Note: Coverage of HCI systems has been discontinued. The HPE Simplivity Product Analysis is no longer receiving updates as of September, 2023.

HPE SimpliVity

HPE SimpliVity Hyperconverged Infrastructure (HCI) appliances run the proprietary OmniStack software on the HPE DL380, DL325 and Apollo 2000 platforms. Each appliance is connected via Ethernet to form a scale-out, virtual SAN architecture using internal all-flash storage. Together with an embedded hypervisor, these appliances cluster to create a comprehensive compute platform for running virtual server workloads. Like other HCI solutions, HPE SimpliVity provides a standardized, agile infrastructure for supporting a wide range of potential use cases including ROBO, VDI, server virtualization and private cloud, to name a few.

OmniStack was the HCI solution developed by Westborough, MA start up SimpliVity that HPE acquired in early 2017. Competitive solutions include Nutanix Cloud Platform, EMC VxRail, Cisco HyperFlex and VMware vSAN Ready Nodes. These and the other leading products are included in the Evaluator Group Research Hyperconverged comparison matrix and associated Product Briefs, Analyses and Evaluation Guide.

Highlights

- Characteristics
  - Performance – PCIe off-load card in HPE SimpliVity 380 models removes deduplication overhead from CPU and provides at-ingest data reduction, decreasing data handling overhead. Other SimpliVity models handle data reduction in software without dedicated hardware cards. All models are all-flash storage exclusively
  - Global capacity optimization – OmniStack uses global deduplication and compression across all nodes in a federation
  - Data Protection – Built-in policy-driven protection for VM-level backup/recovery with replication
  - HPE Cloud Volumes Backup allows direct backup from HPE SimpliVity from edge sites.
  - HPE StoreOnce integration protects distributed edge sites by replicating applications and data to a centralized StoreOnce appliance for long-term retention.
  - SQL Server specific backup feature
  - One-step VM recovery for DR
  - Single node (SimpliVity 380), 2-4-node (SimpliVity2600) and edge (SimpliVity 325) models
  - GPUs supported (380 and 2-node 2600)
- Applications
  - An SMB to Enterprise IT infrastructure solution
  - Virtual Desktop Infrastructure
  - Private/Hybrid cloud
  - Remote Office/Branch Office

  - Non-datacenter, “in the field” deployments
  - Test/Dev
- System Environments
  - Datacenter consolidation
  - Server/desktop virtualization environments
- Deployment and Administration
  - Global Federated Management – OmniStack clusters are managed through a single unified VM-centric interface.
  - HPE Composable Fabric integration provides SDN and automated network management
- Container Support
  - Kubernetes CSI Plugin allows SimpliVity HCI to run containers and virtual machine workloads in a unified platform.

Evaluator Group comment:

SimpliVity, founded in 2007, was one of two major independent vendors in the HCI appliance space along with Nutanix, and helped to establish the hyperconverged market segment. The OmniStack software is not an adaptation of an open source SDS solution but designed “from the ground up”, a decision the company claims delayed its release by two years. Headquartered in Westborough, MA, the company built a name for itself in this crowded market, before being acquired by HPE in 2017.

Overview of System

Like other Hyperconverged Infrastructure (HCI) solutions, HPE SimpliVity appliances are designed to provide simple deployment plus scalability, flexibility and ease of use. The HPE SimpliVity 380 is available in single-node models with different capacity points (see Figure 2), each with multiple CPU/memory options. The 380 H is available with 3.5-inch drives for backup and archiving, as well as 2.5-inch drives. The HPE SimpliVity 2600 comes configured with up to 4 x 1U nodes, or 2 x 2U nodes in a 2U chassis. HPE SimpliVity 325 is a 1U single-node appliance for edge deployments, ROBO sites and VDI. The HPE SimpliVity software (which SimpliVity referred to as OmniStack) is the software-defined storage foundation of this product.

Figure 1: HPE SimpliVity 380 HCI Appliance

Hardware Architecture

An HPE SimpliVity hyperconverged cluster is comprised of server chassis, each housing a single node (380 and 325) or up to four nodes (2600) in a scale-out topology. Each node has single or dual, multi-core CPUs and dual 10GbE data ports, and each chassis has redundant primary hardware components, including fans, power supplies and dual ported NICs (dual NIC cards are availalble as well). The hardware platform for the HPE SimpliVity 380 is the HPE DL380 Gen10, for the HPE SimpliVity 2600, the HPE Apollo 2600 Gen10 and for the HPE SimpliVity 325, the 1U DL325 Gen10 with a single AMD EPYC 7003 processor.

Model	Nodes / Appliance	CPU cores/Node	Max Memory/Node	Storage Devices per Node
HPE SimpliVity 325 Edge Node	1	16 - 64 cores*	2048GB	4, 6 or 8 SSD
HPE SimpliVity 380 H	1	8 – 56 cores	3072GB	4 SSD, 8-20 HDD
HPE SimpliVity 380 G	1	8 – 56 cores	3072GB	6/8/12/16/24 SSD
HPE SimpliVity 2600	2-4	16 - 52 cores	1024GB	6 SSD

* Single AMD EPYC 7003 CPU

Figure 2: HPE SimpliVity HCI Model Configurations

Scalability

HPE SimpliVity clusters normally have a minimum of two nodes and maximum of sixteen although HPE enables a federation of clusters under common management with up to 96 nodes. In a two-node cluster an Arbiter software running on a separate Windows system is used to facilitate communication between nodes and resolve state conflicts. While it's not required, using an Arbiter in larger clusters is recommended to serve as a "tie breaker," providing additional resiliency in even numbered configurations. A single node can also be used in small deployments or remote locations, as data is still RAID protected within each node.

Storage capacity in an HPE SimpliVity cluster can't be scaled independently from compute resources but nodes can be configured with more/larger flash drives and minimal CPU and memory. HPE SimpliVity doesn't support storage that's direct-attached to nodes in the cluster and they don't offer "storage-only" nodes, dedicated storage-heavy configurations with limited compute capabilities. That said, HPE SimpliVity does expose an NFS mount point to "compute nodes" or "legacy hosts" running non-virtualized applications and makes these external servers limited members of the cluster.

Scaling a cluster involves adding or removing nodes. Like most HCIs, the HPE SimpliVity 325, 380 and 2600, don't support adding individual drives to a partially full chassis. And, like other HCIs, HPE SimpliVity uses their data abstraction software to move data onto the new nodes and balance the available capacity in the background, making it a non-disruptive scaling process. When nodes are removed, the OmniStack cluster handles that process the same way, moving data away from the nodes being decommissioned as a background process. (see "Global Unified Management").

PCIe Accelerator Card

The original SimpliVity design included an on-board, PCIe-connected accelerator card that performs data reduction at ingest to eliminate duplicate data and minimize capacity consumption and data handling overhead throughout the cluster. This card provides in-line deduplication, compression and write optimization, offloading these functions from the system CPU in order to minimize impact on workload performance. It contains an FPGA, NVRAM and flash storage. In addition to data reduction, this card also aggregates random writes into full RAID stripes which are then sequentially written to SSD in a RAID 5 or RAID 6 format.

However, HPE has converted the data reduction and other functions of the accelerator card to run on the system CPUs through software optimization. This change was essential for the smaller form-factor nodes in the 2600 (4 nodes per 2U chassis) and the 1U 325 node, which don't have space to support a PCIe card on each node. The company has also eliminated the PCIe accelerator card from the SimpliVity 380 H Backup and 380 G models to save cost or to enable the support of more GPUs as the 380 G model does.

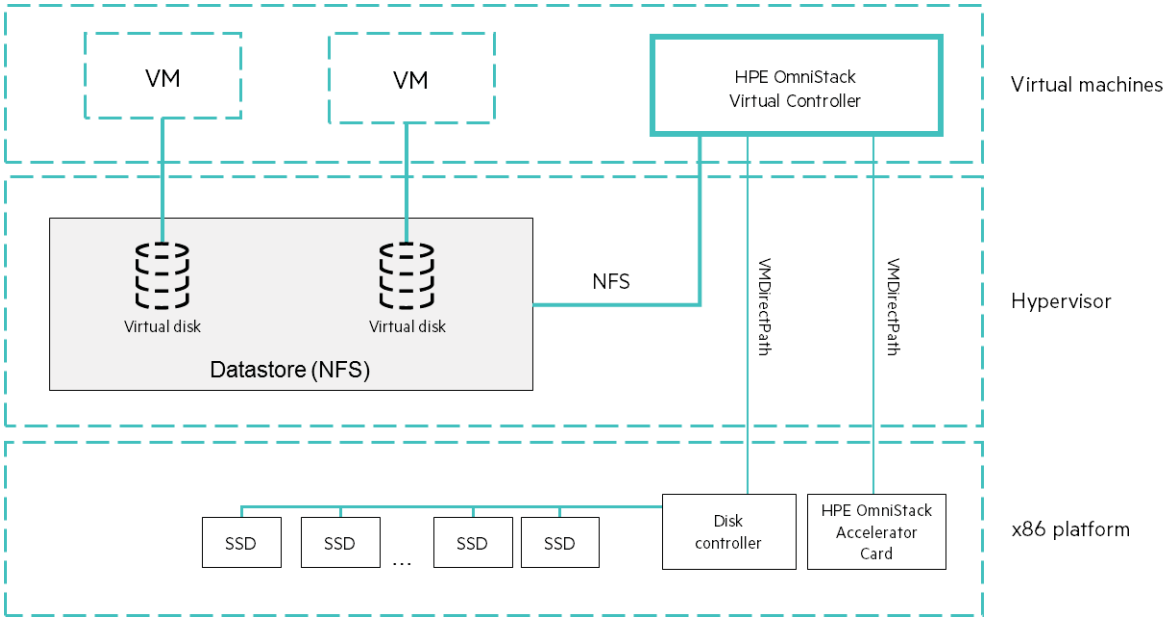

Figure 3: Components of HPE SimpliVity Node
(graphics in this report courtesy of HPE SimpliVity)

Software Architecture

The OmniStack Virtual Controller (OVC) is the foundational software element for the HPE SimpliVity HCI appliance and runs as a VM on every node. The OVC controls all aspects of data operation within each node and manages the communication with other nodes in the cluster (see Figure 3 above).

Data Virtualization Platform

The Data Virtualization Platform (DVP) is the virtual SAN architecture that creates a fabric across HPE SimpliVity nodes in a single cluster or across separate sites (see Figure 4 below). Using the OVC and data reduction functions (in software or on the accelerator card), it combines and abstracts physical storage devices within each node and between nodes in the cluster, presenting that capacity as a single pool of storage to the hypervisor via NFS. The DVP is responsible for maintaining data efficiency with in-line data reduction and data mobility, and replicating VM data sets between nodes to maintain resiliency (see "Data Resiliency"). This software also provides built-in data protection, leveraging metadata-based copies without creating clones, and a VM-centric policy-based data management across the entire environment, local and remote.

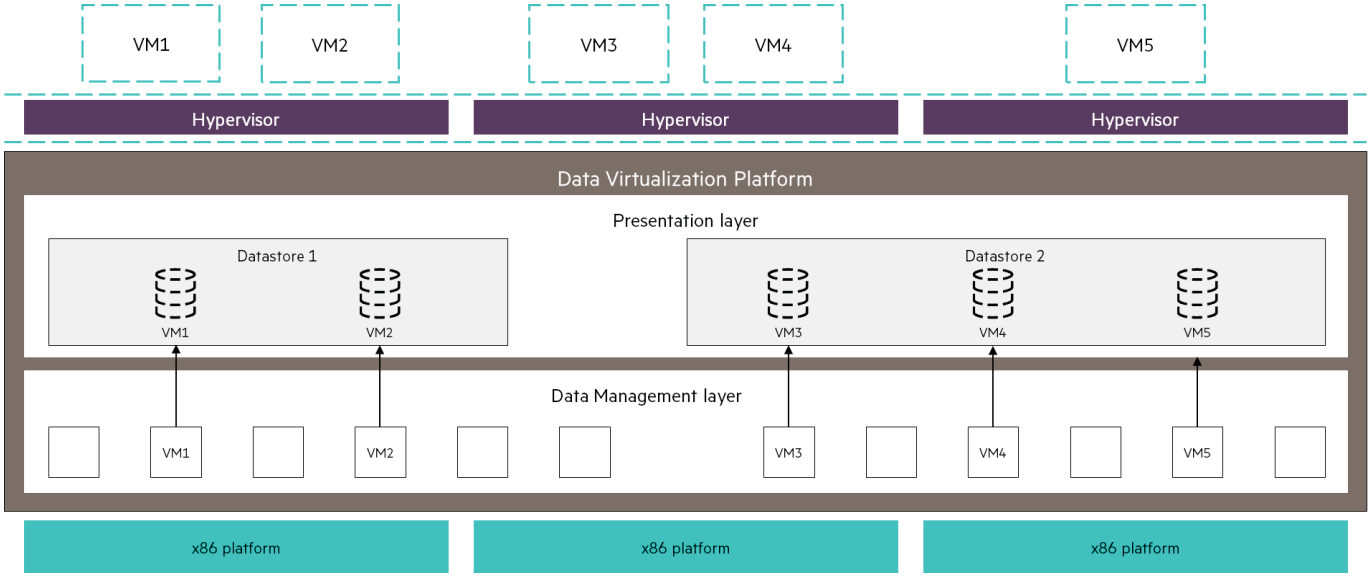

Figure 4: Data Virtualization Platform

The DVP is comprised of the Presentation Layer, which includes NFS, vStorage API for Array Integration (VAAI) and a hypervisor interface, plus the Data Management Layer, comprised of a file system on top of an object storage system. The Presentation Layer maintains data stores that support virtual disk drives for each host's VMs. These data stores are connected to the OVC which provides access to physical storage via the Data Management Layer. The Presentation Layer also interfaces with the hypervisor, of which only VMware vSphere is currently supported. HPE SimpliVity has its own globally-aware file system and object store to accommodate the small files and I/O intensive activity common in HCIs.

The DVP architecture puts the OmniStack Virtual Controller in front of physical storage in the data path. This enables it to do all the metadata operations - abstracting blocks for VMs and containers, creating pointers to unique data blocks, even synchronizing data with other nodes - before writing any data to physical storage. Then, the Data Virtualization Platform runs all data thru the OmniStack Acceleration Card, which does the data reduction process described below.

Performance

The HPE SimpliVity Data Virtualization Platform uses a caching process with DRAM plus an all-flash-only storage architecture to increase system performance. In the Data Management Layer, the OVC controls read and write activity, using the accelerator card to optimize data with deduplication, compression and write aggregation. This process improves the efficiency of storage media and increases overall system performance.

Data locality (assuring that a VM and its supporting data are physically stored on the same node) is an important factor in maintaining consistent performance and something that's addressed by HPE SimpliVity's Intelligent Workload Optimizer feature (see "Advanced Features and Functions").

Evaluator Group comment:

HCIs include compute as well as storage so performance testing is more complicated than with traditional storage systems, and results even more dependent on workloads. Also, the most pertinent performance metrics are those that measure quantity of VMs or virtual desktops supported using real-world tools (such as IOmark.org) and not simply report IOPS or throughput statistics.

Data Resiliency

HPE SimpliVity uses two levels of internal data protection to maintain data availability in the event of a failure, protection at the node level and at the disk level. Node-level resiliency is achieved by replicating the data associated with each VM to another node in the cluster, creating a "replica set". In this way, the VM's data is available on a second node should the primary node fail.

These replica sets are essentially mirrors created at the VM level, with the OVC on primary and secondary nodes establishing a peering relationship. During the write process they send data to the OAC in each node which completes the write and provides an acknowledgement.

Each HPE SimpliVity appliance or server chassis also provides drive-level resiliency by setting up the disk drives in a RAID 5 or RAID 6 configuration, allowing up to two drive failures to occur without losing data. Metadata is stored on the same disk as actual data blocks using RAID 5/6, depending on the model. This local RAID protects against the most common cause type failure, at the storage device (SSD) level, and makes single-node operation feasible.

Replica sets are established at the VM level, not at a volume level. Each node can be the primary for one set of VMs and serve as a secondary for VMs on other different nodes as well, creating a replication factor of 2 (RF2), meaning the system can sustain one complete node failure and not lose data. When the local RAID protection is figured in, it produces a higher level of resiliency one that can sustain the loss of 3 drives (comparable to an RF4 protection).

Evaluator Group Comment:

In most scale-out architectures a single drive failure can trigger a rebuild process that severely impacts performance. This doesn't cause data loss but makes the system vulnerable to another failure, so most spread data blocks around the cluster to shorten this rebuild time. HPE SimpliVity's local RAID can sustain up to three drive failures and run the rebuild process on the internal PCIe card, eliminating the need to spread data out to multiple nodes. The net of this two-level process is increased data resiliency that's more efficient and a failover process that doesn't require data movement that can also impact performance.

Data Optimization

When data is first ingested the DVP looks at all blocks and skips writing those already resident in storage - this is done in the accelerator card or in software. While typically called "deduplication", the process is more like "single instancing" or "duplicate prevention", since it keeps the system from writing duplicate data blocks in the first place, rather than removing duplication after the fact. This technology is unique to HPE SimpliVity.

Running deduplication and compression results in less storage capacity consumed in the HPE SimpliVity cluster, as data reduction does in all storage-related systems. But running that reduction at ingest enables the DVP to minimize the "footprint" of each data object the system encounters and do that before those data have been handled by other processes down the line. This creates a data-reduced global storage pool that supports all functions in the HCI, eliminating the need for second-pass deduplication or "rehydration" of data further down the data path.

These data blocks, which may have multiple data objects associated to them (as an example, representing different files that share some of the same data) are stored by the DVP. But the DVP won't delete a block until it has determined that block has no data objects associated with it (sometimes called a "use factor of 0"). Deletion will only occur during a separate process, similar to the "garbage collection" phase in flash devices. This applies to metadata-based copies as well as files, essentially making every point-in-time copy immutable, a capability that drives SimpliVity's data protection features (see "Snapshots and Clones"). Another benefit of reducing the data footprint at ingest is to reduce WAN bandwidth consumption and speed up off-site data transfers.

Evaluator Group Comment:

Effective data reduction is an important characteristic of storage systems, since it reduces the amount of capacity needed to support a given workload. In an HCI system this means buying fewer nodes, not just fewer drives. While most HCIs offer some form of deduplication and compression, with its DVP technology, HPE SimpliVity guarantees a 10:1 data reduction ratio, which can lower the cost of an HCI deployment. (The company actually claims that 50:1 ratios are typical, which obviously depend on data type, among other things. But capacity is usually an important characteristic in these systems, so data reduction features of HCIs being considered should be understood.

Hypervisor Support

Hypervisors are an integral part of HCI operation, running the primary software defined storage program as a VM (for most HCIs), as well as providing the platfrom that supports the compute function of the system. Most HCIs also embed hypervisor management functionality (typically VMware's vCenter) into their product's control plane. HPE SimpliVity currently supports VMware vSphere and Microsoft Hyper-V hypervisors.

Management Functions

"Federation" is a term HPE SimpliVity uses to describe a networked collection of multiple HPE SimpliVity clusters that can be configured in either a hub-and-spoke or full-mesh topology. This allows for simplified management of multiple clusters from a single site (see "Global Unified Management").

Global Federated Data Space

Within the Federation, data can move between nodes in separate sites while in a compressed and deduplicated state. To do this, the DVP tracks metadata through a system of inter-node communications such that each node knows which data blocks are on every other node. This enables nodes to physically transfer only the blocks required to maintain replica sets, support disaster recovery, etc., minimizing the traffic between nodes. Since data gets deduplicated at ingest, the only duplicates that exist in the environment are those created by replica sets.

In addition to maintaining data efficiency, Global Federation enables cloud integration by managing data mobility between HPE SimpliVity clusters and the cloud. It also facilitates cache accelerated performance and data sharing for high availability, while providing an automated, single point of global management that's organized around the VM. These Global Unified Management features are implemented within existing infrastructure management frameworks, such as VMware vCenter and vRealize Automation.

Composable Fabric

HPE's Composable Fabric is software-defined networking (SDN) that's integrated into the HPE SimpliVity software via APIs. It provides automated network management and simplified configuration and expansion through a vSphere plugin. It allows IT to customize network performance through path isolation, helping to manage the internode "east-west" traffic that scale-out system generate. HPE bundles the Composable Fabric management software and Top of Rack switches with the HPE SimpliVity cluster to provide a comprehensive networking solution.

VM-Centric Management

Storage needs to be organized in such a way that it supports the environment users will be working in. Originally, shared storage systems were block-based or file-based, but now in the virtual server world, VMs are the most pertinent way to organize and manage data.

HPE SimpliVity is designed aroud VMs with all storage services performed at the VM-level, including data protection, performance monitornig and management, troubleshooting and administrative tasks. After setting policies for each VM, users can skip all of the underlying tasks traditionally associated with storage management. All the details around LUNs, IP addresses, networking ports, etc., that come up when VMs are created, files are migrated, or even when nodes are added or removed, are abstracted from the user and handled by DVP.

Infosight

Infosight is a cloud-based "predictive analytics" platform that pulls systems data from the SimpliVity installed base to improve reliability and optimize resource utiliztion. Originally developed by Nimble

Storage, before being acquired by HPE, Infosight was one of the first products to analyze product telemtry data points and use AI/ML techniques to derive management insights. These data, designed to help IT anticipate maintenance events and predict resource requirements, are shared with users to improve their operational experience.

Native Hypervisor Management Applications

HPE SimpliVity doesn't provide a native management GUI but instead, provides plug-ins for native hypervisor management tools, such as VMware vCenter. This enables users to control backups and restores, manage resources, troubleshoot problems, etc., using the vSphere client they're familiar. Through vCenter they can also add, move and delete nodes from HPE SimpliVity clusters and configure certain node settings. The vSphere client provides configuration information, status displays, plus resource inventories about specific HPE SimpliVity nodes, even providing front and back panel graphics to confirm some of these data (see Figure 5).

HPE SimpliVity also integrates with automation and orchestration tools such as Cisco's UCS platform and VMware vRealize. These programs allow users to manage their own infrastructures with self-service portals and service catalogs, under IT administrative control.

An HPE SimpliVity system can generate and forward information about the health and status of the infrastructure, such as an alarm from a disk failure or notification that a backup completed successfully. These data can be displayed on the integrated hypervisor management system or forwarded to a third party orchestration tool. Critical events can even trigger a "phone home" response, which can also be forwarded to the HPE SimpliVity support organization.

Evaluator Group Comment

Increasing IT agility and reducing infrastructure management overhead is a growing need in enterprises, as well as smaller organizations. Self-service operation can help achieve this and HCIs are a popular way to get started, since they're relatively easy to set up and run.

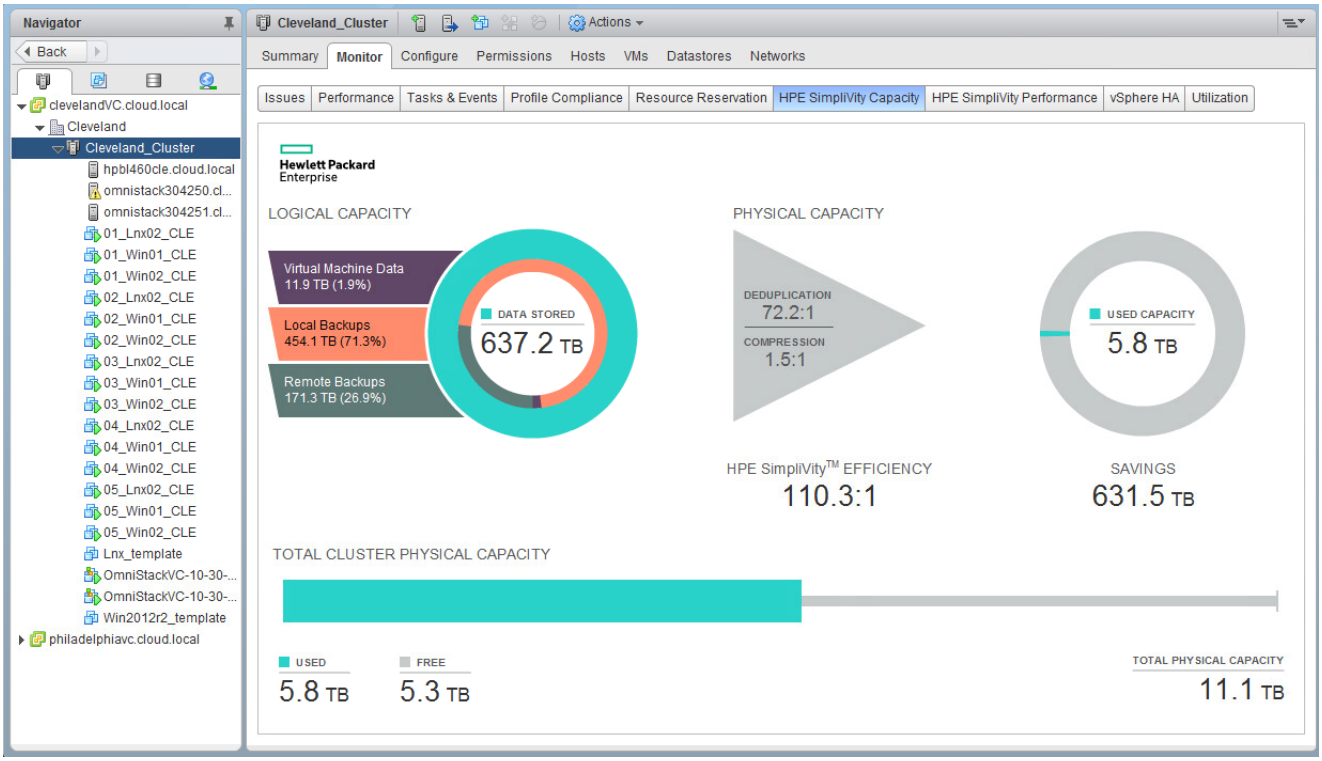

Figure 5: vSphere Client Screen

The vCenter GUI plug-in provides resource utilization data, such as used and available storage capacity, plus VM performance data (throughput, IOPS and latency). In addition, HPE SimpliVity has a REST API that's designed to work in conjunction with the hypervisor plug-ins described in the previous section. This API supports a range of operations including configuration, backup, recovery and replication, as well as routine maintenance and diagnostics. HPE SimpliVity also has programmatic interfaces to VM management functions utilizing a command line interface (CLI).

Data Services

Flash, Caching and Tiering

The OmniStack Accelerator Card handles the data I/O activity in each node, using write caching with DRAM, plus NVRAM and flash storage to maximize performance. As an all-flash solution, HPE SimpliVity doesn't support any storage tiering functions.

HA and Failover

High availability (HA) within each node is maintained through the use of redundant components common on industry-standard server hardware (power supplies, NICs, fans, etc.) and by RAID protecting storage media (See Data Resiliency). HA between nodes is achieved by the creation of replica sets, mirrored copies of VM data stores that are maintained on a separate node, and automated failover software. When a node fails (including a RAID controller or any subsystem component) the secondary

node that holds the replica set becomes the primary, enabling affected VMs to continue operation, uninterrupted. It's important to note that this process occurs without any rebuilding of data sets or reassembling of data from other nodes that is required with many HCIs.

Snapshots and Clones

Snapshots create a metadata tag or "pointer" to the data blocks associated with a particular file. Since copying these pointers can be done in the background, without actually touching the data itself, pointers allow a point-in-time copy (or many copies) of that file to be captured, with little or no impact to the applications using the original data object. But in order for snapshots to be used as backups, the data blocks they represent must be made immutable. In most storage systems this involves creating a "clone", either copying these blocks to another storage area or simply marking them as read-only, so they can't be overwritten. This eventually consumes more storage or maintains extra copies of file blocks that must be managed and eventually deleted.

DVP simplifies this entire process by essentially marking every data block written as immutable and deleting them only when no other file or copy is associated with them. This way, every metadata-based copy is basically a clone of the original data object and is safe to use for backups. DVP runs what amounts to regular deletion cycles, like flash garbage collection, where orphaned data blocks are cleaned out.

These blocks are subject to the same internal data protection, via replica sets discussed earlier, so additional copies are not needed to create a protected backup, unless a copy is made for off-site replication to satisfy DR requirements (see "Built-in Data Protection").

Evaluator Group Comment

While snapshots and clones in general are fairly standard features of every HCI, this method is unique to HPE SimpliVity. Making data blocks immutable as the default state eliminates the cloning step, simplifies the use of point in time copies for data protection and enables their built-in backup feature.

Replication

HPE SimpliVity uses asynchronous replication to satisfy requirements to get backups off-site. The OVC at the primary and remote site figure out which blocks are lacking and only transmit those, reducing bandwidth and storage consumption. And, as detailed above (see Data Optimization), the DVP maintains all data blocks in their deduplicated and compressed state throughout the replication process. HPE SimpliVity also uses synchronous replication to maintain the mirrors or "replica sets" that are established based on policies for each VM. (see Data Resiliency). These replica sets can be local or in two remote locations, connected by a WAN, setting up a stretch cluster.

Stretch Clusters

An HPE SimpliVity cluster can be split between two sites that are physically separated, as long as an adequate bandwidth connection is available to support the synchronous replication between nodes, typically a 10Gbps connection with <1ms latency. This allows the use of VMware HA functionality for automated VM recovery after a site failure. In a stretched cluster configuration the Arbiter should be located at a third site.

Advanced Features and Functions

Built-in Data Protection

HCI appliances are well-suited for data protection, since they typically include the snapshots, clones and replication services that underlie most backup and disaster recovery features. HPE SimpliVity leverages its unique snapshot-like process described above to create "full logical backups", instead of creating the physical copies or clones that are required with traditional snapshots. With this technology, full backup and disaster recovery protection can be implemented without impacting performance. And, HPE SimpliVity HCIs can be fully protected without requiring third-party backup applications, dedicated local storage or special DR infrastructure.

Backup policies are set at the VM level, specifying backup frequency, retention, the backup window and the destination cluster. For example, more critical VMs can have policies to run backups twice a day and have them replicated off-site every night, where other VMs may be backed up once a day and sent off-site weekly. Backups can be local, kept on the same cluster, or replicated to a remote cluster for disaster recovery protection. Since the remote cluster is part of the Federation there's no complex set-up involved, like defining backup targets or configuring shares or mount points.

Backups can also be made application-consistent, using VMware Tools and Microsoft VSS functions, to synchronize the file system and applications like MS SQL, Exchange or Active Directory. Restores, like backups, can be initiated from the CLI or REST API, from the integrated hypervisor management applications or orchestration tools. Users or admins can also manually restore a full VM or specific files or folders from an existing backup.

HPE Simplivity HCI also supports backups to HPE Cloud Volumes and HPE StoreOnce appliances. These targets were added mainly to provide centralized backup for remote sites.

Evaluator Group Comment:

HCIs simplify much of the operation of a virtual compute environment but most don't provide a turnkey data protection solution. For this reason many companies still use a dedicated backup

infrastructure. Incorporating comprehensive data protection into the base product, the way HPE SimpliVity does, is a unique feature.

Intelligent Workload Optimizer

Data locality, the ability to assure that data is stored on the physical node that supports the VMs using that data, is an important characteristic for maintaining predictable performance in scale-out architectures. HPE SimpliVity's use of replica sets described above (see "Data Resiliency") maintains data locality on separate nodes for the primary and secondary copies of each VM.

OmniStack has a feature called "Intelligent Workload Optimizer" (IWO) that works with vSphere DRS, providing awareness of the location of VM data sets for each VM and specifying that each VM stay on the same node as its data. Maintaining these data in replica sets and providing this information to DRS assures that VMs and their supporting data sets remain on the same nodes. IWO assures data locality by moving VMs to follow their data instead of moving virtual machine data to follow the VMs.

Global Unified Management

Most hyperconverged solutions include some kind of remote management capability, a feature that allows multiple nodes and clusters in multiple locations to be controlled from a "single pane of glass". HPE SimpliVity's Global Unified Management takes this a few steps further, utilizing its VM-level abstraction to simplify the management tasks that need to be done.

In the Global Federated Data Space that was described above, each node is continually updating its metadata store to maintain a record of the data blocks that are on every other node. But in addition to being self-updating, HPE SimpliVity nodes are "self-learning", maintaining configuration information from other nodes as well as metadata about their contents as the environment changes and evolves.

This enables each node to manage the changes between nodes without requiring users to do any of this more complex work themselves. This means no updating of replication parameters or network information when a VM and its data set are migrated, or a node is added to or removed from a cluster, etc. The result is a simpler management experience and fewer chances for mistakes or omissions to occur in a multi-site environment.

HPE GreenLake

HPE offers its GreenLake Infrastructure-as-a-Service on HPE SimpliVity HCI clusters. With GreenLake, IT departments can turn the capital expense of procuring and deploying IT assets into operational expense that's based on consumption. As with GreenLake on other HPE products, HPE PointNext provides professional services and customizes the payment structure to the company's needs. GreenLake with HPE SimpliVity uses HPE Infosight to optimize capacity and predict future requirements.

Evaluator Group Comments

While Nutanix is credited with largely establishing the hyperconverged infrastructure product space, SimpliVity's early market success was also a driver for this technology's impact on IT before it was acquired by HPE in 2017. But SimpliVity has not taken off as part of HPE and in early 2019, HPE partnered with Nutanix, offering the Acropolis OS on HPE ProLiant DX and Apollo servers. With this move, HPE is emphasizing SimpliVity as an SMB and mid-market solution, preferring to sell the Nutanix solution in the enterprise, where HPE has a significant installed base. Nutanix and SimpliVity are also sold as services through HPE GreenLake.

The on-board PCIe data reduction and acceleration card provided some unique functionality, and created some primary technology differentiation, but did add to the base hardware cost and didn't fit in the 4-node 2600 chassis or the 1U DL325 chassis. While this card is still available on the DL380, HPE's direction is to do this data reduction in software.

The ability to reduce the amount of storage needed (through data reduction at ingest) for a given deployment can mean buying fewer nodes. This can have a bigger impact in HCI (compared with the effect deduplication has on a traditional storage system), as each node added for capacity adds more compute and memory resources – along with the associated hypervisor licenses.

The built-in backup feature is another important differentiator that leverages the HPE SimpliVity architecture and one that has been used to capture market share, since data protection seems to have universal appeal and few of the other HCI vendors have a comparable solution. That said, HPE needs to develop a more comprehensive cloud strategy and provide features that support native container environments beyond the addition of a CSI plugin.

The SimpliVity 2600 provides a higher-density option, while maintaining support for a single GPU. The 1U DL325 chassis with AMD CPU is an interesting option for smaller companies and environments with space constraints. HPE's GreenLake Infrastructure-as-a-Service offering provides additional options as well.

More detailed information is available at http://evaluatorgroup.com