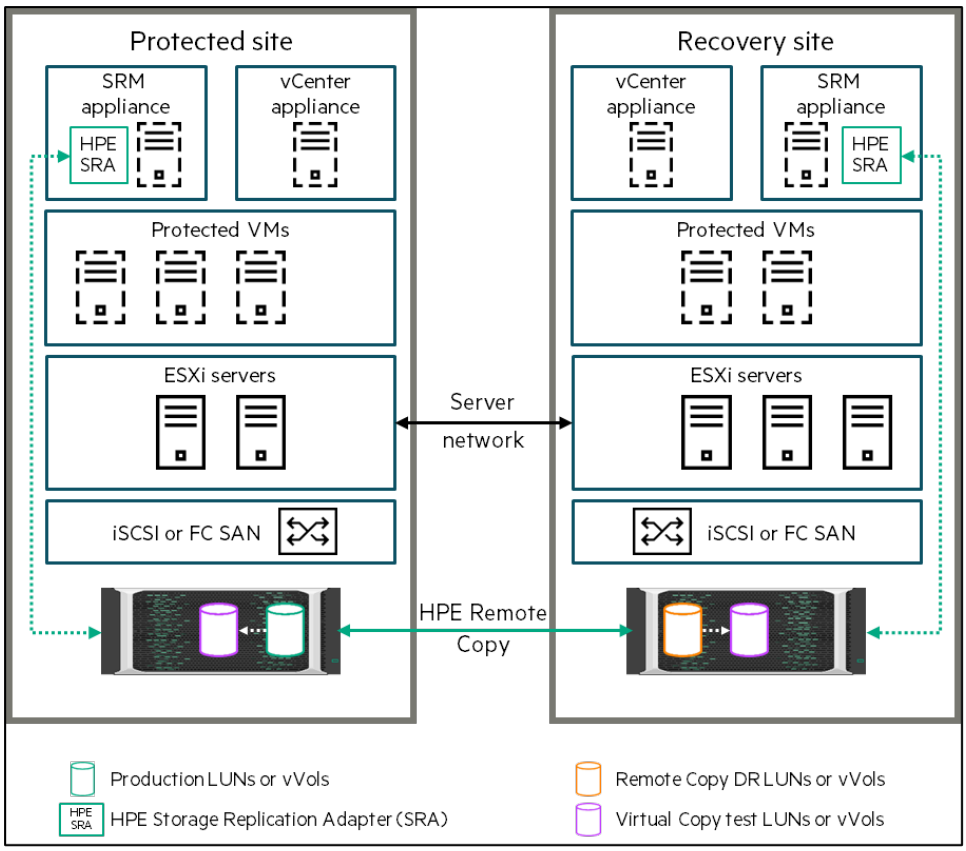

Figure 10. Site Recovery Manager

For more information about VMware Site Recovery Manager, see HPE Storage Replication Pack 8.4 for VMware Site Recovery Manager Virtual Appliance User Guide.

VMware Cloud Foundation

VMware Cloud Foundation™ (VCF) is a hybrid cloud platform that combines VMware software products (vSphere, VMware Aria (previously known as vRealize), VMware vSAN™, and VMware NSX®) into a centrally automated and managed platform: VMware EVO™ SDDC Manager™. EVO SDDC Manager automates the entire system lifecycle — from initial startup to configuration and provisioning, to upgrades and patching — and simplifies day-to-day management and operations.

VCF environments contain management domains and workload domains. A management domain is a special-purpose workload domain that is used to host the infrastructure components needed to manage VCF. It leverages vSAN for storage.

As Figure 11 shows, VCF has two types of storage options: principal storage and supplemental storage. Although vSAN is the only principal storage option for the management domain, the workload domains can use vSAN or supported principal storage. HPE Alletra Storage or HPE Primera Storage can be used for principal and supplemental storage through vVols or VMFS. Additional shared storage types can be added to a cluster in the management domain or a VI workload domain after it has been created. The additional supported shared storage options include:

- vSAN
- Fibre Channel (FC)
- iSCSI Network File System
- Network File System (NFS) protocol version 3 or 4.1
- VMware vSphere Virtual Volumes (vVols)

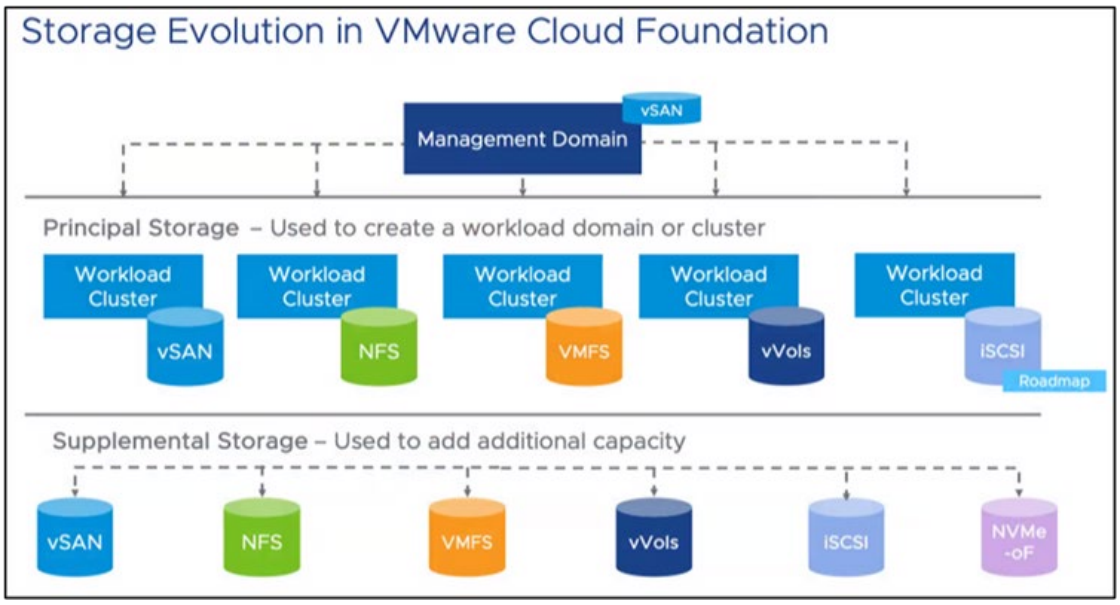

Figure 11. VCF storage options (graphic courtesy of VMware)

HPE plug-ins for VMware

HPE offers several free plug-ins that integrate with various VMware products. These integrations help VMware deployments with key functions:

- **Management** through HPE Storage Integration Pack for VMware vCenter (SIP4VC)
- **Monitoring** through HPE Storage Management Pack for VMware vRealize® Operations Manager™ (VMware Aria Operations)
- **Protection** through HPE Storage Replication Pack for VMware Site Recovery Manager™
- **Automation** through HPE Storage Automation Pack for VMware vRealize® Orchestrator™ (VMware Aria Automation Orchestrator)

The plug-ins can be downloaded from the HPE Free Software page.

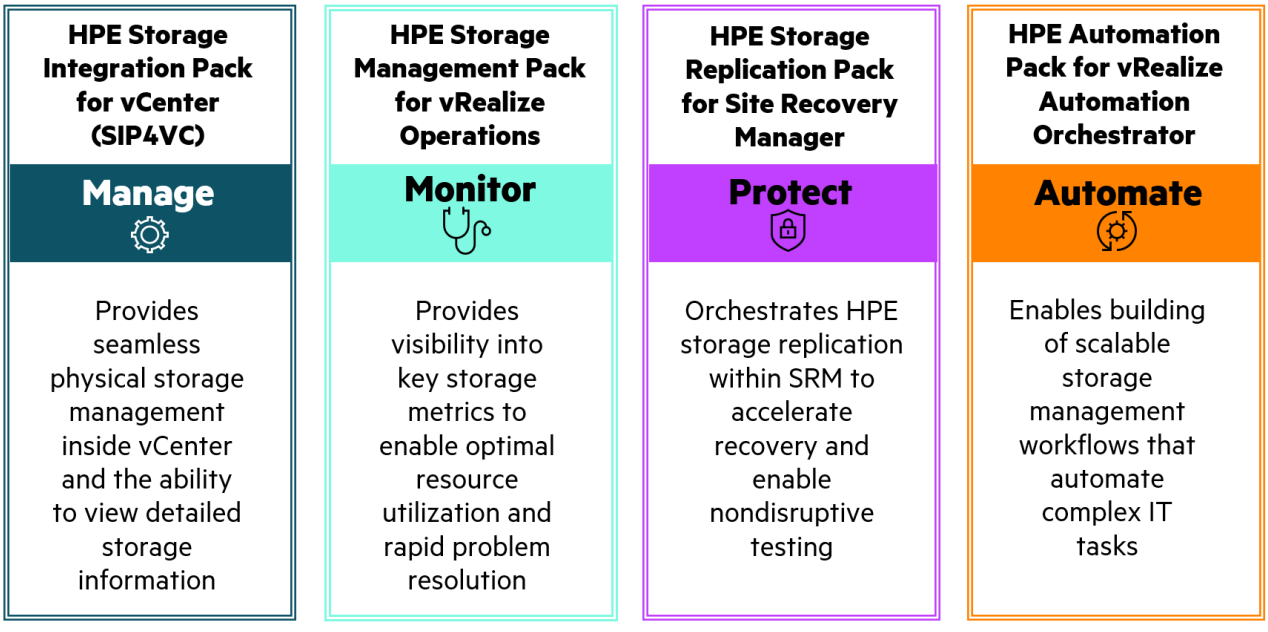

Figure 12. HPE value-added software for VMware

Summary

Deploying HPE Alletra MP/9000 or HPE Primera Storage in VMware vSphere environments helps avoid the issues associated with traditional storage platforms in relation to performance, efficiency, availability, and management. The built-in architecture differentiators of the HPE Storage platform — such as the ASIC, active-active controllers, full-mesh backplane, system-wide striping, thin provisioning, and so on — have a direct and positive impact on VMware deployments through features such as these:

- **Adaptive Queue Depth Throttling** provides automatic LUN queue depth adjustments.
- **Multipath configuration** (incorporated in the release distribution as of ESXi 7.0U2) no longer requires manual SATP/PSP settings.
- **VMware vVols** provide granular flexibility to VMs that are not available for VMFS VMs.
- **QoS support for vVols and AI QoS** automatically reign in bully volumes.
- **VCF with HPE Storage** is used for principal and supplemental storage through vVols or VMFS.
- **HPE Storage Integration Pack for VMware vCenter** simplifies administration.

Resources

Technical deep dive: HPE GreenLake for Block Storage built on HPE Alletra Storage MP
community.hpe.com/t5/around-the-storage-block/technical-deep-dive-hpe-greenlake-for-block-storage-built-on-hpe/ba-p/7187957

HPE Alletra 9000 Architecture
h20195.www2.hpe.com/v2/getdocument.aspx?docname=a00113422enw

HPE Primera Architecture
hpe.com/psnow/doc/a00115999en_us

HPE GreenLake Block Storage: VMware ESXi Implementation Guide
support.hpe.com/hpesc/public/docDisplay?docId=sd00002428en_us&page=index.html

HPE Alletra 9000 VMware ESXi Implementation Guide
support.hpe.com/hpesc/public/docDisplay?docId=a00115274en_us&docLocale=en_US

HPE Primera VMware ESXi Implementation Guide
support.hpe.com/hpesc/public/docDisplay?docId=sd00001341en_us&docLocale=en_US

VMware and HPE
vmware.com/partners/work-with-partners/global-partners/hpe.html

HPE Active Peer Persistence
hpe.com/psnow/doc/a00115612enw?

Implementing vSphere Metro Storage Cluster (vMSC) using HPE Primera Peer Persistence
kb.vmware.com/s/article/77061

VMware vVols QoS
hpe.com/psnow/doc/a50002145enw

Learn more at
hpe.com/storage

Explore HPE GreenLake

Chat now (sales)

a00115204ENW, Rev. 3

Hewlett Packard Enterprise