HPE SimpliVity hyperconverged infrastructure for VMware vSphere

Technical white paper

Contents

Introduction

Purpose

This technical white paper reviews the features and benefits of HPE SimpliVity hyperconverged infrastructure and describes the hardware and software components and subsystems that make up the solution.

The document explains the functions and inner workings of the solution. It decomposes the HPE SimpliVity solution into physical and logical constructs and describes the various internal and external interfaces and data flows. Particular attention is given to the native data efficiencies of the HPE SimpliVity architecture as well as the VM-centric, unified management approach that simplifies system administration, data protection, and service automation.

The document also reviews supported configurations and deployment models, and provides sample use cases for HPE SimpliVity hyperconverged infrastructure.

Audience

This technical white paper is intended to assist IT planners, architects, system administrators, system integrators, and other technology professionals who are researching, evaluating, or implementing HPE SimpliVity hyperconverged infrastructure solutions. It explains the key performance, scalability, availability, and manageability aspects of the hyperconverged solution and provides a brief overview of the basic system design, deployment, and integration guidelines to assist with project planning and best practices development.

Introduction to HPE SimpliVity hyperconverged infrastructure

Background

Many enterprises are hampered by legacy IT infrastructure that isn't well suited for today's cloud-based services and on-demand applications. These legacy IT architectures are composed of silos of compute, storage, network, and data protection platforms with distinct administrative interfaces. Each platform requires support, maintenance, licensing, power, and cooling—not to mention a set of dedicated team members capable of administrating and maintaining the system. Deploying a new application can be a manually intensive, time-consuming proposition involving multiple different platforms, management interfaces, and operations teams. Turning up new IT services can take days or even weeks. Troubleshooting problems and performing routine data backup, replication, and recovery tasks can be just as inefficient.

The data problem

Enterprise data has historically grown exponentially, resulting in ever-increasing storage capacity needs. At the same time, businesses are placing pressure on IT to back up data more frequently, restore that data more quickly, and support more aggressive SLAs. Unfortunately, performance of individual hard disk drives (HDDs) has not increased relative to the capacity increases, causing a major disconnect between storage capacity and storage performance. The solution to this problem has traditionally required either the utilization of solid-state drives (SSDs) or the creation of oversized pools of HDDs to provide the necessary performance, both of which increase the cost of the storage system.

Enterprise IT organizations have invested in a variety of discrete data storage, data efficiency, and data protection solutions to cope with this ever-increasing data growth. Various point technologies were introduced over time—bolted on in a reactive manner to tackle a particular symptom of the data problem. As a result, many organizations are now held hostage to a multitude of disparate data storage, management, and protection products. Each product is often supplied by a different vendor and features a unique administrative interface and distinct APIs.

A legacy IT stack might include

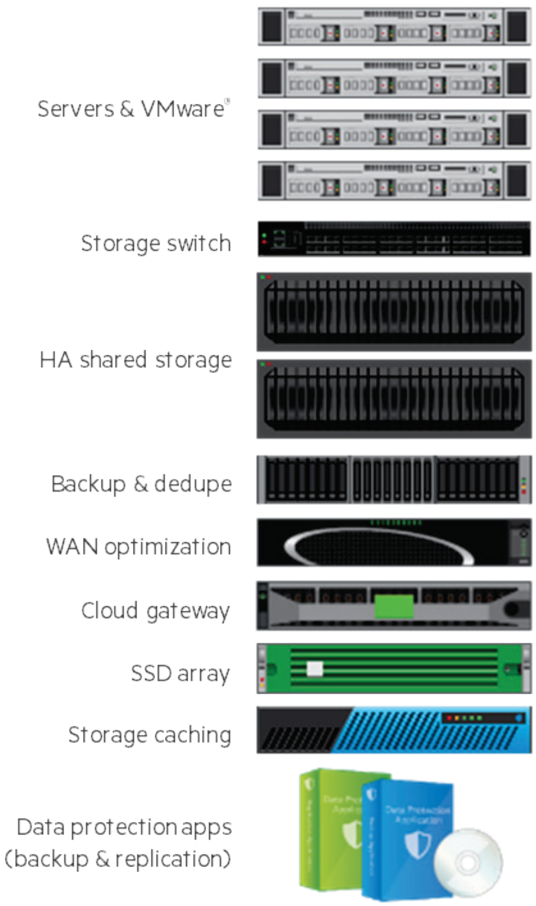

Figure 1. Siloed IT architectures are inherently inefficient, cumbersome, and costly.

Legacy IT stack limitations

Siloed IT architectures can't meet the stringent price-performance, scalability, and agility demands of today's highly virtualized IT environments and data-intensive applications. Legacy IT stack disadvantages and constraints include

- **High CAPEX**—A variety of independent technology platforms to purchase and scale results in resources that are often wasted, due to the chronic difficulty in scaling and integration of the individual platforms. This leads to regular overbuying of capacity and repeat periods of underutilized IT assets.
- **High OPEX**—With so many independent technology platforms, IT operational expenses can be high, with multiple points of management, additional training requirements, hardware and software maintenance fees, power, cooling, and rack space for each silo.
- **Inadequate data protection**—Many IT organizations are challenged to meet backup and recovery goals with existing data protection solutions. This is particularly true in smaller data centers or remote offices due to cost constraints. VM backup and restore jobs can take hours or days when using legacy solutions that store and manage the different stages of the data lifecycle in separate, independent silos.
- **Inefficient resource utilization**—Conventional data storage, optimization, and protection solutions were designed around physical servers and these storage systems aren't well suited for virtualized IT. To ensure peak performance within each silo, these solutions overprovision storage resources and network bandwidth, causing customers to overbuy in every silo.

- **Poor service agility**—Turning up new applications and IT services can take weeks, due to the complexities of coordinating multiple technology platforms and points of administration. This impairs the organization's service velocity and time to market.
- **High barriers to innovation**—IT teams spend the bulk of their time provisioning, maintaining, extending, and refreshing underlying IT infrastructure instead of working on strategic business initiatives.

The solution: HPE SimpliVity hyperconverged infrastructure

HPE SimpliVity hyperconverged infrastructure is designed from the ground up to meet the increased efficiency, management, and data protection demands of today's data-intensive, highly virtualized IT environments. The HPE SimpliVity solution provides a scalable, modular, 2U building block of x86 resources that offers all the functionality of traditional IT infrastructure—in one device. It assimilates storage; compute; hypervisor; and real-time deduplication, compression, and optimization; along with comprehensive data management, data protection, and disaster recovery capabilities.

Designed to work with any hypervisor or industry-standard x86 server platform, the hyperconverged solution provides a single, shared resource pool across the entire IT stack, eliminating point products and inefficient siloed IT architectures. The solution is distinguished from other converged infrastructure solutions by three unique attributes: guaranteed data efficiency; built-in resiliency, backup, and disaster recovery; and global VM-centric management and mobility.

Guaranteed data efficiency

HPE SimpliVity hyperconverged infrastructure performs hardware-assisted inline deduplication, compression, and optimization on all data at inception across all phases of the data lifecycle (primary, backup, WAN, archive, and in the cloud), across all tiers within a system (DRAM, SSD, and HDD), all handled with fine data granularity of just 8 KB. By driving efficiencies at the point of origin and by offloading processor-intensive data handling functions on to purpose-built data acceleration hardware, the solution optimizes the use of system resources and minimizes I/O and network traffic, accelerating data clone, backup, restore, and mobility operations.

Built-in resiliency, backup, and disaster recovery

Another core intention of the HPE SimpliVity platform is data protection. Maintaining availability of data is the most critical function of any platform that stores data. Within most HPE SimpliVity hyperconverged nodes (except for the Remote Site nodes and Small All-Flash nodes), data is protected from the loss of multiple disk drives. All data is also synchronously persisted on two local HPE SimpliVity hyperconverged nodes to protect against the loss of any single node. These two levels of protection work together so that the loss of multiple disks in multiple nodes, even when one node is offline due to failure or during maintenance, will not cause data to become unavailable.

The HPE SimpliVity architecture includes built-in VM-centric data protection functionality, eliminating the need for purpose-built backup and recovery products. The ultra-efficient approach for data management accelerates data replication, backup, and restore functions, improving recovery-point objectives (RPOs) and recovery-time objectives (RTOs). VMs can be backed up or recovered in seconds or minutes instead of hours or days as with legacy data protection solutions. Backups can be created manually or automatically based on policy, and can be stored either locally or remotely by simply choosing the appropriate HPE SimpliVity cluster.

Global VM-centric management and mobility

Simplifying the complexities and administrative overhead of managing multiple silos of IT infrastructure distributed across multiple data centers was one of the intentions of this hyperconverged solution from the very beginning. Rather than managing individual silos at the device level, a globally distributed HPE SimpliVity implementation is managed as a unified system across multiple sites. The entire system, called an HPE SimpliVity Federation, is centrally managed through a single administrative interface and common APIs.

Normally, managing a global IT infrastructure requires tracking and updates of IP addresses to ensure cross-site connectivity is maintained when introducing or upgrading any devices. The global features of HPE SimpliVity eliminate the need to maintain cross-site IP configurations in any site when adding or removing nodes.

Another way HPE SimpliVity has simplified IT administration is to place focus on managing data at the most important level in today's data center: the VM. A VM-centric management approach simplifies the monitoring and administration of virtual applications, while shielding administrators from the complexities of the underlying IT infrastructure. In an HPE SimpliVity-based IT environment, administrators no longer need to worry about disk groups, LUNs, or replication groups because the administration actions are abstracted away from the underlying infrastructure, and instead, are based on actions and policies applied to the VM.

All of this is implemented within the confines of existing virtual infrastructure management frameworks (for example, VMware® vCenter™, VMware vRealize Automation, Cisco UCS Director), further reducing the number of management interfaces and IT staff training needs.

Technology overview

Terminology

The following terms are used to describe certain HPE SimpliVity architectural elements and constructs:

- **HPE SimpliVity hyperconverged node**—An x86 server that is the basic hardware building block of the HPE SimpliVity hyperconverged infrastructure solution
- **HPE OmniStack Virtual Controller (OVC)—**The software stack, implemented as a single VM per node, which controls all aspects of HPE SimpliVity hyperconverged infrastructure
- **HPE OmniStack Accelerator Card (OAC)—**PCIe-based device that offloads and provides acceleration of writes and data management functions within the HPE SimpliVity hyperconverged infrastructure solution
- **HPE SimpliVity Data Virtualization Platform (DVP)—**A globally aware File System and Object Store with data optimization techniques that enables a coordinated collection of scalable compute and storage resource pools across multiple sites, and provides highly efficient data storage, management, and mobility
- **HPE SimpliVity Cluster**—A collection of one or more HPE SimpliVity hyperconverged nodes typically located at the same physical site connected over a standard Ethernet network collectively providing a single storage pool to the hypervisor on each node. An HPE SimpliVity Cluster can also be extended across two physical sites, commonly known as a stretched cluster, over low latency metro networks for disaster recovery and business continuity
- **HPE SimpliVity Federation**—A collection of one or more HPE SimpliVity Clusters and the main construct within which data is managed
- **Intelligent Workload Optimizer**—A feature of an HPE SimpliVity Cluster that provides a multidimensional approach to balancing workloads across HPE SimpliVity hyperconverged nodes based on CPU, memory, and storage metrics to keep applications running at their peak. Avoids costly data migrations through greater intelligence
- **OmniView**—A cloud-based web application that uses predictive analytics to deliver actionable intelligence, enabling customers to quickly understand system performance, optimize resource utilization, and efficiently plan and manage future growth within their own HPE SimpliVity infrastructure
- **OmniWatch**—A proactive support service for all HPE SimpliVity support plans that complements our support portal for comprehensive care. OmniWatch continuously monitors your HPE SimpliVity infrastructure, checking over 100 indicators to continuously evaluate system health and automates support case creation

Scale-out architecture

The HPE SimpliVity solution features a scale-out architecture that minimizes upfront investments and provides a high degree of flexibility and extensibility. HPE SimpliVity hyperconverged nodes are installed in an incremental fashion to accommodate growth, enable new applications, or extend system availability. Two or more HPE SimpliVity hyperconverged nodes can be combined across one or more data centers into an HPE SimpliVity Federation to create a massively scalable collection of shared resources that are administered as a single global system. This global Federation is managed through a single administrative interface or with a unified programmatic interface for easy integration with external management applications and service orchestration platforms.

HPE SimpliVity hyperconverged node

The basic building block of the HPE SimpliVity solution is referred to as an HPE SimpliVity hyperconverged node. It includes three fundamental components.

A compact hardware platform

This hyperconverged infrastructure was designed to be server hardware agnostic and capable of running on any industry-standard x86 platform supporting a hypervisor and containing compute, memory, PCIe slots, and 1GbE and 10GbE network interfaces.

HPE SimpliVity nodes are available as either hybrid nodes or all-flash nodes. Hybrid nodes include a combination of performance-optimized solid-state drives (SSDs) and capacity-optimized hard disk drives (HDDs). All-flash nodes include only performance-optimized SSDs. In both cases, all disks are protected in hardware RAID configurations to provide greater resiliency to disk failure.

In the hybrid nodes, the HDDs are utilized for permanent virtual machine data and hash table index storage, and are configured to support at least two drive losses per node (all models except Remote Site models). The SSDs are utilized for permanent metadata storage, as well as a caching tier for virtual machine data, and are configured to support a single drive loss. This design philosophy is based on the historical failure probabilities of HDD and SSD components. Since HDDs with moving parts have a significantly higher annualized failure rate and much longer rebuild time, double-drive failures on HDD are much more likely than on SSD.

In the all-flash nodes, the SSDs are utilized for permanent virtual machine data, hash table index storage, and permanent metadata storage. The drives are configured to support at least one drive loss per node (depending on model).

When hardware-based data-at-rest encryption is required, self-encrypting disks (SEDs) are supported on specific server platforms as identified in the most recent support matrix.

Depending on model, one or two additional disks are utilized to host the hypervisor installation and the HPE OmniStack Virtual Controller storage. Please see the latest data sheets for specifics on drive configurations.

Running on the bare metal of every node is a virtual machine hypervisor. HPE SimpliVity was designed to be hypervisor agnostic to provide customer choice of which hypervisor best fits their environment. This is made possible by the implementation of separate Presentation and Data Management layers, which allows data to be managed agnostic to the hypervisor being used. This document specifically addresses implementation of HPE SimpliVity with VMware vSphere®.

The HPE SimpliVity 380 hyperconverged node is a turnkey appliance based on the industry's best-selling server, the HPE ProLiant DL380.[1] Prior to the HPE acquisition of SimpliVity, legacy solutions were integrated with a number of other x86 servers under the name OmniStack Integrated Solutions. All of the deployment options run the same HPE SimpliVity components.

[1] Based on the CQ1 2016 IDC Server tracker, 2016

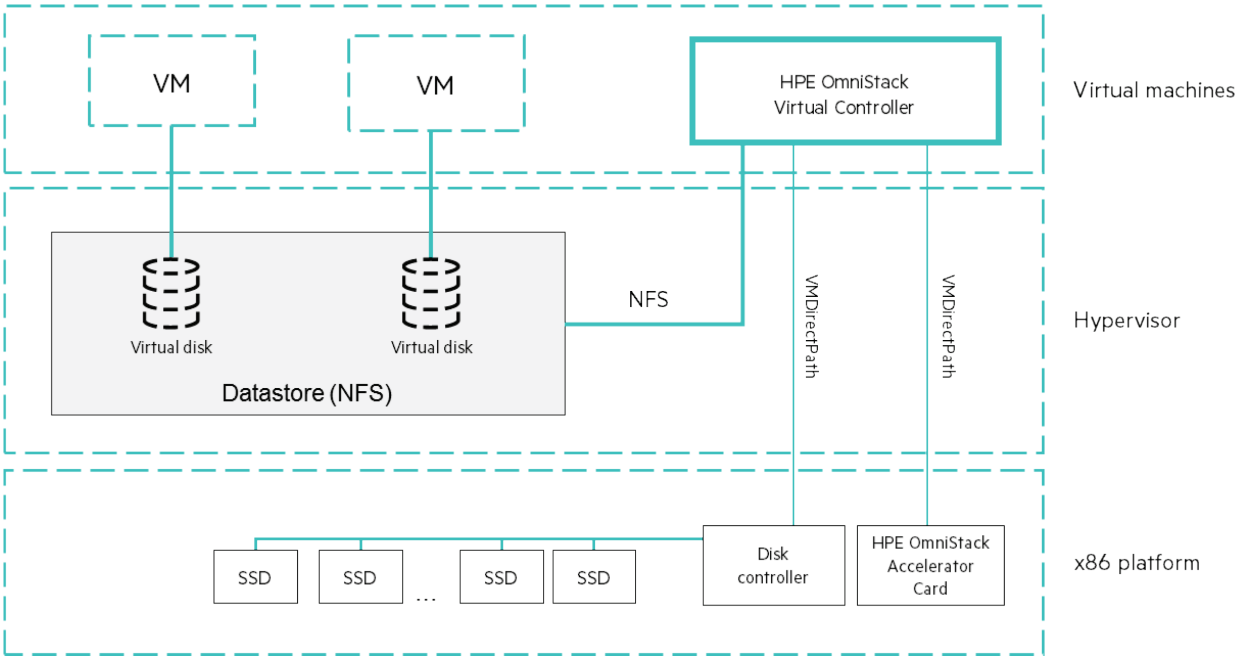

Figure 2. Components of an HPE SimpliVity hyperconverged node.

HPE OmniStack Virtual Controller

A VM-based controller running on the hypervisor is the foundational software of the HPE SimpliVity hyperconverged infrastructure. Every HPE SimpliVity hyperconverged node will have a single HPE OmniStack Virtual Controller running on it.

HPE OmniStack Accelerator Card

The HPE OmniStack Accelerator Card is a purpose-built PCIe-based accelerator card designed by Hewlett Packard Enterprise (HPE). The HPE OmniStack Accelerator Card consists of

1. Field-programmable gate array (FPGA)
2. NVRAM, which consists of
   a. DRAM
   b. Flash storage used to persist the data in DRAM in a nonvolatile state after power loss
   c. Super capacitors to provide power to the DRAM and flash upon a power loss

HPE OmniStack Accelerator Card offloads CPU-intensive functions from the x86 processors, performing inline deduplication, inline compression, and inline write optimization. By offloading these functions, HPE SimpliVity nodes can provide these core features with improved performance and minimal impact to the resources available to business application VMs.

HPE SimpliVity Cluster

HPE SimpliVity Cluster is the primary logical construct that defines a single grouping of HPE SimpliVity hyperconverged nodes. The HPE OmniStack Virtual Controllers on these nodes aggregate the storage local to each node and present one or more datastores to the hypervisor instance on each node.

An HPE SimpliVity Data Center is automatically defined in a VMware environment based on the membership of a VMware vSphere datacenter. No further configuration is necessary.

VMware® vCenter Server™

All enterprise-class VMware infrastructures are reliant on vCenter Server for efficient management. HPE SimpliVity builds off this concept by providing daily HPE SimpliVity management functionality directly within the vSphere Client. These tasks are available due to the integration of the HPE SimpliVity plug-in to either the vSphere C# Client or VMware vSphere® Web Client.

With the 3.5.1 release of the HPE OmniStack software, HPE SimpliVity will support vCenter Server—both Windows®-based or VMware® vCenter™ Server Appliance™ (VCSA)—deployed either inside or outside of the HPE SimpliVity infrastructure it will manage. To support full data protection of a vCenter Server deployed within the HPE SimpliVity infrastructure, emergency restore capabilities have been added to the CLI to enable backup restoration of the vCenter Server VM.

HPE SimpliVity Arbiter

A split-brain scenario occurs when two independent systems configured in a cluster lose network connectivity and assume they have exclusive access to resources. The HPE SimpliVity Arbiter ensures that a two-node or larger cluster survives a hyperconverged node failure without service disruption or loss of access to data.

The HPE SimpliVity Arbiter is software residing on a Windows system that can communicate with the hyperconverged nodes, and facilitates communication between nodes and resolves state conflicts to ensure service continuity. The Arbiter is what enables a smaller data center of only two HPE SimpliVity nodes versus the need for a minimum of three nodes—which reduces initial CAPEX costs for small high availability configurations.

Under normal operations, HPE SimpliVity hyperconverged nodes and the HPE SimpliVity Arbiter are in constant communication through heartbeats that are sent across the network. However, it is important to note that the HPE SimpliVity Arbiter is not required for the environment to run. If the Arbiter is taken offline, the two-node environment will continue to operate without interruption, though VM availability will be at risk if one of the nodes were to fail.

For the Arbiter to properly provide this independent witness, it must exist outside of the HPE SimpliVity Cluster in which it is a participant, and therefore cannot exist on an HPE SimpliVity node or datastore it is witnessing. The HPE SimpliVity Arbiter is commonly installed on the vCenter Server, but is fully supported on a separate unrelated server and must remain separated if choosing to run vCenter on the HPE SimpliVity hyperconverged nodes.

HPE SimpliVity datastores

When configuring an HPE SimpliVity Data Center, one or more datastores must be created to present the storage to the VMware® ESXi™ hypervisor. Each datastore is presented as an NFS datastore with a common namespace across all the HPE OmniStack Virtual Controllers in an HPE SimpliVity Cluster. These datastores are logical constructs within the HPE SimpliVity Data Virtualization Platform Data Presentation layer, and are simply utilized as a communication channel between the hypervisor and the HPE SimpliVity Data Virtualization Platform. Any number of datastores can be created, with no ties to the actual underlying storage capacity. These datastores can be resized, larger or smaller, online with no disruption to storage availability or performance.

Each datastore has a single policy applied. This policy acts as the default policy for any virtual machines created on the datastore. Creating multiple datastores provides a mechanism for easy application of policies to virtual machines.

Deduplication is not bound by datastore since this happens across all data on an HPE SimpliVity hyperconverged node. Similarly, replication is not limited by datastores since it is defined at the virtual machine level (via policy).

Compute nodes

Compute nodes are nonhyperconverged servers that plug into an HPE SimpliVity solution and provide additional compute capacity while consuming the storage provided by the hyperconverged nodes. Effective use of compute nodes in a properly designed environment can provide greater performance to tier-1 applications, virtual desktops, and other workloads, while driving down costs in your environment. Nearly all major server platforms are supported as compute nodes, including HPE, Dell, Cisco, Lenovo, and Huawei x86 servers. When upgrading your data center to an HPE SimpliVity infrastructure, you can take advantage of your existing server investments to repurpose them as compute nodes within your HPE SimpliVity deployment.

Stretched clusters

An HPE SimpliVity Cluster typically exists within a single physical site but can be extended across two physical sites when the two sites are connected by a high bandwidth (10 Gbps), low latency interconnect, generally < 1 ms. However, higher latencies can be deployed depending on application performance requirements.

A stretched HPE SimpliVity Data Center ensures fully committed synchronous writes between two physical sites by only acknowledging writes back to the VM after nodes at both sites have safely persisted the write. This allows the utilization of VMware HA functionality to automate the recovery of VMs after the failure of an entire site, providing RPOs of zero and RTOs of seconds. To ensure this functionality and avoid split-brain scenarios, HPE recommends the HPE SimpliVity Arbiter be deployed at a third site to avoid losing it along with one-half of the HPE SimpliVity Cluster.

This functionality is implemented via Availability Zones. By placing the nodes in each site into one of two Availability Zones, the HPE SimpliVity Data Virtualization Platform will ensure that nodes used for each VM's replica set will be pulled from both Availability Zones. (See HPE SimpliVity Cluster availability section for more details on replica sets.) This takes advantage of the HPE SimpliVity Data Virtualization Platform's native functionality to synchronously persist all data to two nodes.

An HPE SimpliVity stretched cluster can be implemented either as part of a new (greenfield) deployment or enabled on an existing (brownfield) deployment. In both cases, both sites must have an identical number of nodes.

HPE SimpliVity Federation

An HPE SimpliVity Federation is a networked collection of one or more HPE SimpliVity Clusters. The HPE SimpliVity Federation is a logical construct that allows the management of the total IT infrastructure across sites as a single entity, rather than at the individual component or cluster level. With this model, data can move between sites without the need to remove the data from its deduplicated and compressed format. This allows an HPE SimpliVity infrastructure to efficiently move only unique data between HPE SimpliVity Clusters when performing Remote Backup and VM Move operations.

In a VMware vSphere environment, an HPE SimpliVity Federation is automatically established when creating multiple vSphere clusters under a single vCenter Server, or by connecting multiple vCenter Servers, each with one or more clusters, with vSphere Linked Mode. vSphere Linked Mode is used to join multiple vCenter Servers for unified monitoring and management purposes.

This capability enables centralized HPE SimpliVity administrators to perform data protection and management functions transparently across sites, regardless of vCenter Server and data center boundaries, including VM-level backups, restores, and full VM movement.

While not mandatory, HPE recommends implementing one vCenter Server for each data center that contains one or more HPE SimpliVity Clusters to ensure each site can maintain full vSphere management capabilities even if the other site is unavailable.

Supported Federation topologies

HPE SimpliVity Federation can be configured in either a full-mesh or a hub-and-spoke topology to support diverse deployment scenarios. The full-mesh topology is primarily intended for data center, campus, and metro network deployments. It enables any-to-any mobility, where VMs can migrate, back up, and restore from any HPE SimpliVity Cluster to any other HPE SimpliVity Cluster, without restrictions. The hub-and-spoke topology is intended for geographically distributed deployments such as remote office/branch office implementations. VMs can migrate, back up, and restore only between a spoke site (that is, a branch office) and a hub site (that is, a central data center or corporate headquarters). The hub-and-spoke topology minimizes communications traffic between spokes.

When deploying additional HPE SimpliVity Clusters, the software will automatically detect the network topology and join either a hub-and-spoke or mesh topology with no user input, eliminating the overhead and potential misconfigurations that come from individually configuring the intercommunication between each data center and remote office.

HPE SimpliVity Data Virtualization Platform

At the heart of the HPE SimpliVity solution lies the HPE SimpliVity Data Virtualization Platform—a fabric that extends across multiple HPE SimpliVity hyperconverged nodes in multiple globally distributed data centers that abstracts the VM data away from the underlying hardware. In a single data center, the HPE SimpliVity Data Virtualization Platform abstracts the underlying hardware across multiple HPE SimpliVity hyperconverged nodes and presents this unified storage to the hypervisor on each node as a single pool of storage. HPE SimpliVity Data Virtualization Platform performs inline data deduplication, compression, and optimization functions across all data ingested, and provides the administrator the ability to execute hardware-accelerated backup, restore, clone, and move operations of any VM within the infrastructure. Across data centers, this platform also allows for a single point of management, including unified policy definitions for data protection and efficient movement of data between physical locations.

HPE SimpliVity Data Virtualization Platform is a logical construct realized by HPE OmniStack Virtual Controller software and HPE OmniStack Accelerator Cards working in concert. It is comprised of multiple logical layers.

Presentation layer

The Presentation layer is responsible for the creation and presentation of datastores to the hypervisor and maps datastore objects to the underlying Data Management layer. These datastores are presented to vSphere via NFS using a unified namespace across all HPE OmniStack Virtual Controllers. This is the layer that requires the majority of customization when supporting other hypervisors (for example, SMB will be utilized for Hyper-V hosts). Each HPE OmniStack Virtual Controller maintains its own IP and piece of the datastores, but coordinates with the others to present a single datastore to the hypervisors in the Cluster.

Data Management layer

The Data Management layer is responsible for the tracking and storage of all data and metadata. Each instantiation of a virtual machine is tracked as a data container, as represented in Figure 3 by the squares in the Data Management layer. These data containers represent both active virtual machines and backup versions of the virtual machines. The data containers that represent active virtual machines are presented through the Presentation layer.

Each HPE OmniStack Virtual Controller maintains its own record of metadata independent of all other nodes. All data and metadata is deduplicated within the node regardless of datastore location or content type (that is, ISO file, VMDK, VHD, and so on).

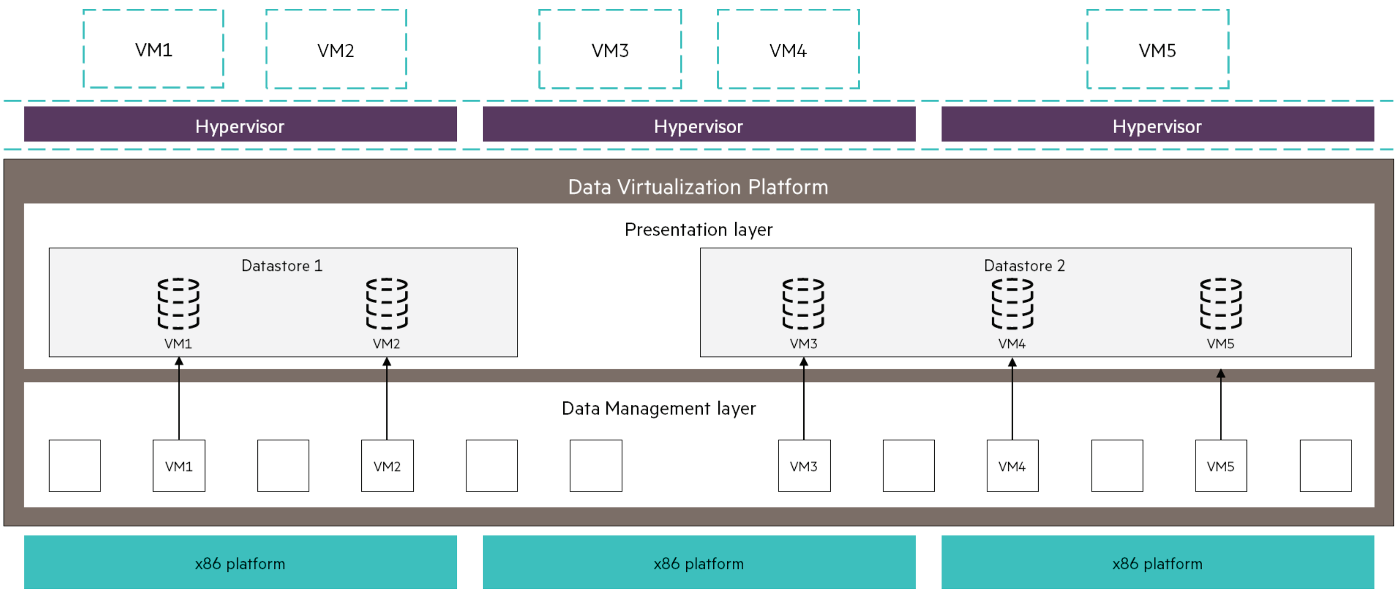

Figure 3. HPE SimpliVity Data Virtualization Platform in a single HPE SimpliVity Cluster spanning three HPE SimpliVity hyperconverged nodes.

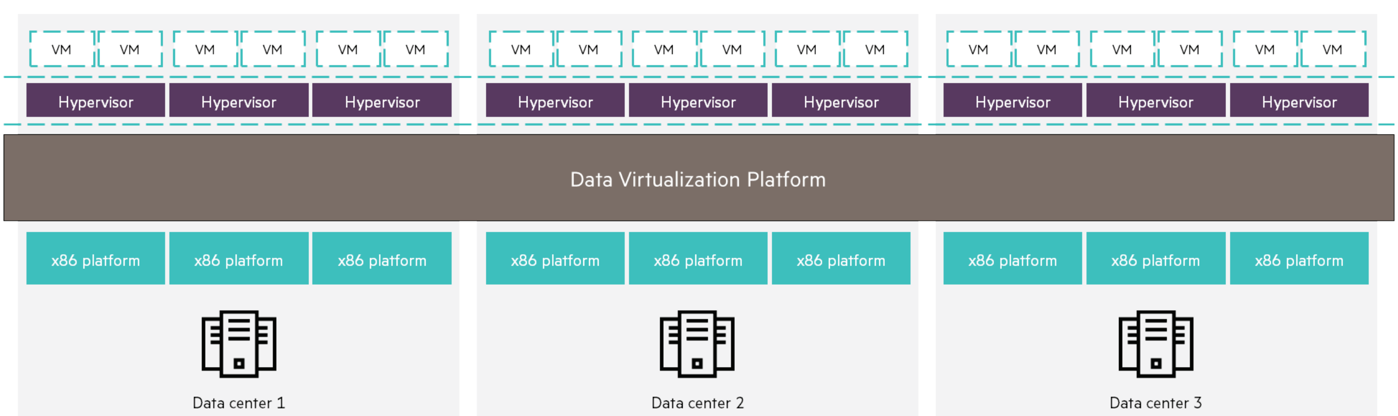

Figure 4. HPE SimpliVity Data Virtualization Platform in an HPE SimpliVity Federation spanning three data centers.

The Data Management layer is further subdivided into two other layers: the File System and Object Store. They work together to track and manage all metadata and blocks of physical data.

File System

The primary job of the File System is tracking the relationships between different pieces of metadata and between metadata and data. It does not actually store any data or metadata, but does know how to read the metadata and interpret the relationships that metadata forms. The result is a logical hierarchy of metadata laid out like an upside-down tree, where the root at the top represents the entire VM and the leaves at the bottom are the actual blocks of data.

Object Store

All the metadata is stored and managed as individual objects within the Object Store. The metadata defines its own subordinates, which could be either metadata or physical blocks of data, which allows the File System to be able to reconstruct the virtual machine files. By tracking the relationship of all the metadata objects, it is easy and very efficient to create logical copies of virtual machines to create clones and backups. When creating a clone of a virtual machine, only the highest-level metadata object needs to be copied. This can be a significantly quicker operation for larger virtual machines than having to copy the underlying metadata that directly describe the actual blocks.

Data mobility

Creation of the metadata objects and the indexes that point to them follows the same formula on all nodes, regardless of location within the Federation, which ensures that comparisons between nodes and sites are consistent. When moving data between sites (via Remote Backup or VM Move operations), a comparison of existing objects occurs between the sending and receiving nodes. This is executed by investigating the relationship of the data defined within the metadata. Every data container is a hierarchy of metadata that defines the structure of the underlying data for that VM or backup. Through this hierarchy, the entire sections of data can be found to be duplicate between sites, and therefore avoiding data transfer of that data to the receiving site, and also not even needing to be compared. Once a section of the metadata hierarchy is found to be duplicate between sites, the HPE SimpliVity Data Virtualization Platform knows that everything within that section will also be duplicate. This leads to replication of both data and metadata that can be as efficient as possible since only truly unique objects are transferred between sites, and not all metadata needs to be investigated to determine what is duplicate between sites.

VM-centric design

Most of today's data center equipment was originally designed around nonvirtualized physical servers and shared storage arrays. When virtualization was introduced into the data center, the virtualized data center utilized this legacy equipment instead of starting with a clean slate that would have allowed virtualization to realize its full potential. In a typical legacy infrastructure, data associated with multiple VMs is pooled together in a single datastore on a LUN on a shared storage array. In this scenario, storage-based functions (that is, backup, restore, replication, performance monitoring) are either performed at the LUN level, with little or no visibility to the VM level, or outside of the storage system, with little or no visibility at the storage layer. Both approaches are inherently inefficient from a monitoring, management, and resource utilization perspective as well as inherently complex and impractical from an operations perspective. In an ideal implementation, system administrators would institute policies and manage performance, data protection, and SLAs on applications and workloads, and not physical devices.

HPE SimpliVity Data Virtualization Platform was designed from the ground up with virtualization in mind. All administrative tasks, including managing data protection policies, analyzing performance, and troubleshooting problems, are performed at the VM level. All monitoring of VM performance and alerting occur at the VM level. With everything centered on the VM, it only makes sense that this all happens within the same tool used to manage all other virtualization infrastructure tasks (for example, VMware vSphere Client).

Inline deduplication, compression, and optimization

HPE SimpliVity Data Virtualization Platform Data Management layer performs inline data deduplication, compression, and optimization on all data, at inception, across all tiers (DRAM, SSD, and HDD) and phases of the data lifecycle (primary, backup, WAN, archive, and in the cloud). All data is deduplicated, compressed, and optimized in 8 KB granularity by the HPE OmniStack Accelerator Card inline before it is written to disk. With HPE SimpliVity, deduplication and compression does not require any intervention by the administrator since it is always on and accelerated by the HPE OmniStack Accelerator Card. Utilizing dedicated hardware ensures that deduplication can be performed 100% of the time no matter the workload and without performance impact to the application.

By deduplicating writes before they go to the disks, the I/O to commit the blocks of data are also eliminated. All of the processing is offloaded to the HPE OmniStack Accelerator Card, with writes first committed to NVRAM to ensure inline deduplication actually improves performance instead of impeding it (see the Lifecycle of a write I/O section for more detail).

This approach improves both performance and economics by conserving storage resources and minimizing I/O. By offloading processor-intensive data efficiency functions to the HPE OmniStack Accelerator Card, x86 CPU resources are reserved for business applications.

Resiliency and high availability

HPE SimpliVity solution is designed to be highly resilient with no single point of failure. The solution supports node-level resiliency via RAID and other hardware resiliency technologies, cluster-level redundancy via data and metadata mirroring and other techniques, and site-level protection via HPE SimpliVity-enabled DR functionality.

Node availability

HPE SimpliVity utilizes standard x86 servers for the compute platform, gaining all the availability features inherent in the selected platform, including redundant power supplies, redundant fans, multiple NIC ports, and error-detection and recovery within system RAM.

Hardware RAID is also used for protection of the disk subsystem. RAID 1, RAID 5, RAID 6, or RAID 60—depending on the number of disks—is used to protect data on the disks. In hybrid arrays, HDDs and SSDs are grouped into two separate RAID sets. In all-flash configurations, SSDs are grouped into a single RAID set. This ensures disk-level resiliency so that the loss of disks will not impact the availability of data on an individual node.

Table 1. RAID levels utilized in HPE SimpliVity hyperconverged nodes.

Drives	RAID level	Survivable failures
2	1	1
4	5	1
5	5	1
8	6	2
9	6	2
10	6	2
12	6	2
14	60	2 per RAID 6 set
20	60	2 per RAID 6 set

If a node loses an HPE OmniStack Accelerator Card, disk controller, or HPE OmniStack Virtual Controller, the virtual machines on that node continue to run. All storage functions will continue on another node by failing over the storage IP address of the affected HPE OmniStack Virtual Controller. This failover occurs in less than a second, avoiding any interruption of storage access for the VM. This approach creates better uptime of VM than the failure of a motherboard or CPU, which would result in VM downtime and a vSphere HA event, if configured. More on that in the next section.

HPE SimpliVity Cluster availability

In a high availability implementation, the Data Management layer statefully maintains the complete set of data associated with a VM on two distinct hyperconverged nodes within an HPE SimpliVity Cluster. These two instances are referred to as a **replica set**. Should one hyperconverged node fail, the VM's data is still available on another node so vSphere HA can restart all the affected VMs.

By simultaneously utilizing disk-level resiliency and node-level resiliency, an HPE SimpliVity Cluster can withstand the loss of disks in every node in a cluster while an entire node is offline for maintenance without losing a single VM.

Lifecycle of a write I/O

The replica set associated with a particular VM is managed on two separate HPE SimpliVity hyperconverged nodes within a cluster. The HPE OmniStack Virtual Controllers establish this peering relationship on a per VM basis. The primary HPE OmniStack Virtual Controller on the same node as the VM will share every write operation with the secondary HPE OmniStack Virtual Controller for that VM. Each HPE OmniStack Virtual Controller will then calculate and execute deduplication, compression, and optimization independently. With this approach, the persistence of every block of data is guaranteed across two HPE SimpliVity hyperconverged nodes without the overhead of all nodes tracking the details of how or where each block is stored within the cluster. In other words, disk-level resiliency is abstracted from node-level resiliency. All of this occurs automatically, including the peering process, requiring no administrative configuration. See the section on Intelligent Workload Optimizer for more details on the initial VM creation process.

This approach remains the same when implementing a stretched cluster, and is key to ensuring synchronous writes between physical sites. See the section on Stretched clusters for more detail.

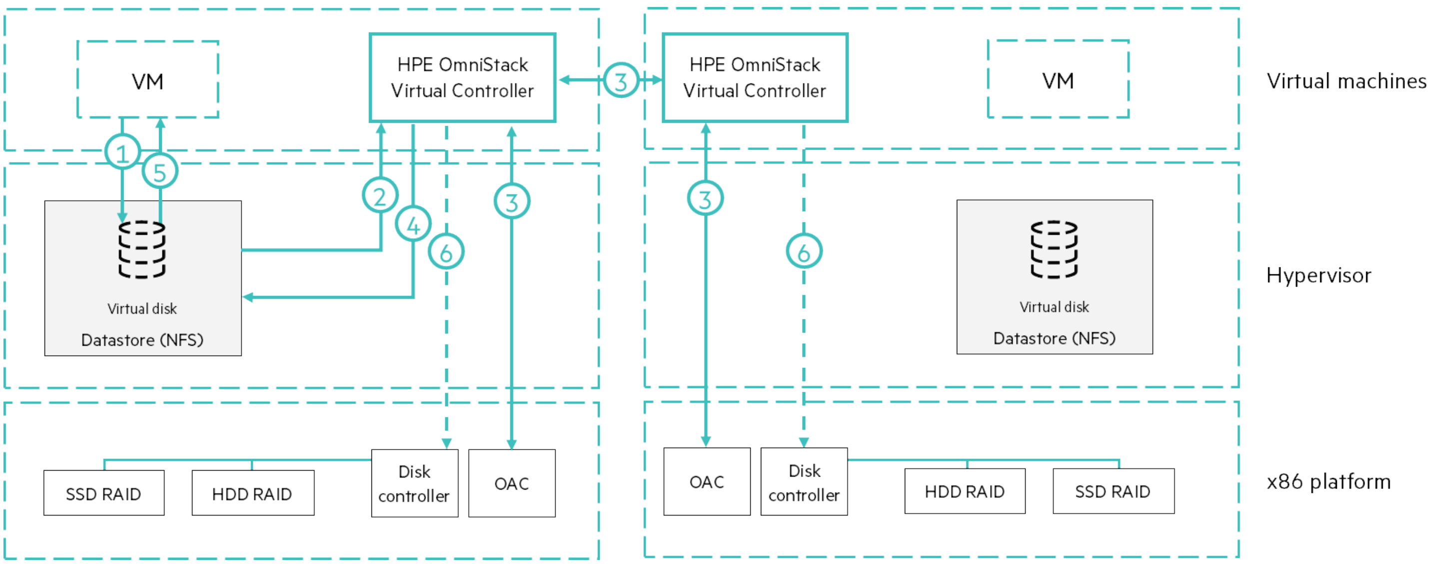

Figure 5. High availability write path.

The Data Management layer performs the following steps to synchronize VM data across nodes:

1. The virtual machine issues a write I/O to the hypervisor.
2. The hypervisor kernel directs the I/O to the datastore via NFS, which is presented by the local HPE OmniStack Virtual Controller.
3. The primary HPE OmniStack Virtual Controller replicates the I/O to the OVC on the VM's secondary node. In parallel, each HPE OVC places the data in NVRAM on its HPE OmniStack Accelerator Card. This NVRAM consists of DRAM that is backed with super capacitors and a flash module so the data can be persisted in the event of power failure to prevent data loss. Once staged in NVRAM, each HPE OmniStack Accelerator Card independently processes the data for deduplication and compression, then acknowledges back to the primary HPE OmniStack Virtual Controller.
4. Immediately after the primary HPE OmniStack Virtual Controller receives acknowledgements from both HPE OmniStack Accelerator Cards, it acknowledges back to the hypervisor that the write is complete.
5. Immediately after Step 4, the hypervisor acknowledges back to the VM that the write is complete. The I/O operation is now complete from the VM's perspective.
6. Independent of the acknowledgment steps, each HPE OmniStack Virtual Controller serializes the data into a full RAID stripe of unique data. When complete, the full stripe is sent to the disk controller to permanently store to disk.

Ensuring high availability

Depending on the type of loss the HPE SimpliVity hyperconverged node suffers, the response will be different. Components core to the operations of the x86 hardware, like CPUs or the motherboard, will cause a complete outage of the node and result in a vSphere HA event. Storage access for the affected virtual machines will failover to the other HPE OmniStack Virtual Controller in the replica set. When VMware vSphere® Distributed Resources Scheduler™ (DRS) is enabled on the cluster, Intelligent Workload Optimizer (see Intelligent Workload Optimizer for more details) will ensure that vSphere HA will restart the VM on the other node in the replica set. This ensures data locality and an optimal and consistent storage performance, even with a failed node. This is due to the HPE SimpliVity architecture, which stores a complete copy of the VM on two different nodes, instead of distributing either the secondary copy or both copies across all the nodes, thus impacting the performance of the VM after the restart. This approach also eliminates unnecessary disk and network I/O that would be created to either cache or recreate a complete copy of the VM on the recovery node.

Each HPE OmniStack Virtual Controller has an IP address that the hypervisor uses to access the NFS service. If an HPE SimpliVity hyperconverged node has an HPE OmniStack Virtual Controller, disk controller, or HPE OmniStack Accelerator Card fail, the IP address and NFS mount of the HPE OmniStack Virtual Controller automatically fail over to a surviving node's HPE OmniStack Virtual Controller. Since ESXi uses this IP address to access the datastore, all VMs and applications continue to run without interruption because no vSphere HA event occurs or is necessary. This prevents these components from being a single point of failure for virtual machine or data availability. The net effect is that the loss of the HPE SimpliVity components within a node has less impact to the virtual machines than the loss of the motherboard, and should not be considered a single point of failure within the cluster.

To protect from the loss of an entire physical data center, customers can choose to utilize a stretched cluster or remote backups. Utilizing stretched clusters, customers can achieve zero data loss recovery if a site were to fail because all writes are written to both sites synchronously, as described previously in the Lifecycle of a write I/O section. In this case, the loss of one site would be detected and VMs restarted by vSphere HA, providing for a minimal amount of downtime.

Federation availability

To protect data across an HPE SimpliVity Federation, remote backups can provide RPOs as short as 10 minutes by utilizing policy-based backups. The recovery of this data requires the restoration of the most recent backup utilizing the techniques in the next section. These restore operations can occur in seconds, significantly speeding up RTOs. This recovery operation can be automated through a robust REST API by using scripts or an orchestration tool to further enhance the RTOs. HPE offers a DR automation product called RapidDR that utilizes the HPE SimpliVity REST APIs to automate the failover tasks of an HPE SimpliVity hyperconverged infrastructure, including the recovery of the VMs, changes to VM networking, and other guest OS-based tasks.

Highly efficient data backup, replication, and recovery

With conventional backup products, full backup solutions squander storage resources and network bandwidth by creating complete backup images. When a full backup executes, every block for that VM must be read off the disk, transferred across the network, and written to some remote system. Incremental and differential backup solutions are more efficient than full backup solutions because they only copy and transfer data that has changed since a previous backup. But they still introduce appreciable storage I/O and networking overhead since that changed data may still be duplicate data. For example, think about installing a 1 GB patch on 20 Windows Server® VMs. Traditional backup products would back up that 1 GB 20 times, including all the associated reads, writes, and network bandwidth because that 1 GB has changed in each VM. This heavy utilization of storage resources during backup periods can negatively impact the performance of other VMs on the same storage and is the reason why backup windows are traditionally only allowed once a day, late at night.

Because all data is deduplicated, compressed, and optimized at inception within HPE SimpliVity hyperconverged infrastructure, data management operations like creating backups, clones, restores, or moves of VMs can be completed locally or globally much more efficiently than with traditional data management solutions. The HPE SimpliVity Data Virtualization Platform maintains full logical backups without impairing performance or squandering storage capacity. Each backup is a complete, stand-alone image of a specific virtual machine, taken at a specific point of time. These backups do not have dependencies on previous backups or connections to a root image or disk, like typical storage snapshots. Individual blocks are not owned by any specific entity, so recovering a VM that's been deleted, infected with ransomware, or from a different site is not a problem.

By tracking individual blocks, the HPE SimpliVity Data Virtualization Platform ensures that only unique blocks are written to disk. This makes it possible to maintain full logical copies of VM backups in a highly efficient and cost-effective manner. It conserves storage resources, minimizes storage I/O and network traffic, and accelerates data backup and recovery operations. This is inherently more efficient than approaches that rely on differential or incremental backup approaches that have dependency mappings from one backup to the next.

When creating remote backups or moving virtual machines between data centers, the HPE SimpliVity hyperconverged nodes in the source data center send a hierarchy of metadata that represents the data needed to recreate the VM. This metadata is then compared to the metadata the receiving hyperconverged nodes already have stored, and request only the necessary unique blocks be transmitted. By creating this hierarchy of metadata, entire sections of common data, like common OS or application files, can be quickly identified and unnecessary I/O, WAN bandwidth, and disk consumption completely avoided. If common OS and application installs are written in the Remote Site prior to the first remote backup, then those files will never need to traverse the WAN links. All of this happens by simply connecting hyperconverged nodes together into a single HPE SimpliVity Federation, and it all happens underneath the hypervisor with no additional technologies to set up, configure, or maintain.

With three simple clicks, administrators can manually restore a full VM or specific files and/or folders from a previous working state to recover from application problems or administrative mishaps. A full VM restore allows the administrator to choose to restore to a new VM anywhere in the HPE SimpliVity Federation or to replace the existing VM in place. A file-level restore allows the administrator to choose the relevant files and folders, and the system presents those files and folders on a virtual DVD image to any VM in the HPE SimpliVity Federation. This gives the administrator total control over the restoration process. Since none of this restored data is unique, there is no actual data to read or write. This means that no disk I/O is utilized and the restores can be completed very fast.

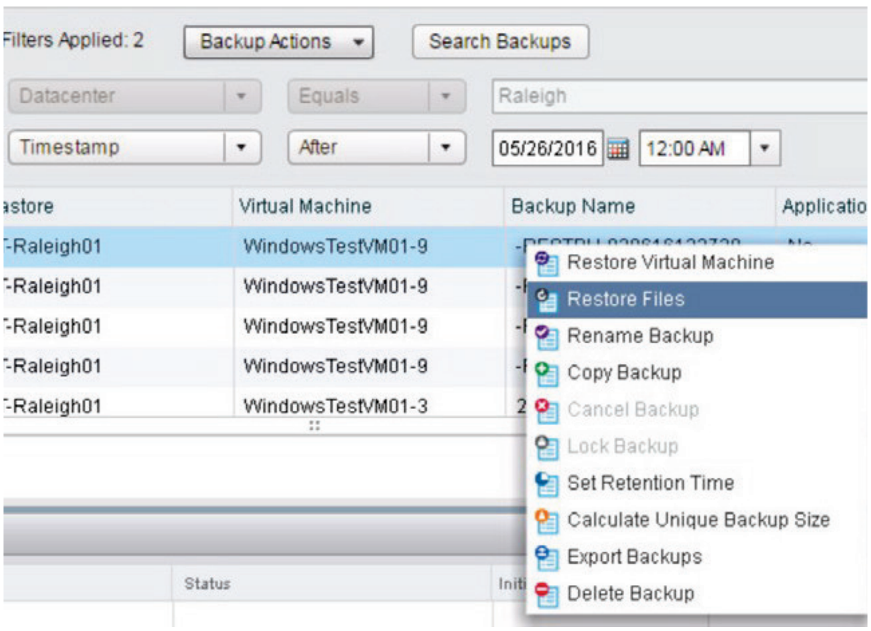

Figure 6. Simple restoration of an entire VM or specific files/folders.

The HPE SimpliVity approach delivers better economics and performance by reducing storage requirements, I/O, and WAN traffic. And it improves data protection by eliminating the constraints that the data problem imposes on the legacy IT stack, including eliminating the need to limit backup to nonproduction hours, thus accelerating backup and recovery functions (improving RPOs and RTOs).

Management capabilities and interfaces

The HPE SimpliVity hyperconverged infrastructure includes a complete set of management capabilities and interfaces. From a management perspective, HPE SimpliVity Clusters and Federations are treated as a cohesive system with a single interface. The entire system is administered in a unified manner from a centralized management application.

HPE SimpliVity offers GUI plug-ins for native hypervisor management applications, such as VMware vCenter, and third-party service orchestration tools, such as VMware vRealize Automation and Cisco UCS Director. The hyperconverged solution also supports a native command-line interface (CLI) and a programmatic interface for securely performing management functions via a terminal or scripts.

Management functions

Data protection and recovery

The HPE SimpliVity solution supports a full set of data protection and recovery capabilities. They can be initiated via the HPE SimpliVity CLI, programmatic interfaces, or GUI plug-ins for native hypervisor system management applications or service orchestration tools. Automated backup policies can be established to address specific recovery-point objectives (RPOs). Manual backups are also supported, as are interactive restores and VM cloning functions.

Backup policies can be configured as either a default policy at the datastore level, or a specific policy at the VM level. Data can be backed up locally within an HPE SimpliVity Cluster or remotely (to another HPE SimpliVity Cluster). Local backups are typically used to recover from equipment failures, application problems, or administrative mishaps. Remote backups are typically used for disaster recovery, creating isolated test versions of production environments, and long-term off-site retention.

Policies for scheduled backup operations

Backup policies are based on administratively defined rules. Each rule only requires the following at a minimum:

- **Backup frequency**—How often should backups be created
- **Backup retention**—Maximum timeframe an individual backup should be maintained
- **Backup destination**—Local or remote HPE SimpliVity Cluster where backup is to be maintained

If additional control needs to be defined, the following properties can also be defined for each rule:

- **Backup window**—Start and end times when the backup can be created
- **Backup days**—Days of the week or specific dates when the backup should be created
- **Backup type**—Whether the backup is application consistent (see the following)

Individual rules are combined to create diverse policies. For example, a single policy can contain rules that define daily, weekly, and monthly backups to multiple data centers, each with a separate retention period and destination. This allows a single policy to define local backups for point-in-time file restores, remote backups for disaster recovery, and cloud-based backups for long-term retention and compliance.

Policies and schedules can be viewed, modified, or deleted via the CLI, programmatic interfaces, native hypervisor system management, or service orchestration applications.

Application-consistent backups

HPE SimpliVity environments optionally support application-consistent backup for VMware vSphere environments to adequately protect transactional applications like Microsoft® SQL, Active Directory, or Exchange. The solution utilizes VMware Tools, which in turn use standard Microsoft Volume Shadow Copy Service (VSS) functions to synchronize the File System and applications. For Linux®-based VMs or applications that lack VSS integration, VMware Tools can be utilized to run scripts designed to prepare the application for backup.

For backing up Microsoft SQL Server instances, HPE SimpliVity provides app-aware backups that integrate directly with the VSS function of SQL Server to reduce the time and storage I/O overhead required to take properly quiesced databases backups. This approach to taking application-consistent backups also eliminates the need to rely on VMware Tools and vSphere-based snapshots.

Troubleshooting and diagnostics

HPE SimpliVity hyperconverged infrastructure provides a number of troubleshooting tools including alarms and events, component and connection heartbeats, and diagnostic files.

Alarms and events

The HPE SimpliVity solution can generate a wide variety of alarms and events for reporting the health and status of the hyperconverged infrastructure components and processes (for example, a backup operation). Alarms are automatically forwarded to and displayed on native hypervisor management system and third-party service orchestration tools via plug-ins. For example, the HPE SimpliVity GUI plug-in for VMware vSphere includes over 130 custom alerts for tracking VM-level events in the vSphere Client alarms panel, further enhancing the VM-centric design. HPE SimpliVity events, such as an HPE SimpliVity rapid clone, manual backup, or creation of a backup policy are tracked through the VMware vCenter task panel and logged in the vSphere logs. This streamlined approach allows existing upstream monitoring utilities (such as

VMware vRealize Operations) to rapidly begin ingesting status information from vCenter to begin monitoring the HPE SimpliVity systems, and to allow system administrators to be alerted to vSphere and HPE SimpliVity events through a common interface.

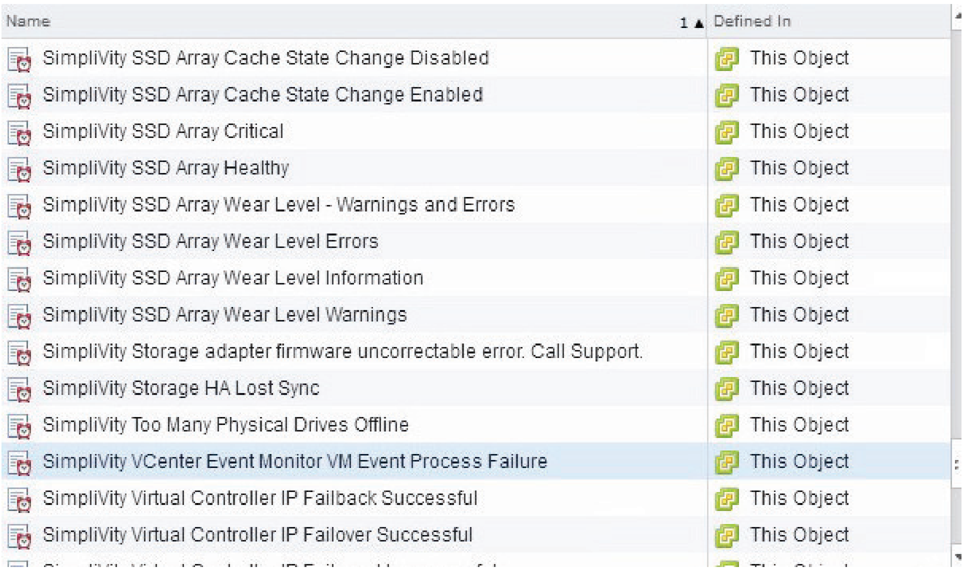

Figure 7. HPE SimpliVity alarms within the vSphere Web Client.

Diagnostic files

HPE SimpliVity hyperconverged nodes automatically generate detailed diagnostic files and logs in response to system events. These files are used by HPE Pointnext support and engineering personnel to isolate hardware or software issues. Diagnostic support capture bundles can also be generated on demand and shared with HPE SimpliVity support and engineering personnel for troubleshooting purposes. These support capture bundles include

- Status of hardware components (disks, controllers, power supplies, and so on)
- Logs from the HPE OmniStack Virtual Controllers that comprise the HPE SimpliVity Cluster
- Logs from the ESXi hosts that comprise the HPE SimpliVity Cluster
- Logs from vCenter relevant to the HPE SimpliVity Cluster

OmniWatch

OmniWatch, is the phone home service for HPE SimpliVity, which will forward critical alarms and event messages to HPE Pointnext. The feature improves customer service and system availability by proactively notifying support personnel of issues that may affect system performance or reliability. Based on these alerts, HPE Pointnext customer support cases may be automatically created, depending on the type of alert and criticality.

Alarms, events, and messages can be automatically forwarded to the HPE Pointnext customer support organization. Health and status messages provide information about the hyperconverged systems: the triggering alarm and alarm description, device name, serial number, current statuses of HPE OmniStack Accelerator Card, and so on.

OmniWatch capabilities and settings are fully configurable via the CLI, programmatic interfaces, or supported HPE SimpliVity plug-ins for native hypervisor system management.

Deployment and upgrades

HPE has created a set of tools and processes to simplify the deployment and upgrade of HPE SimpliVity hyperconverged nodes. These tools can be utilized by Remote Deployment Services and Customer Support teams (requiring no on-site presence by HPE personnel), HPE field personnel, certified partners, and end customers.

Deployment

The deployment process of a new HPE SimpliVity hyperconverged infrastructure starts with the preflight checklist, where a customer will indicate environment details necessary for deploying hyperconverged nodes. These details are entered into HPE SimpliVity Deployment Manager to create a configuration file to be used for the actual deployment. All the customer needs to do is have the server powered on and connected to the network. Deployment Manager automatically detects the nodes on the network, executes a series of predeployment checks on the environment, and can push down all the configuration necessary, including indicating the hypervisor to be installed. Predeployment checks include NIC connectivity and speed, Jumbo Frame settings, and TCP/IP port connectivity, with more being added with upcoming releases.

At the end of the Deployment Manager process, the node will have been added to the vCenter Server and will be ready for virtual machines.

Upgrades

HPE SimpliVity upgrades are completed either through the vSphere Client (pre 3.5) or through HPE SimpliVity Upgrade Manager (3.5+). Utilizing Upgrade Manager, HPE SimpliVity infrastructures can be upgraded with a simple process that allows the upgrade of an entire HPE SimpliVity Federation at one time in parallel with no application downtime.

Plug-ins for native hypervisor system management applications

Core to the platform of HPE SimpliVity is the concept of global unified management. This concept minimizes the number of disparate interfaces required to manage an environment, and distills it into a single pane of glass. HPE SimpliVity delivers this streamlined singular global unified management platform by interfacing with existing management frameworks used to manage the virtualized environment.

HPE SimpliVity is managed through a GUI plug-in for the VMware vSphere Web Client. Central administrators can implement data protection policies, monitor resources, and troubleshoot problems utilizing the same vSphere Client they already use to manage their vSphere environment. All administrative functions are performed across an HPE SimpliVity Cluster or Federation in a holistic manner. The VMware vSphere Web Client plug-in simplifies operations and reduces the learning curve by allowing IT organizations to extend existing administrative systems and practices to HPE SimpliVity technology.

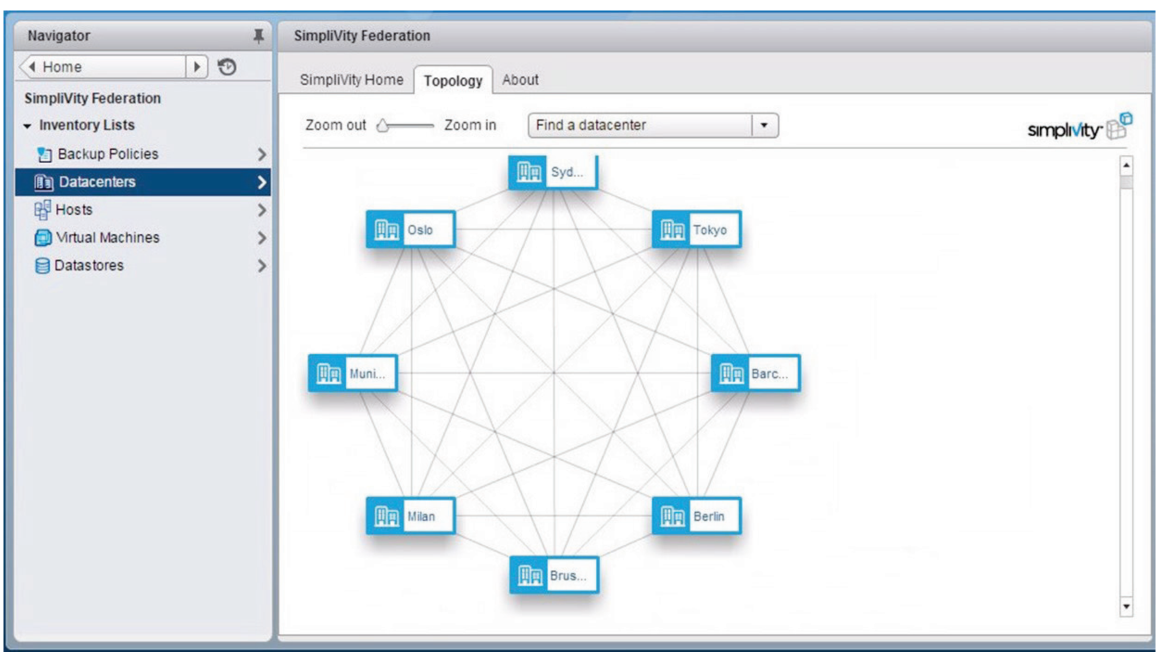

Figure 8. vSphere Web Client integration.

Managing HPE SimpliVity Clusters, Federations, and nodes

The HPE SimpliVity vSphere Web Client plug-in includes tools for defining and managing HPE SimpliVity Clusters and Federations. Administrators can add, move, and delete nodes from HPE SimpliVity Clusters and configure certain node settings.

Federation-level management is enabled by simply adding hosts into different vSphere clusters. These clusters can be managed by a single vCenter Server or by multiple vCenter Servers joined together via vSphere Linked Mode. There is no additional configuration necessary to create an HPE SimpliVity Federation that can span across the globe.

The vSphere Web Client provides configuration, inventory, and status displays for individual HPE SimpliVity hyperconverged nodes.

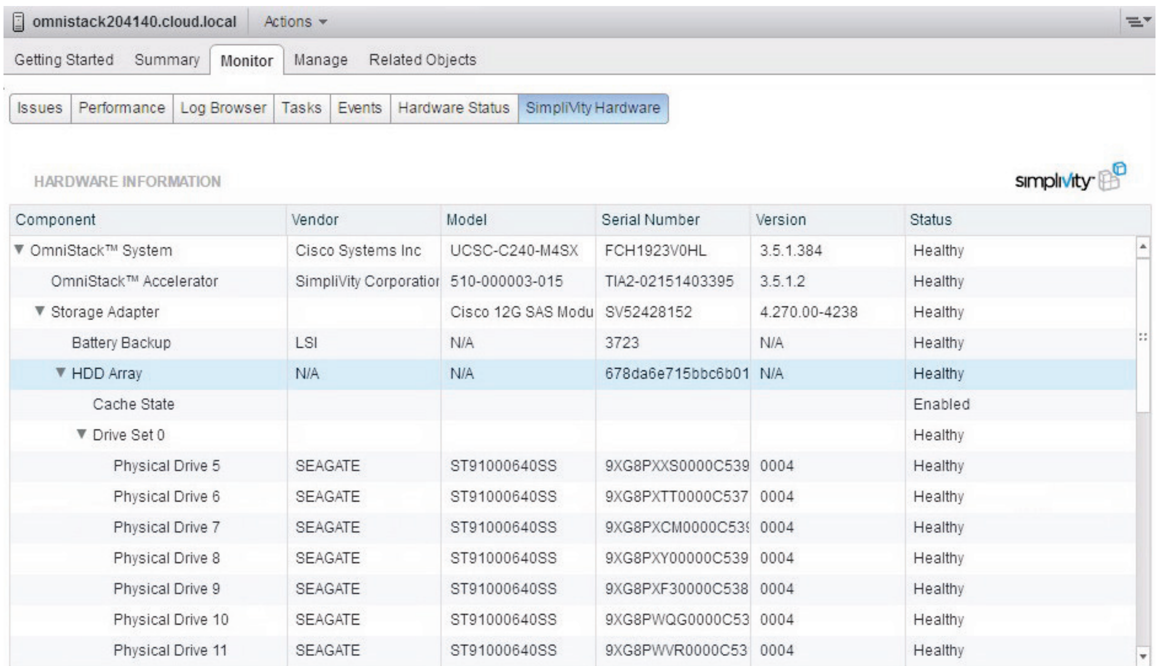

Component	Vendor	Model	Serial Number	Version	Status
▼ OmniStack™ System	Cisco Systems Inc	UCSC-C240-M4SX	FCH1923V0HL	3.5.1.384	Healthy
OmniStack™ Accelerator	SimpliVity Corporatior	510-000003-015	TIA2-02151403395	3.5.1.2	Healthy
▼ Storage Adapter		Cisco 12G SAS Modu	SV52428152	4.270.00-4238	Healthy
Battery Backup	LSI	N/A	3723	N/A	Healthy
▼ HDD Array	N/A	N/A	678da6e715bbc6b01	N/A	Healthy
Cache State					Enabled
▼ Drive Set 0					Healthy
Physical Drive 5	SEAGATE	ST91000640SS	9XG8PXXS0000C539	0004	Healthy
Physical Drive 6	SEAGATE	ST91000640SS	9XG8PXTT0000C537	0004	Healthy
Physical Drive 7	SEAGATE	ST91000640SS	9XG8PXCM0000C53(	0004	Healthy
Physical Drive 8	SEAGATE	ST91000640SS	9XG8PXY00000C539	0004	Healthy
Physical Drive 9	SEAGATE	ST91000640SS	9XG8PXF30000C538	0004	Healthy
Physical Drive 10	SEAGATE	ST91000640SS	9XG8PWQG0000C53	0004	Healthy
Physical Drive 11	SEAGATE	ST91000640SS	9XG8PWVR0000C53	0004	Healthy

Figure 9. HPE SimpliVity hyperconverged node view within vSphere Web Client.

Capacity and performance management

The HPE SimpliVity vSphere Web Client plug-in provides summary and detailed performance and utilization displays for HPE SimpliVity Clusters and Federations including physical and logical storage capacity views (free, used, total space, and so on) and VM performance views (throughput, IOPS, and latency strip charts).

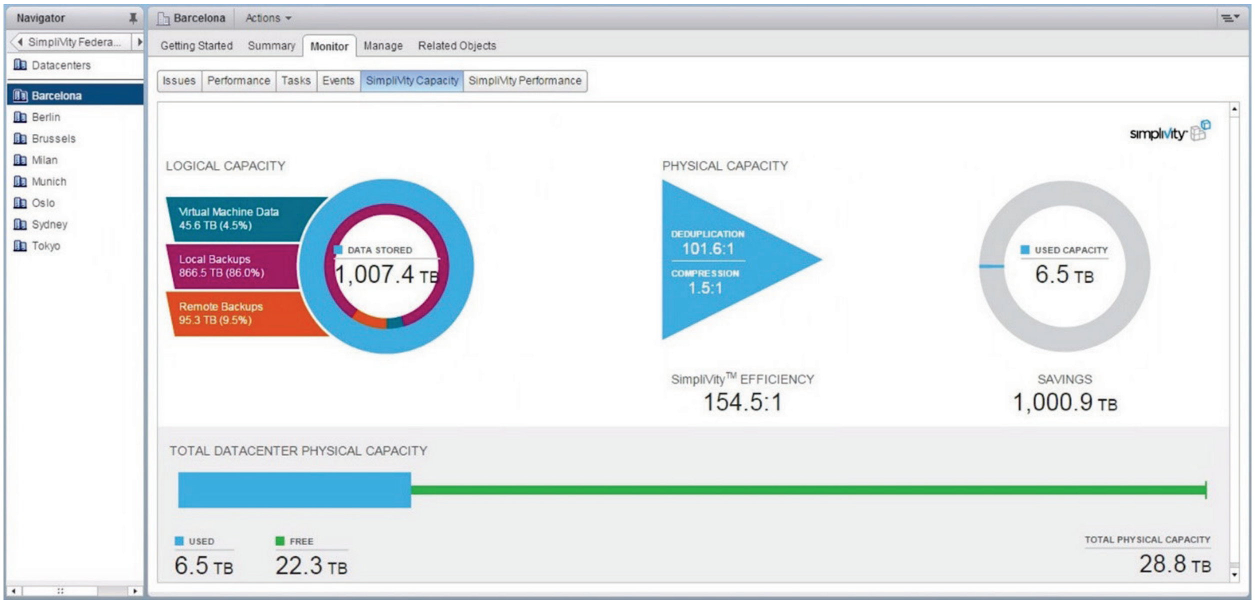

Figure 10. HPE SimpliVity capacity view within vSphere Web Client.

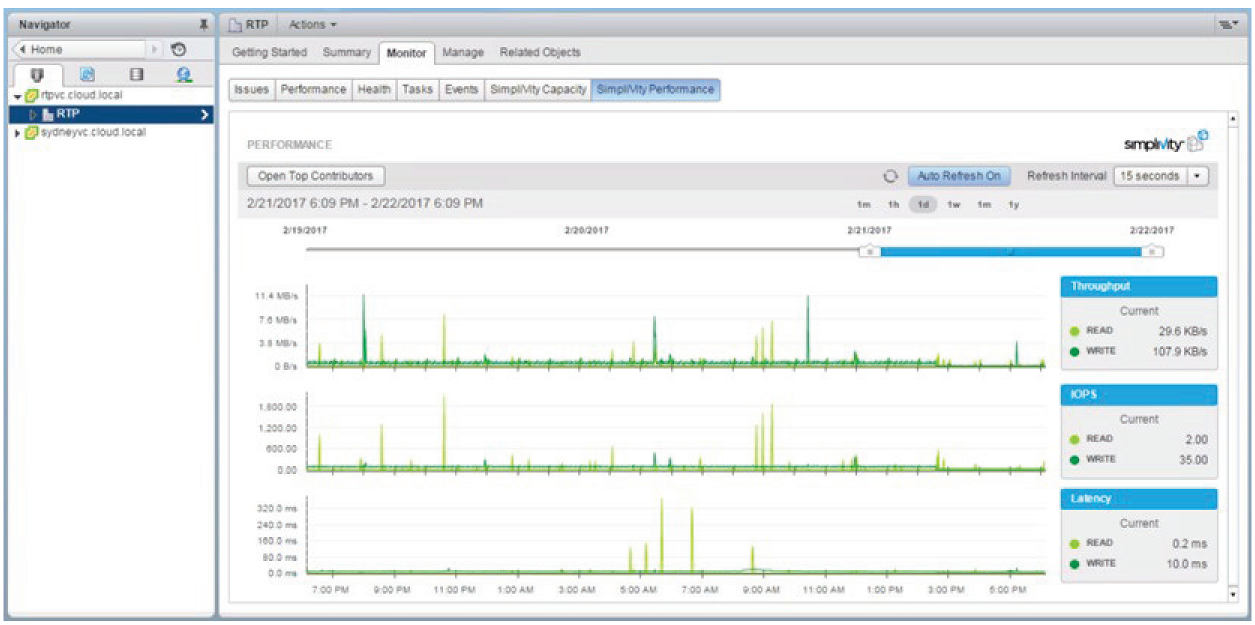

Figure 11. HPE SimpliVity performance view within vSphere Web Client.

Intelligent Workload Optimizer

To provide the best application performance across an HPE SimpliVity Cluster, a component called Intelligent Workload Optimizer will intelligently place the replica set of a VM across HPE SimpliVity nodes when it is created and integrate with vSphere Distributed Resources Scheduler (DRS) to ensure optimal placement of the VM compute resources.

When a VM is created, Intelligent Workload Optimizer will review the storage capacity and storage I/O on all HPE SimpliVity hyperconverged nodes over last seven days to determine the initial placement of the two parts of the replica set. If multiple Availability Zones exist, Intelligent Workload Optimizer will ensure the two selected HPE SimpliVity hyperconverged nodes are in different Availability Zones to maintain separation of the replica set for stretch cluster configurations.

The integration with DRS solves a fundamental flaw with hyperconvergence in a DRS-controlled vSphere cluster. DRS is designed to optimize the overall performance of the vSphere cluster by balancing CPU, memory, and network resources across vSphere hosts to ensure every VM receives the resources they require. Unfortunately, this was designed for legacy infrastructure and assumes that the storage is hosted externally to the hosts, and therefore does not consider one of the most important advantages of hyperconverged infrastructure: data locality.

By bringing awareness of the location of the data within the HPE SimpliVity Cluster to DRS, Intelligent Workload Optimizer can help DRS to maximize the performance of storage access in addition to the CPU, memory, and network resources. Intelligent Workload Optimizer creates and maintains DRS Host and VM Groups that are combined into **Suggested Affinity** DRS Rules so that DRS will prefer, but not require, that VMs remain on the two nodes that host the VM's data. Intelligent Workload Optimizer maintains these rules to ensure that they are kept up to date throughout the life of the VM. On a regular cadence, Intelligent Workload Optimizer will check the validity of the groups and rules, and correct any manual changes that were inadvertently made by an administrator.

This integration into DRS ensures that VMs and their data remain collocated. This delivers maximum storage performance relative to other approaches of managing the data in the cluster in response to changes made by DRS. The first of these alternative approaches is to not maintain any centralized copy of the VM, which avoids any need for data locality knowledge. However, this approach does require all storage retrieval to be network-based at all times, thus eliminating the advantages of data locality. The second approach is to have the data follow the VM every time it is migrated by DRS. This follow the VM approach maintains data locality without having to integrate with DRS, but results in inconsistent performance until all data is migrated locally and produces unnecessary CPU, network, and storage usage that can hurt the overall performance of the cluster.

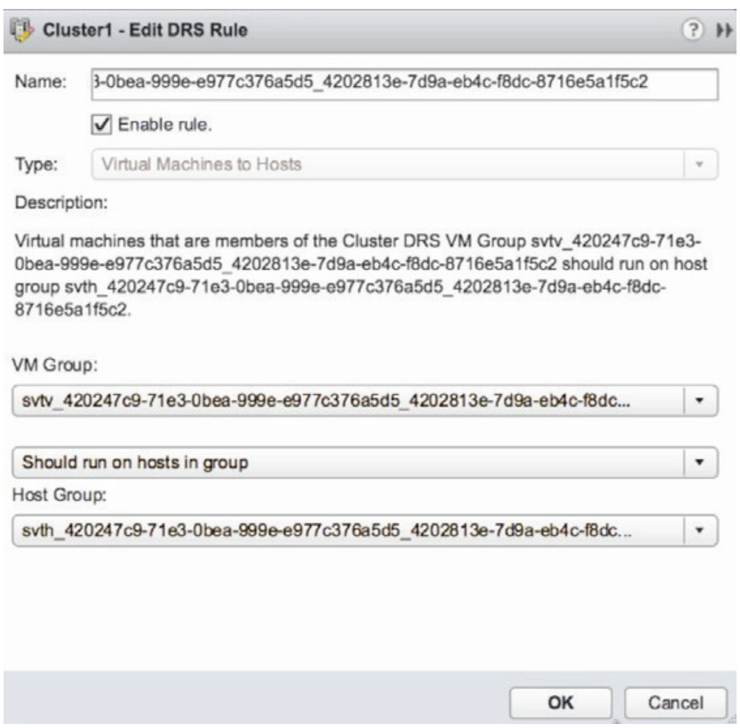

Figure 12. HPE SimpliVity performance view within vSphere Web Client.

Data protection and recovery tasks

The vSphere Web Client plug-in supports a full set of data protection and recovery management capabilities. Administrators can use the GUI to establish and manage automated backup policies, to clone VMs and to manually back up and restore VMs (see the <u>Data protection and recovery</u> section for more details).

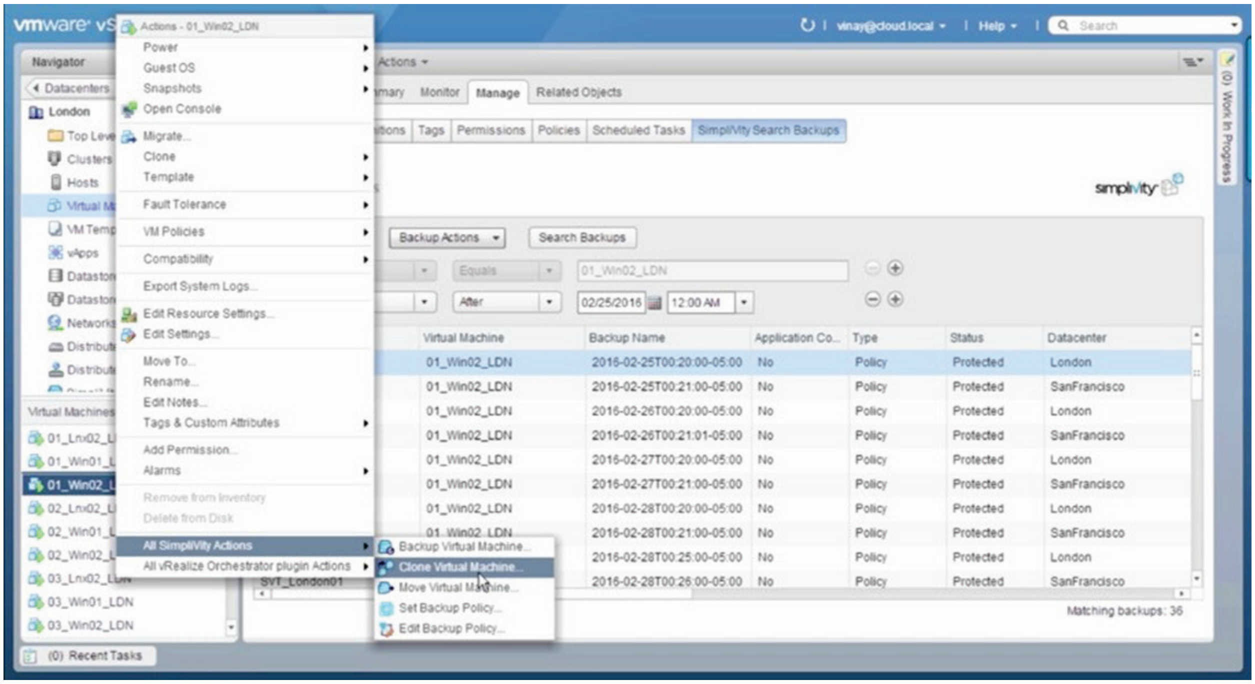

Figure 13. HPE SimpliVity manual data protection in the vSphere Web Client.

Plug-ins for third-party service orchestration applications

HPE offers plug-ins for third-party service orchestration applications such as VMware vRealize Automation and Cisco UCS Director to better integrate HPE SimpliVity functionality. With service orchestration applications, users (typically application developers and application administrators) can create and/or allocate compute resources and applications without IT involvement, using self-service portals and service catalogs.

HPE SimpliVity Data Virtualization Platform simplifies the configuration of automated workflows by reducing the amount of disparate infrastructure technologies from as many as 12 to 1. With the integrated solution with HPE SimpliVity, users can also back up, restore, clone, and move VMs on demand using the service orchestration solution's user interface—all under strict administrative control of the central IT organization. The integrated solution helps IT organizations accelerate service velocity, improve business agility, and reduce operational expenses by automating IT service deployment and lifecycle management tasks.

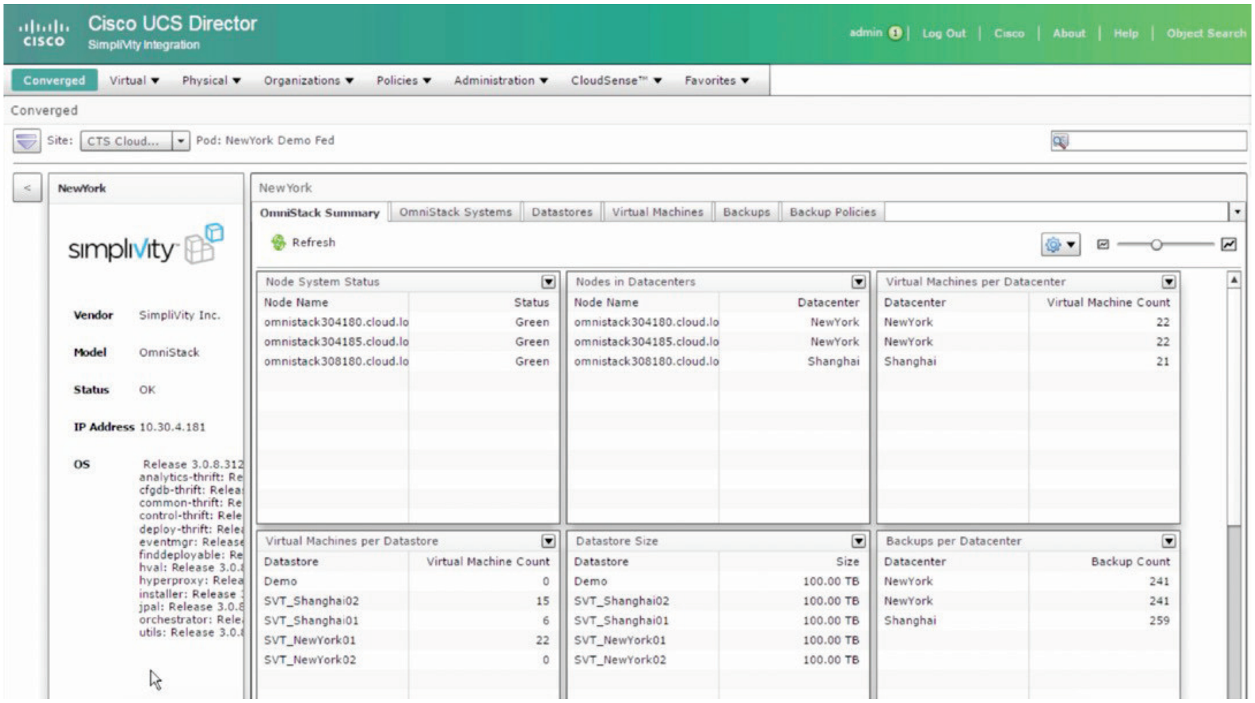

Figure 14. HPE SimpliVity integration with Cisco UCS Director.

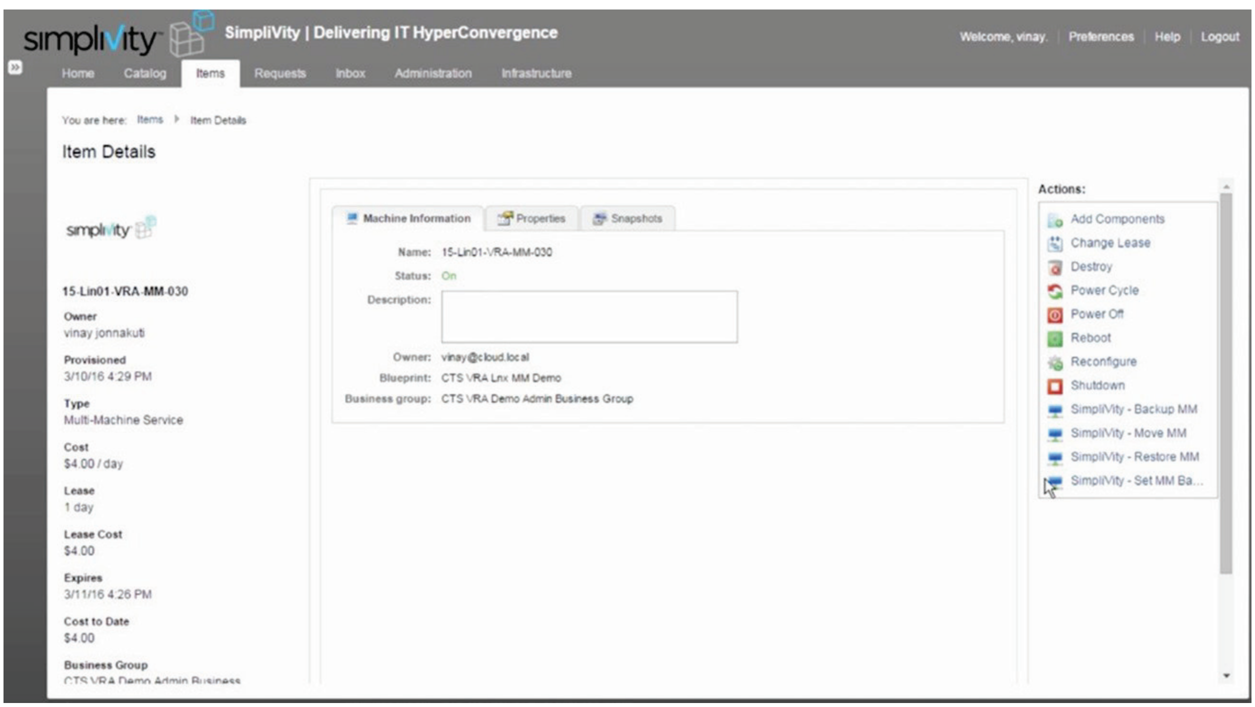

Figure 15. HPE SimpliVity integration with VMware vRealize Automation.

Native management interfaces and APIs

Command-line interface

HPE SimpliVity hyperconverged infrastructure supports a complete command-line interface (CLI) for configuring system options and settings; for initiating data protection, recovery, and replication functions; and for performing routine maintenance and diagnostic functions. The CLI can be executed interactively via a terminal interface or programmatically via scripts.

The CLI is designed to work in concert with HPE SimpliVity GUI plug-ins for native hypervisor system management applications, such as VMware vCenter, or service orchestration tools, such as VMware vRealize Automation and Cisco UCS Director. Any configuration changes made using the CLI are immediately visible in the native hypervisor management applications. Conversely, any configuration changes made using native hypervisor management applications are immediately visible in the CLI.

The CLI runs on the HPE OmniStack Virtual Controller and is accessed via a secure shell (SSH) session from a terminal emulator. A single CLI session (to one HPE OVC) provides full visibility and management across an entire HPE SimpliVity Federation.

Examples of CLI command sets include

- **HPE SimpliVity Federation and Cluster commands**—Organize and manage nodes within an HPE SimpliVity Federation or Cluster
- **Backup policy commands**—Create, view, and manage backup rules and policies
- **Datastore commands**—Create, manage, remove, and apply backup policies to datastores
- **VM commands**—Interactively back up, clone, restore, and move VMs; apply backup policies to VMs
- **Software management commands**—Manage HPE OmniStack software releases (upgrade, roll back, commit)
- **Support commands**—Configure phone home settings and other maintenance options

Programmatic APIs

Built on top of the existing CLI, the HPE SimpliVity programmatic interfaces provide access to the full suite of VM management functions. Important functions like creating HPE SimpliVity clones and backups can be automated utilizing existing orchestration tools or by developing all new applications. Currently, both an XML API and a REST API are available. The XML API is currently utilized for prebuilt automation plug-ins to Cisco UCS Director and VMware vRealize Automation and in custom applications. The REST API provides access to HPE SimpliVity functionality through a simple HTTP-based interface that works easily with almost any modern automation and orchestration tools. Included with every HPE OmniStack Virtual Controller is a web-based tool that allows access to the full set of REST API functions, including the ability to execute operations and access to the code necessary to call these operations via HTTP and cURL. This allows administrators and automation engineers to be able to rapidly prototype automated HPE SimpliVity functionality.

Please see the official HPE SimpliVity documentation for details on utilizing these APIs. Interactive documentation for the REST API is available on every HPE SimpliVity OmniStack Virtual Controller at https://<OVC_mgmt_IP>/api/index.html and at api.simplivity.com.

OmniView

To improve visibility of metrics within an HPE SimpliVity-based vSphere environment, HPE created a tool called OmniView. This tool was originally developed for use within HPE Pointnext to help with the monitoring and diagnosis of large customer environments but is now available for direct customer utilization. Once customers with the proper support contract opt in to the service, both HPE SimpliVity and vSphere data will be collected by the HPE OmniStack Virtual Controller and forwarded to a support cloud database. Analytics are run on this data and the resulting analytics can be accessed through a secure web interface on the HPE SimpliVity Online Support Center.

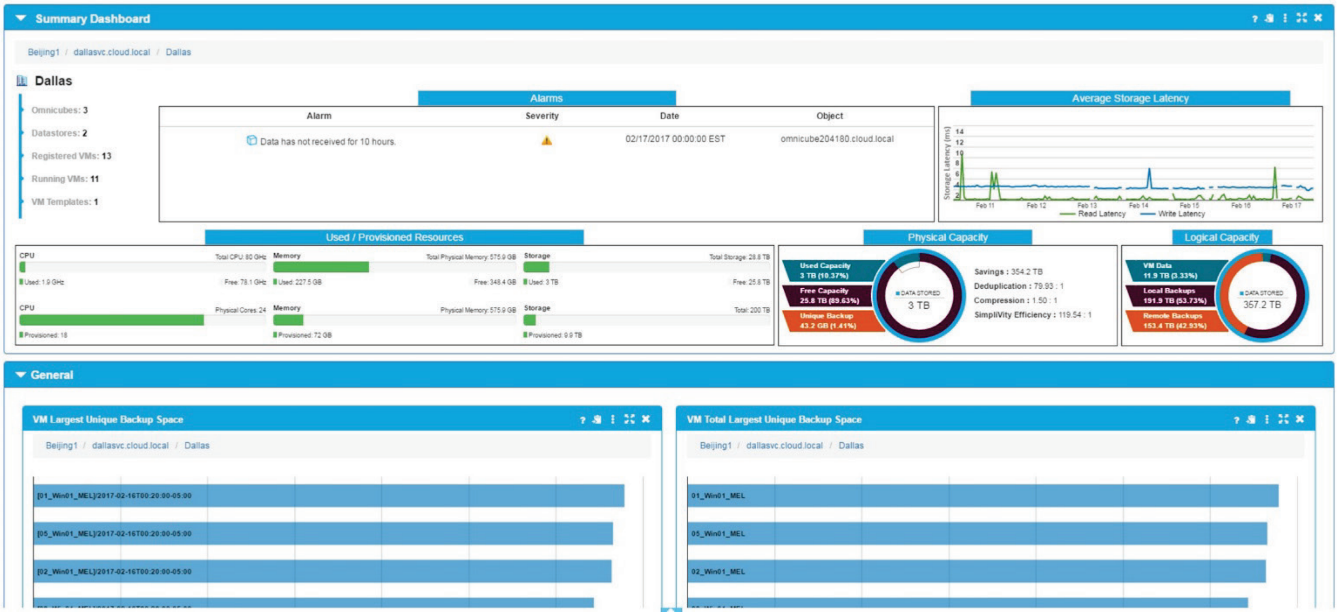

Figure 16. OmniView data center dashboard.

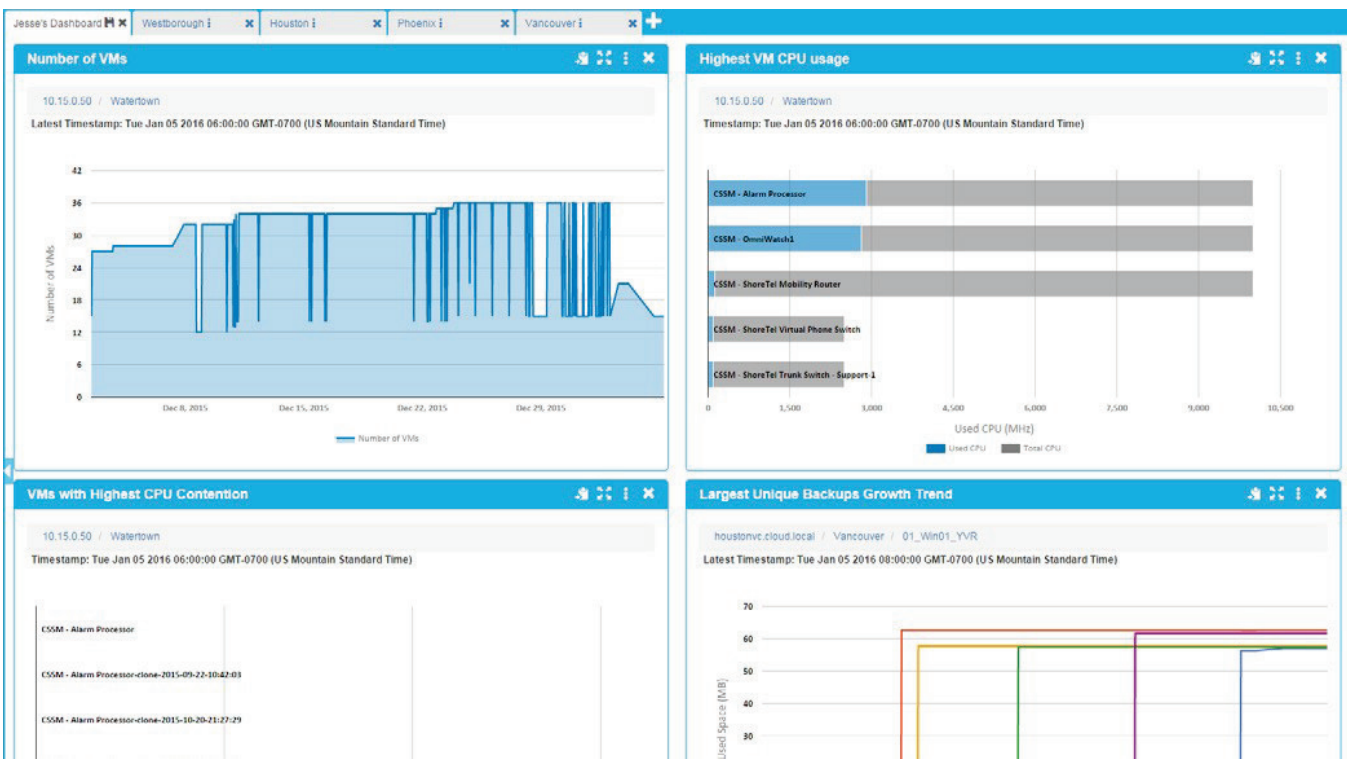

Figure 17. OmniView custom dashboard.

The collected data is analyzed by a cloud-based analytics engine that provides improved insight into a customer's environment without requiring additional resources from the customer's on-premises infrastructure. The resulting graphs, charts, and trends provide a robust yet simplified view into the customer's global hyperconverged infrastructure, by displaying combined data from HPE SimpliVity and vSphere metrics collected across multiple data centers and HPE SimpliVity Federations.

This data can be utilized to better monitor current and historical data, diagnose issues, and view predicted utilization trends to improve plans for future expansion.

HPE SimpliVity hyperconverged infrastructure benefits

HPE SimpliVity systems are built to provide the best of both worlds—enterprise-grade capabilities and cloud-like economics and flexibility. These benefits include the following:

Enterprise capabilities

- **Improved business continuity and disaster recovery**—Highly resilient architecture and rapid data protection features reduce risks and uncertainty and improve availability for business-critical applications.
- **Improved service velocity**—Integration with system management and service orchestration tools improves time to value and time to market.
- **Better application performance**—The unique data management approach of HPE SimpliVity coupled with hardware-assisted deduplication, compression, and optimization functionality minimizes latency, frees up x86 CPU resources, and provides both peak application performance as well as very predictable performance as application workloads increase.
- **Simplified administration**—HPE SimpliVity architecture unifies global management and integrates into established administrative systems, thereby streamlining operations.

Cloud economics

- **Lower CAPEX**—HPE SimpliVity hyperconverged infrastructure reduces hardware expenses by consolidating the entire IT stack below the hypervisor onto a single hardware platform and by making optimal use of storage, compute, and networking resources.
- **Lower OPEX**—HPE SimpliVity hyperconverged infrastructure and unified management enables lower ongoing power, cooling, and rack space costs; lower recurring operational expenses; and lower recurring product maintenance and support fees.
- **VM-centric management**—Management designed around the VM, utilizing policy-based controls and abstraction from the underlying infrastructure allow simple VM management and easy VM mobility.
- **Better business innovation**—HPE SimpliVity infrastructure allows IT organizations to spend less time managing underlying IT infrastructure and more time focusing on strategic business initiatives.
- **Greater flexibility**—HPE SimpliVity provides choice in hardware platforms, hypervisors, management systems, and deployment options.
- **Invest-as-you-grow economics**—The scale-out architecture minimizes upfront capital investments by right-sizing the initial investment, and it avoids lopsided business models with long payback periods.

Hyperconverged infrastructure use cases

HPE SimpliVity hyperconverged infrastructure supports a wide variety of use case scenarios including the following:

- **Private cloud/data center consolidation**—Customers can contain TCO and modernize infrastructure by converging all IT below the hypervisor in a single 2U hardware platform with a scale-out architecture and unified, cohesive management.
- **Data protection**—Customers can simplify operations and mitigate risks with policy-driven VM-level data protection as well as execute data backup and recovery functions more efficiently, quickly, and frequently. On average, HPE SimpliVity customers achieve 40:1 data efficiency while fully protecting their data.[2]
- **Data migration**—The VM-centric approach to data management and efficient data architecture enables faster time to deployment and time to market.
- **Tier-1 applications**—Customers can optimize performance, availability, and economics of business-critical applications.
- **Dev/test**—Customers can eliminate cost and inefficiencies by converging separate development and test/QA environments, reducing the administrative overhead of separate technology silos, and reducing the time to clone workloads.
- **Unified protected remote office/branch office (ROBO)**—Customers can centralize administration and improve data protection for small sites and remote locations.
- **Virtual desktop infrastructure**—Customers can optimize performance and user experience for VDI initiatives with low cost of entry and scaling in small or large increments.
- **Data center migration**—Utilizing the unique data migration techniques provided by HPE SimpliVity, customers can drastically reduce the time and effort necessary to migrate from one data center to another. Automation through cloud management or automation tools can further accelerate this time.

[2] Real Data on Data Efficiency; "The True Hyperconvergence Ticker Explained," SimpliVity (now Hewlett Packard Enterprise), March 2015

Index

Appendix: List of abbreviations

API—Application programming interface

CLI—Command-line interface

CPU—Central processing unit

DRAM—Dynamic random-access memory

DRS—Dynamic Resource Scheduler

DVP—Data Virtualization Platform

FPGA—Field-programmable gate array

GUI—Graphical user interface

HDD—Hard disk drive

IOPS—Input/output operations per second

LUN—Logical unit number

NFS—Network file system

NIC—Network interface card

OAC—HPE OmniStack Accelerator Card

OVC—HPE OmniStack Virtual Controller

RAM—Random-access memory

RAID—Redundant array of independent disks

RAIN—Redundant Array of Independent Nodes

REST—Representational State Transfer

RPO—Recovery-point objective

RTO—Recovery-time objective

SLA—Service-level agreement

SSD—Solid-state drive

SSH—Secure shell

VDI—Virtual desktop infrastructure

VM—Virtual machine

WAN—Wide area network

Learn more at
hpe.com/info/simplivity

a00019351ENW, July 2017