Technology Overview:

# SimpliVity Hyperconverged Infrastructure for VMware vSphere

THE INFORMATION PRESENTED IN THIS PAPER IS SUBJECT TO CHANGE WITHOUT NOTICE.

SimpliVity, the SimpliVity logo, OmniCube, OmniStack, and Data Virtualization Platform are trademarks or registered trademarks of SimpliVity Corporation in the United States and certain other countries. All other trademarks are the properties of their respective owners.

## Table of Contents

# 1. Introduction

## 1.1. Purpose

This Technology Overview reviews the features and benefits of 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 SimpliVity solution into physical and logical constructs and describes the various internal and external interfaces and data flows. Particular attention is given to the inherent data efficiencies of the 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 SimpliVity configurations and deployment models, and provides sample use cases for SimpliVity hyperconverged infrastructure.

## 1.2. Audience

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

# 2. Introduction to SimpliVity Hyperconverged Infrastructure

## 2.1. 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 siloed 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 a number of 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.

## 2.2. 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 in order 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:

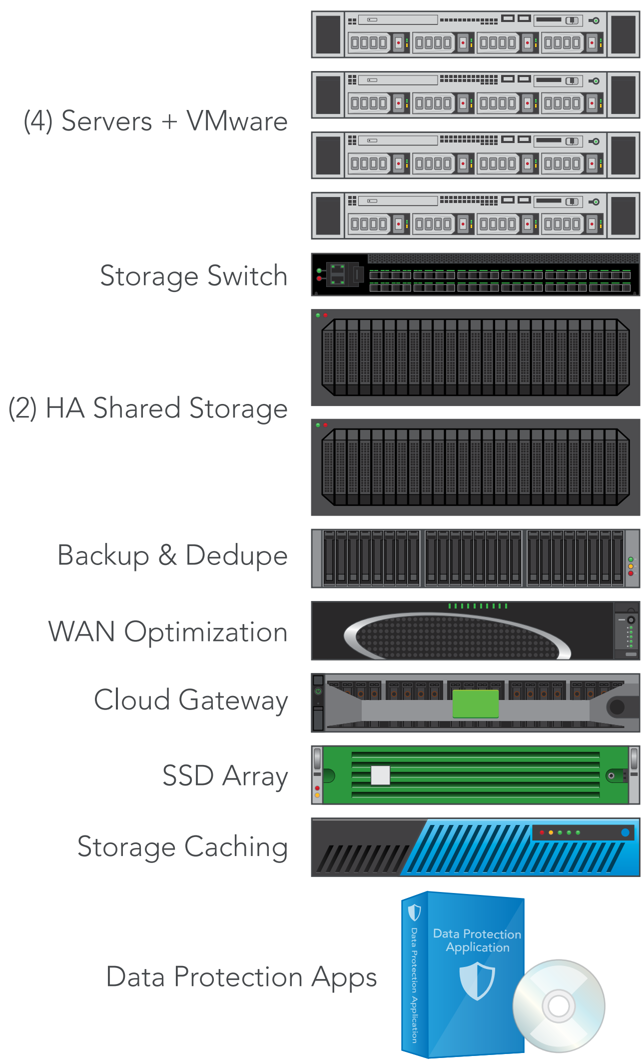


Figure 1: Siloed IT architectures are inherently inefficient, cumbersome, and costly.

## 2.3. 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.

## 2.4. The Solution: SimpliVity Hyperconverged Infrastructure

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 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; 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 SimpliVity 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: accelerated data efficiency, global unified management, and built-in data protection.

### 2.4.1. Accelerated Data Efficiency

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, Flash/SSD, and HDD), all handled with fine data granularity of just 4KB-8KB. 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.

### 2.4.2. Global Unified Management

Simplifying the complexities and administrative overhead of managing multiple silos of IT infrastructure distributed across multiple data centers was one of the intentions of the SimpliVity solution from the very beginning. Rather than managing individual silos at the device level, a globally-distributed SimpliVity implementation is managed as a unified system with

a single pool of resources across multiple sites. The entire system, which SimpliVity refers to as a 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. With SimpliVity's Global Unified Management features, there is no longer any need to maintain IP configurations in any site globally when adding or removing nodes.

Another way 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 a 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 (e.g., VMware vCenter, VMware vRealize Automation, Cisco UCS Director), further reducing the number of management interfaces and IT staff training needs.

### 2.4.3. Built-in Data Protection

Another core intention of the SimpliVity platform is data protection. Maintaining availability of data is the most critical function of any platform that stores data. Within each node, data is protected from the loss of at least two HDDs. All data is also synchronously persisted on two local SimpliVity nodes to protect against the loss of any single node. These two levels of protection work together so that the loss of two HDDs in one node, while a second node is offline for maintenance, will not cause data to become unavailable.

The SimpliVity solution includes built-in VM-centric data protection functionality, eliminating the need for purpose-built backup and recovery products. SimpliVity's ultra-efficient approach to 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 SimpliVity Data Center.

## 3. SimpliVity Technology Overview

### 3.1. Terminology

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

- **OmniStack node** – an x86 server that is the basic hardware building block of the SimpliVity hyperconverged infrastructure solution.
- **OmniStack Virtual Controller (OVC)** – the software stack, implemented as a single VM per node, which controls all aspects of SimpliVity hyperconverged infrastructure.
- **OmniStack Accelerator Card (OAC)** – a SimpliVity-designed PCIe-based device that offloads and provides acceleration of writes and data management functions within the SimpliVity hyperconverged infrastructure solution.

- **OmniStack Data Virtualization Platform (DVP)** – a globally-aware file system 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.
- **SimpliVity Data Center** – a collection of one or more OmniStack nodes typically located at the same physical site connected over a standard Ethernet network collectively providing a single storage pool to the hypervisor. SimpliVity Data Centers 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.
- **Federation** – a collection of one or more SimpliVity Data Centers and the main construct within which data is managed.
- **Intelligent Workload Optimizer** – a multi-dimensional approach to balancing workloads across OmniStack 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 identify performance issues, optimize resource utilization, and effectively plan and manage future growth within their own SimpliVity infrastructure.
- **OmniWatch** – a proactive support service for all SimpliVity support plans, that complements our support portal for comprehensive care. OmniWatch continuously monitors your SimpliVity infrastructure, checking over 100 indicators to continuously evaluate system health and automates support case creation based on thresholds.

**Note:** When capitalized, "Data Center" refers to a SimpliVity architectural construct. When spelled in lower case, "data center" refers to an enterprise or service provider facility housing IT equipment.

## 3.2. SimpliVity Scale-Out Architecture

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

### 3.2.1. SimpliVity OmniStack Node

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

### A compact hardware platform

The SimpliVity 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.


Figure 2: SimpliVity OmniCube

Every node includes performance-optimized solid state drives (SSDs) and capacity-optimized hard disk drives (HDDs) protected in hardware RAID configurations. The HDDs are utilized for permanent virtual machine data and hashtable index storage, and are configured to support at least two drive losses per node (all models except ROBO models). The SSDs are utilized for permanent metadata, 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. When hardware-based data at rest encryption is required, Self-Encrypting Disks (SEDs) are supported on specific SimpliVity platforms as identified in the most recent support matrix.

Running on the bare-metal of every node is a virtual machine hypervisor. SimpliVity was designed to be hypervisor agnostic in order 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 SimpliVity's implementation with VMware vSphere.

SimpliVity offers the OmniStack node in the form of a turnkey appliance utilizing an OEMed server under the name **OmniCube**. The solution is also integrated with a number of market-leading x86 servers, including Cisco UCS and Lenovo System x servers, under the name **OmniStack Integrated Solutions**. All of the deployment options run the same OmniStack components.

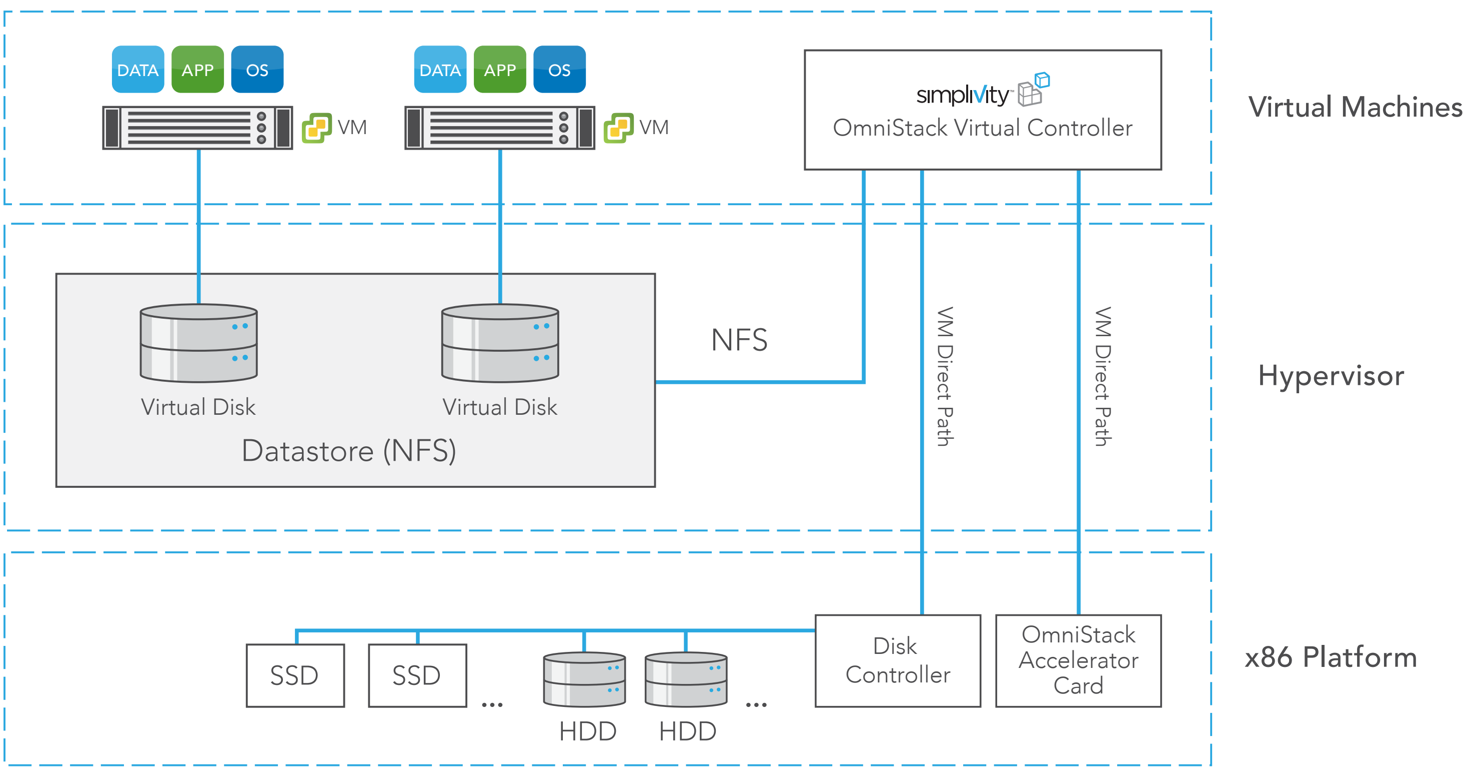


Figure 3: The Components of an OmniStack Node

## OmniStack Virtual Controller

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

## OmniStack Accelerator Card

The OmniStack Accelerator Card is a purpose-built PCIe based Accelerator card designed by SimpliVity. The OmniStack Accelerator Card consists of:

1. FPGA (Field Programmable Gate Array)

2. NVRAM

3. Super capacitors to provide power to the NVRAM upon a power loss

4. Flash storage used to persist NVRAM after power loss

The Accelerator offloads CPU-intensive functions from the x86 processors, performing in-line deduplication, in-line compression, and in-line write optimization. By offloading these functions, SimpliVity is able to provide these core features with improved performance and very minimal impact to the resources available to business application VMs.

## 3.2.2. SimpliVity Data Center

The SimpliVity Data Center is the primary logical construct that defines a single grouping of OmniStack nodes. The 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.

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

### vCenter Server

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

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

### SimpliVity Arbiter

A "split brain" occurs when two independent systems configured in a cluster lose network connectivity and assume they have exclusive access to resources. SimpliVity uses an Arbiter to ensure that a two-node or larger cluster survives an OmniStack node failure without service disruption or loss of access to data.

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

Under normal operations, OmniStack nodes and the SimpliVity Arbiter are in constant communications through heartbeats that are sent across the environment. However, it is important to note that the 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.

In order for the arbiter to properly provide this independent witness, it must exist outside of the SimpliVity Data Center in which it is a participant, and therefore cannot exist on a SimpliVity node or datastore. The 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 OmniStack nodes.

### SimpliVity Datastores

When configuring a SimpliVity Data Center, one or more datastores must be created to present the storage to the ESXi hypervisor. Each datastore is presented as an NFS datastores with a common namespace across all the OmniStack Virtual Controllers in a SimpliVity Data Center. These datastores are logical constructs within the Data Virtualization Platform Data Presentation layer, and are simply utilized as a communication channel between the hypervisor and the 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 a node. Similarly, replication is not limited by datastores since it is definied at the virtual machine level (via policy).

### Stretched Clusters

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

A stretched 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. In order to ensure this functionality and avoid split brain scenarios, SimpliVity recommends the Arbiter be deployed at a third site to avoid the loss of the Arbiter at the same time as one-half of the SimpliVity Data Center.

### 3.2.3. SimpliVity Federation

A SimpliVity Federation is a networked collection of one or more SimpliVity Data Centers. The 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 SimpliVity to efficiently move only unique data between SimpliVity Data Centers when performing Remote Backup and VM Move operations.

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

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

While not mandatory, SimpliVity recommends implementing multiple vCenter Servers for multiple SimpliVity Data Center configurations to ensure each site can maintain full vSphere management capabilities even if the other site is unavailable.

### 3.2.4. Supported Federation Topologies

The SimpliVity Federation can be configured in either a full mesh or 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, backup, and restore from any SimpliVity Data Center to any other SimpliVity Data Center, without restrictions. The hub-and-spoke topology is intended for geographically distributed deployments such as remote office/branch office implementations. VMs can migrate, backup, and restore only between a spoke site (i.e., a branch office) and a hub site (i.e., a central data center or corporate headquarters). The hub-and-spoke topology minimizes communications traffic and is suited for SimpliVity Data Centers connected over a WAN.

When deploying additional SimpliVity Data Centers or remote offices, 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 mis-configurations that come from individually configuring the intercommunication between each data center/remote office.

## 3.3. Data Virtualization Platform

At the heart of the SimpliVity solution lies the Data Virtualization Platform—a fabric that extends across multiple OmniStack nodes in multiple globally-distributed SimpliVity Data Centers that abstracts the VM data away from the under-lying hardware. In a single Data Center, the Data Virtualization Platform abstracts the underlying hardware across mul-tiple OmniStack nodes and presents this unified storage to the hypervisor on each node as a single pool of storage. The 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 residing on any SimpliVity system. Across Data Centers, the Data Virtualization Platform also allows for a single point of management, including unified policy definitions for data protection and efficient movement of data between physical locations.

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

### 3.3.1. Presentation Layer

The Presentation Layer is responsible for the creation and presentation of datastores to the hypervisor, and maps datas-tore objects to the underlying Data Management Layer. These datastores are presented to vSphere via NFS using a unified namespace across all OmniStack Virtual Controllers. This is the layer that requires the majority of customization when sup-porting other hypervisors (for example, SMB will be utilized for Hyper-V hosts). Each OmniStack Virtual Controller main-tains it's own IP and piece of the datastores, but coordinates with the others to present a single datastore to the hypervi-sors in the Data Center.

### 3.3.2. Data Management Layer

The Data Management Layer is responsible for the tracking and storage of all data and metadata. The Data Management Layer is further sub-divided into two other layers: the file system and object store. They work together to track and man-age all metadata and blocks of physical data. All of the metadata is stored and managed as individual objects, as are the relationships of all the objects. By tracking the relationship of all the metadata object, 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 higher level objects representing the relationship of data need to be copied. This can be a significantly quicker operation for larger virtual machines than having to copy the underlying metadata that describe the actual blocks.

Understanding the relationship between the different objects is critical to the Data Virtualization Platform's global data efficiency as well. When moving data between sites (via Remote Backup or VM Move operations), a comparison of existing objects occurs between the sending and receiving nodes. By comparing the indexes of objects that represent the relation-ships between other objects, replication of both data and metadata can be as efficient as possible since only truly unique objects are transferred.

Each OmniStack Virtual Controller maintains it's own record of metadata independent of all other nodes. All of this data and metadata is deduplicated within the node regardless of datastore location or content type (i.e. ISO file, VMDK, VHD, etc.).

Figure 4 depicts the DVP in a single SimpliVity Data Center composed of three OmniStack nodes.

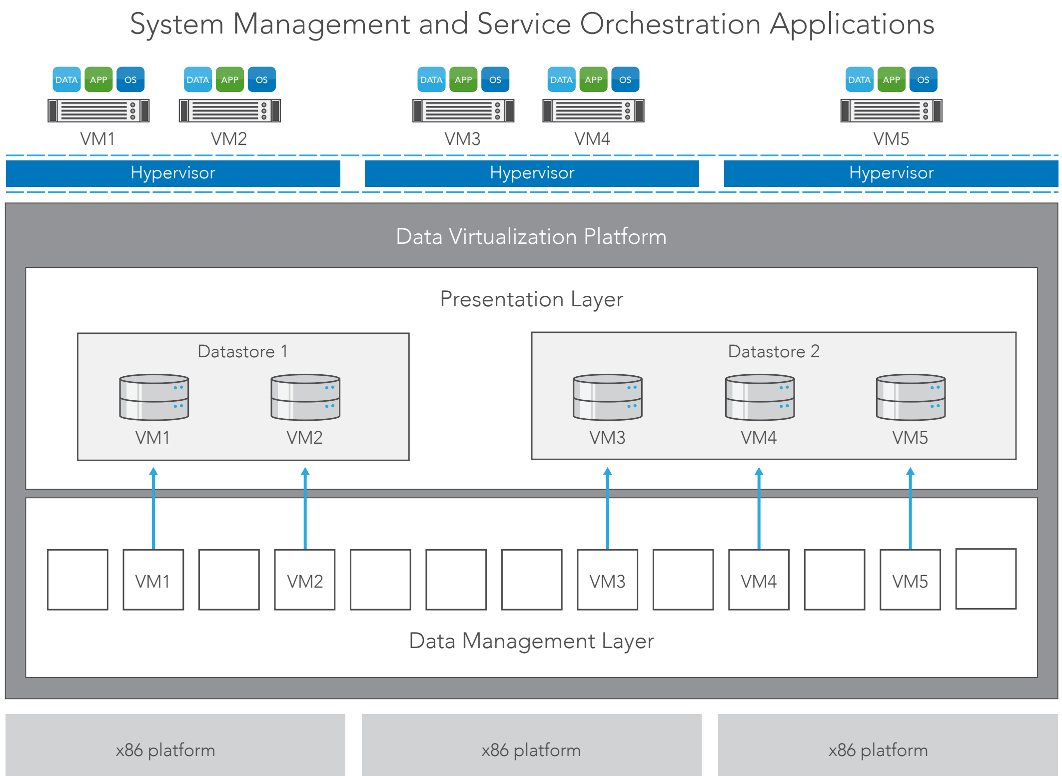


Figure 4: Data Virtualization Platform spanning three OmniStack nodes

Figure 5shows the DVP in a Federation of three SimpliVity Data Centers.

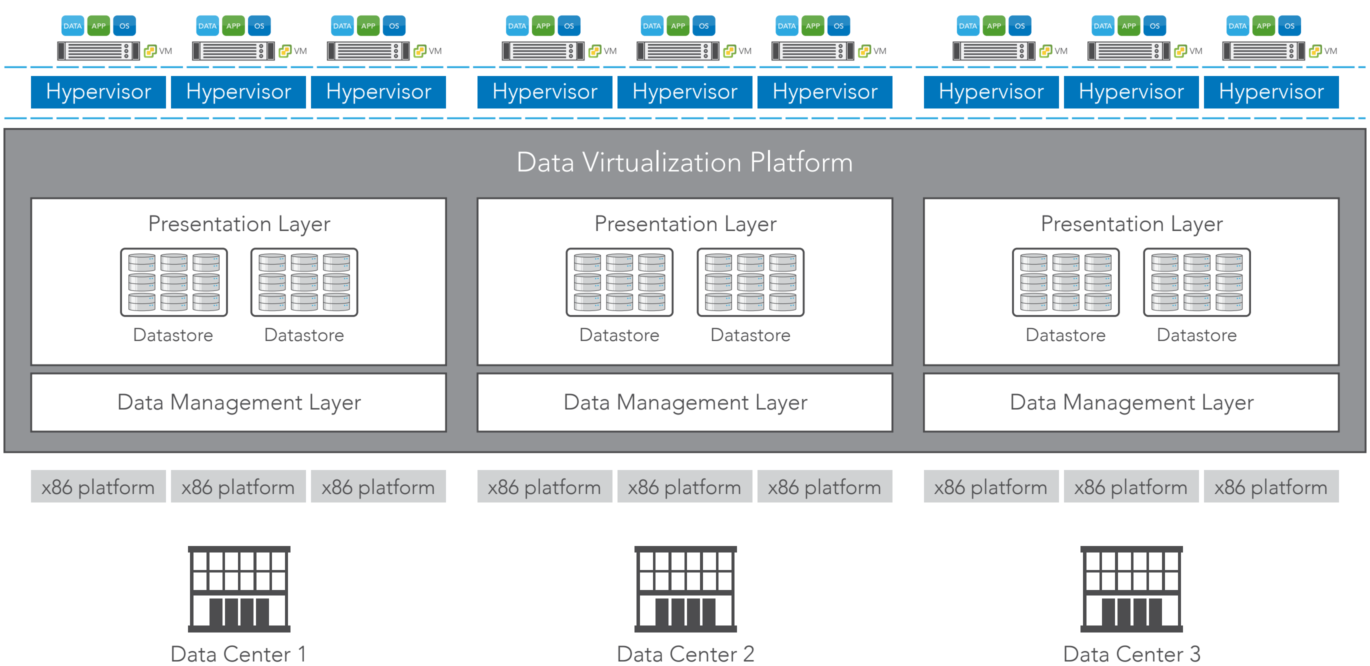


Figure 5: Data Virtualization Platform spanning three Data Centers

### 3.3.3. VM-Centric Design

Most of today's data center equipment was originally designed around non-virtualized 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 (i.e., 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, not physical devices.

SimpliVity's Data Virtualization Platform was designed from the ground up with virtualization in mind. All SimpliVity 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 also 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 (e.g., VMware vSphere client).

### 3.3.4. Inline Deduplication, Compression, and Optimization

The Data Virtualization Platform Data Management layer performs inline data deduplication, compression and optimization on all data, at inception, across all tiers (DRAM, Flash/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 4KB-8KB granularity by the OmniStack Accelerator Card inline, before it is written to hard disk. SimpliVity's deduplication and compression does not require any intervention by the administrator since it is always on and accelerated by the OmniStack Accelerator Card, ensuring 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 OmniStack Accelerator Card, with writes first committed to NVRAM to ensure inline deduplication actually improves performance instead of impeding it.

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

## 3.4. Resiliency and High Availability

The 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 SimpliVity-enabled DR functionality.

### 3.4.1. Node Availability

SimpliVity utilizes standard x86 servers for the compute platform, gaining all the availability features, 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. RAID6 is used to protect the VM data and index tables on the HDDs, and either RAID1 or RAID5 - depending on the number of SSDs - is used to protect the metadata on the SSDs. This ensures disk-level resiliency, so that the loss of disks will not impact the availability of data on an individual node.

Node-level resiliency is utilized for the non-redundant components within a SimpliVity node (e.g. OmniStack Accelerator Card, CPU, Disk Controller, motherboard, etc.). More on that in the next section.

### 3.4.2. 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 OmniStack nodes within a SimpliVity Data Center (these two instances are referred to as a "replica set"). Should one OmniStack node fail, the VM's data is still available on another node so vSphere HA can restart all the affected virtual machines.

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

### Lifecycle of a Write I/O

The replica sets associated with a particular VM is managed on two separate nodes within a Data Center. The OmniStack Virtual Controllers establish a peering relationship on a per VM basis. The OmniStack Virtual Controller on the same node as the VM will share every write operation with its peer for that VM, and each OmniStack Virtual Controller will calculate and execute deduplication, compression, and optimization independently. With this approach, the persistence of every block of data is guaranteed across two nodes without the overhead of all nodes tracking the details of how or where each block is stored within the individual nodes. 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.

Figure 6 details the write path in a high availability implementation.

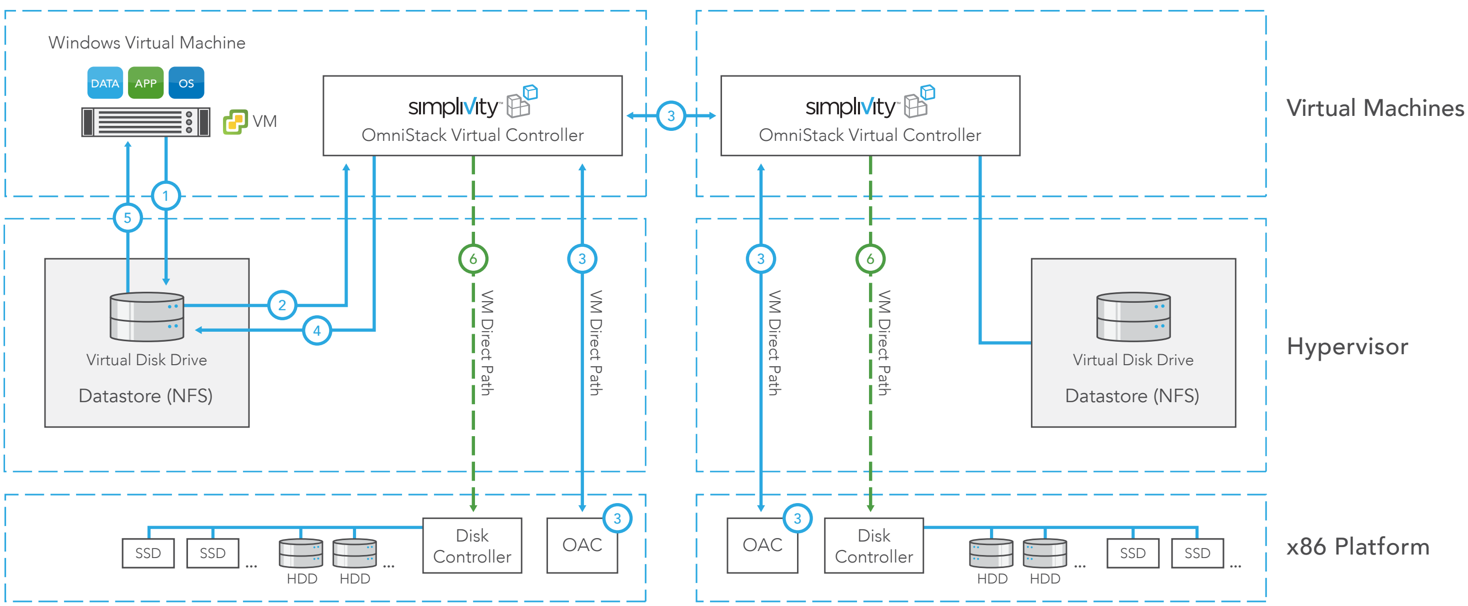


Figure 6: High Availability Write Path

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 OVC.
3. The OVC replicates the I/O to the OVC on the VM's secondary node. In parallel, each OVC places the data in NVRAM on its OmniStack Accelerator Card. The data is now protected by flash and super-capacitors on both nodes. Once staged in NVRAM, each OAC independently processes the data for deduplication and compression, then acknowledges back to the OVC.
4. Immediately after step 3, the primary OVC 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, the OVC 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 OmniStack 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 VMware HA event. Storage access for the affected virtual machines will failover to the other OmniStack Virtual Controller in the replica set.

Each OmniStack Virtual Controller has an IP address that the hypervisor uses to access the SimpliVity NFS service. Should the OmniStack Virtual Controller, disk controller, or the OmniStack Accelerator Card fail, the IP address and NFS mount of the OVC automatically fail over to a surviving node's OmniStack Virtual Controller. Since ESXi uses this IP address to access the datastore, all VMs and applications continue to run without interruption, because no VMware HA event occurs or is necessary. This prevents these components from being a single point of failure for virtual machine or data availability.

To protect from the loss of an entire 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 due to the fact that all writes are written to both sites synchronously, as described above in the Lifecycle of a Write I/O section above. In this case, the loss of one site would be detected and VMs restarted by vSphere HA, providing for a minimal amount of downtime.

### 3.4.3. Federation Availability

To protect data across a SimpliVity Federation, remote backups can provide Recovery Point Objectives 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 Recovery Time Objectives. This recovery operation can be automated through SimpliVity's REST API by using scripts or an orchestration tool to further enhance the Recovery Time Objective.

### 3.4.4. 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 (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 1GB of patches on 20 Windows servers. Traditional backup products would backup that 1GB 20 times, including all the associated reads, writes, and network bandwidth, because that 1GB 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 SimpliVity OmniStack, data management operations like creating backups, clones, restores, or performing move operations can be completed locally or globally much more efficiently than with traditional data management solutions. The Data Virtualization Platform maintains full logical backups without impairing performance or squandering storage capacity. Each backup is a complete, standalone 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 standard storage snapshots. Individual blocks are not owned by any specific entity, so recovering a VM that's been deleted or recovering a VM at a different site is no problem.

By tracking individual blocks, the 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 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 SimpliVity Data Centers, the nodes in the receiving SimpliVity Data Center investigate the blocks needed to recreate the backup or VM, compare them to the unique blocks the nodes already have stored, and request only the necessary unique blocks be transmitted. This approach to data movement maintains data efficiency across the Federation and also reduces the impact to the WAN when compared to traditional data replication methodologies.

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 or to replace the existing VM. 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 the original VM. 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 hard disk drive I/O is utilized and the restores can be completed very fast.

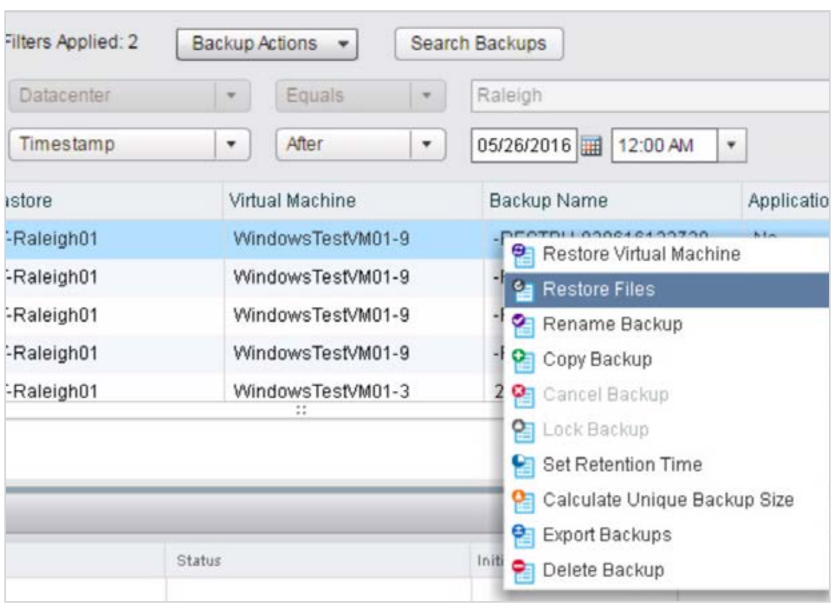


Figure 7: Simple restoration of an entire VM or specific files/folders.

The 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 windows to non-Production hours, thus accelerating backup and recovery functions (improving RPOs and RTOs).

## 4. SimpliVity Management Capabilities and Interfaces

The SimpliVity hyperconverged infrastructure solution includes a complete set of management capabilities and interfaces. From a management perspective, a SimpliVity Data Center or Federation is treated as a cohesive system with a single interface. The entire system is administered in a unified manner from a centralized management application.

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 SimpliVity solution also supports a native command line interface (CLI) and a programmatic interface for securely performing management functions via a terminal or scripts.

### 4.1. Management Functions

#### 4.1.1. Data Protection and Recovery

The SimpliVity solution supports a full set of data protection and recovery capabilities. They can be initiated via the SipliVity 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 a SimpliVity Data Center or remotely (to another SimpliVity Data Center). 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 SimpliVity Data Center 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 for 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 below).

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 for 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

The SimpliVity solution optionally supports application-consistent backup for VMware environments to adequately protect transactional applications like Microsoft SQL, Active Directory, or Exchange. The solution leverages 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 leveraged to run scripts designed to prepare the application for backup.

### 4.1.2. Troubleshooting and diagnostics

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

## Alarms and Events

The SimpliVity solution can generate a wide variety of alarms and events for reporting the health and status of SimpliVity hyperconverged infrastructure components and processes (i.e., a backup operation). Alarms are automatically forwarded to and displayed on native hypervisor management system and third-party service orchestration tools via SimpliVity plugins. For example, the 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. SimpliVity events, such as a SimpliVity rapid clone, manual backup, or creation of a backup policy are tracked through vCenter's 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 SimpliVity systems, and to allow system administrators to be alerted to vSphere and SimpliVity events through a common interface.

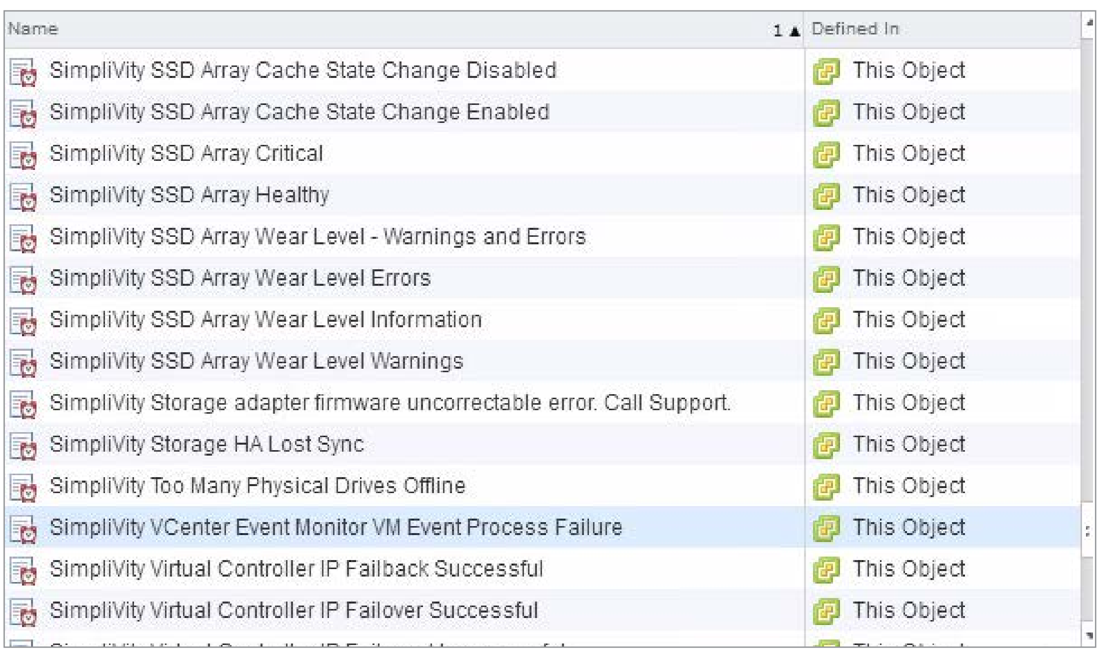

| Name 1 ▲ | Defined In |
|---|---|
| SimpliVity SSD Array Cache State Change Disabled | This Object |
| SimpliVity SSD Array Cache State Change Enabled | This Object |
| SimpliVity SSD Array Critical | This Object |
| SimpliVity SSD Array Healthy | This Object |
| SimpliVity SSD Array Wear Level - Warnings and Errors | This Object |
| SimpliVity SSD Array Wear Level Errors | This Object |
| SimpliVity SSD Array Wear Level Information | This Object |
| SimpliVity SSD Array Wear Level Warnings | This Object |
| SimpliVity Storage adapter firmware uncorrectable error. Call Support. | This Object |
| SimpliVity Storage HA Lost Sync | This Object |
| SimpliVity Too Many Physical Drives Offline | This Object |
| SimpliVity VCenter Event Monitor VM Event Process Failure | This Object |
| SimpliVity Virtual Controller IP Failback Successful | This Object |
| SimpliVity Virtual Controller IP Failover Successful | This Object |

Figure 8:SimpliVity Alarms within the vSphere Web Client

### Diagnostic Files

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

- Status of hardware components (disks, controllers, power supplies, etc.).
- Logs from the OmniStack Virtual Controllers that comprise the SimpliVity Data Center.
- Logs from the ESXi hosts that comprise the SimpliVity Data Center.
- Logs from vCenter relevant to the SimpliVity Data Center.

### OmniWatch

OmniWatch, SimpliVity's "phone home" capability, forwards critical alarms and event messages to SimpliVity's customer support. 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, SimpliVity customer support cases may be automatically created, depending on the type of alert and criticality.

Alarms, events, and messages can be automatically forwarded to SimpliVity's customer support organization. Health and status messages provide information about SimpliVity systems: the triggering alarm and alarm description, device name, serial number, current statuses of OmniStack Accelerator Card/datastores, etc.

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

## 4.2. Deployment and Upgrades

SimpliVity has created a set of tools and processes to simplify the deployment and upgrade of SimpliVity OmniStack nodes. These tools can be utilized by SimpliVity's Remote Deployment Services and Customer Support teams (requiring no onsite presence by SimpliVity personnel), SimpliVity field personel, certified partners, and end customers.

### 4.2.1. Deployment

The deployment process of a new SimpliVity infrastructure starts with the Pre-Flight Checklist, where a customer will indicate environment details necessary for deploying the OmniStack nodes. These details can then be entered into the SimpliVity Deployment Manager wizard 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. The Deployment Manager will automatically detect the nodes on the network and can push down all the configuration necessary, including indicating the hypervisor to be installed.

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.

### 4.2.2. Upgrades

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

## 4.3. Plug-Ins for Native Hypervisor System Management Applications

Core to SimpliVity's platform 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. SimpliVity delivers this streamlined singular global unified management platform by interfacing with existing management frameworks used to manage the virtualized environment.

SimpliVity offers a GUI plug-in for VMware's vSphere Web Client. Central administrators can implement data protection policies, monitor resources and troubleshoot problems leveraging the same vSphere client they already use to manage their ESXi environment. All administrative functions are performed across a Data Center 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 SimpliVity technology.

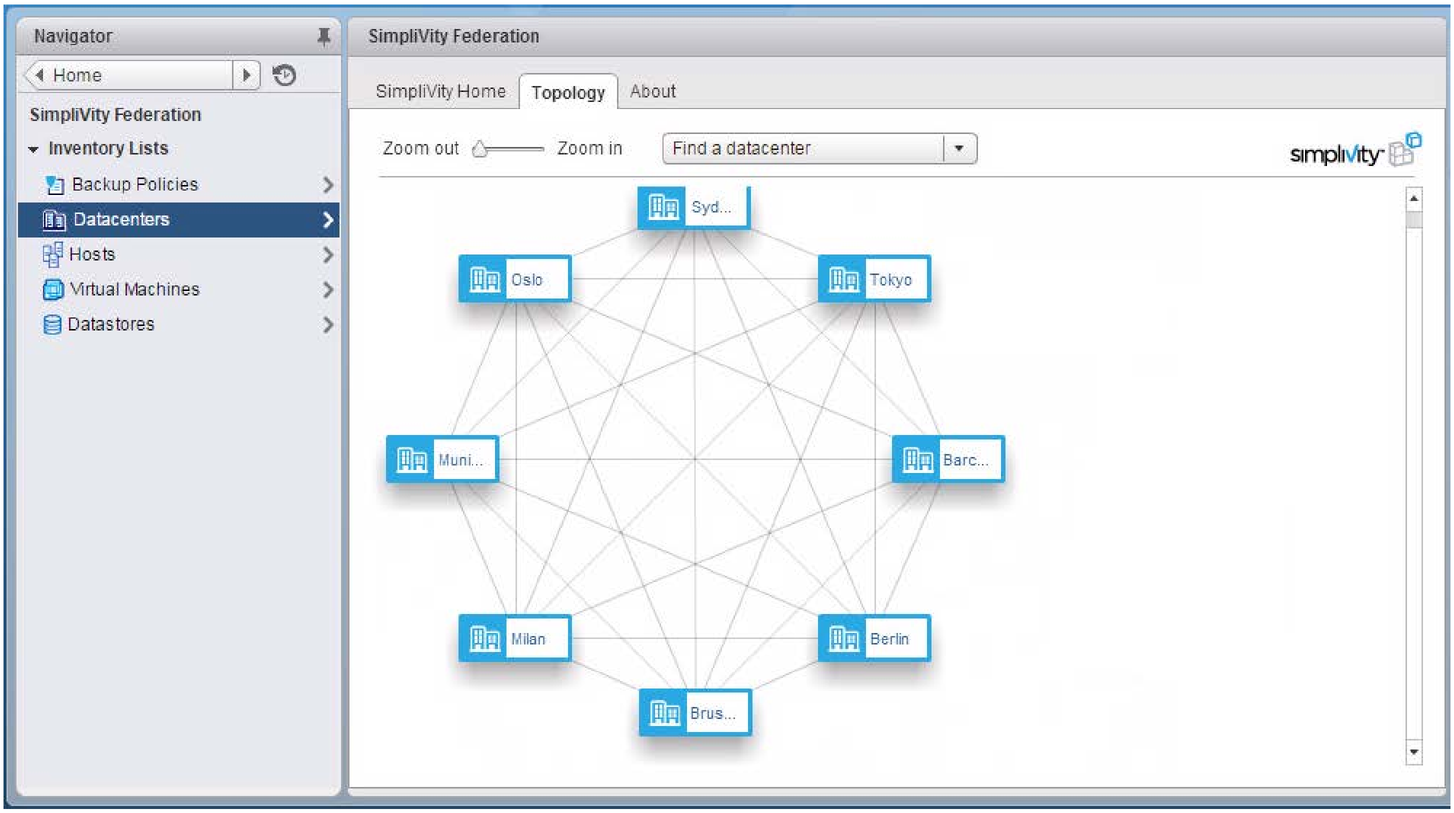


Figure 9: vSphere Web Client integration

## Managing SimpliVity Data Centers, Federations and Nodes

SimpliVity's vSphere Web Client plug-in includes tools for defining and managing SimpliVity Data Centers and Federations. Administrators can add, move, and delete nodes from SimpliVity Data Centers and Federations and configure certain node settings.

Federation-level management is enabled by simply adding hosts into different vSphere Datacenter. These Datacenters 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 a SimpliVity Federation that can span across the globe.

The vSphere Web Client provides configuration, inventory, and status displays for individual OmniStack nodes (see Figure 10).

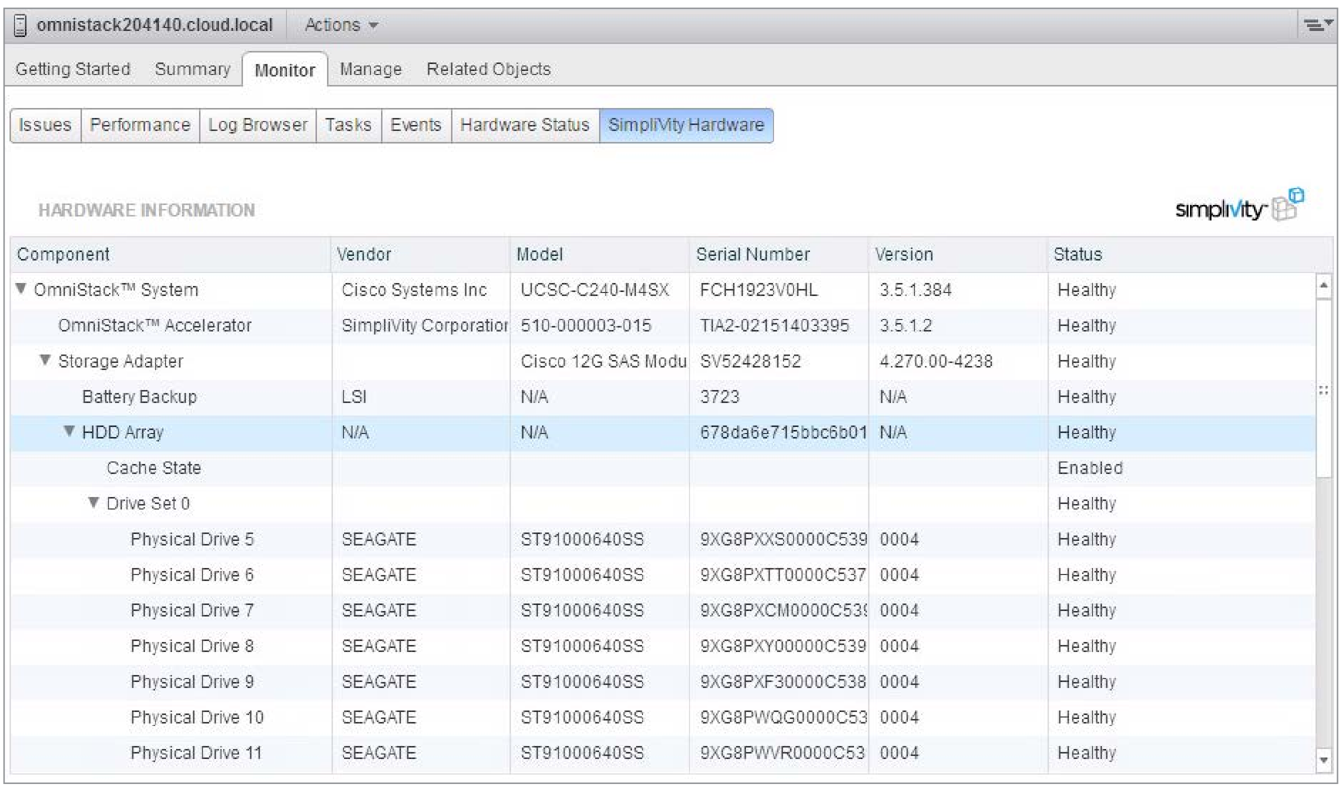


Figure 10: SimpliVity vCenter node view

## Capacity and Performance Management

The vSphere Web Client plug-in provides summary and detailed performance and utilization displays for SimpliVity Data Centers and Federations including physical and logical storage capacity views (free, used, total space, etc.) and VM performance views (throughput, IOPS, and latency strip charts) (see Figure 11).

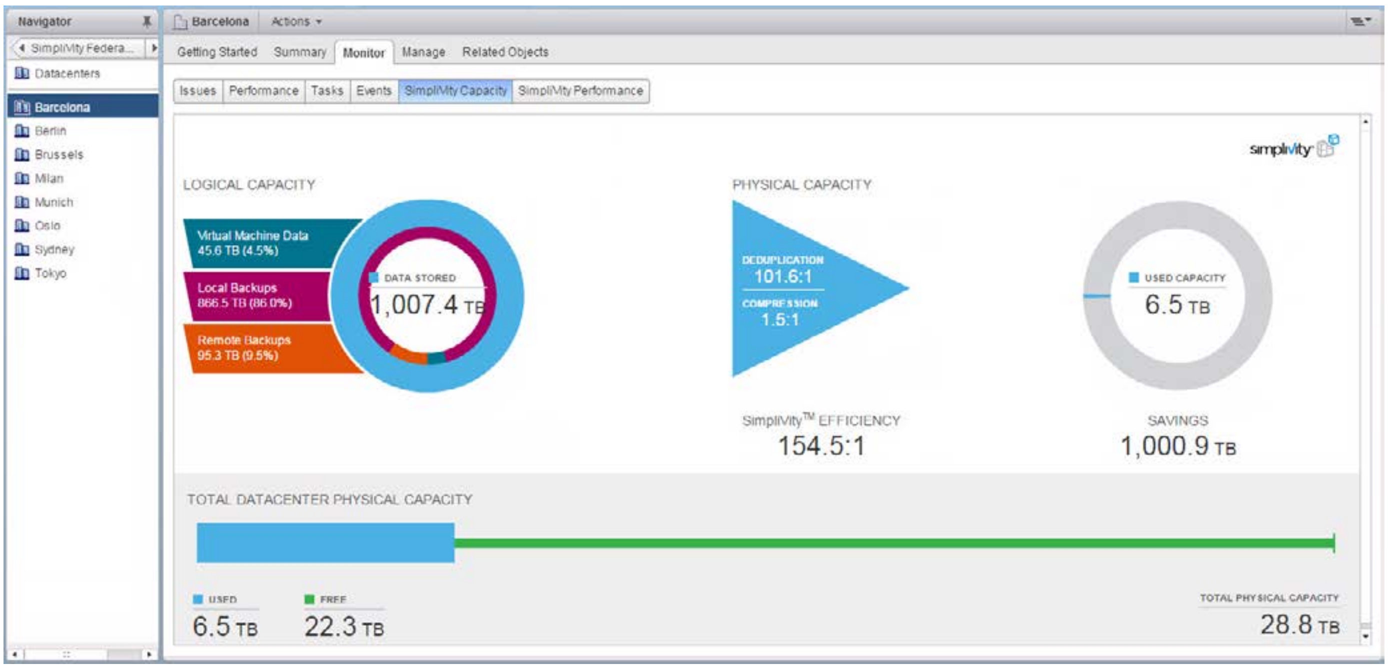


Figure 11: SimpliVity vSphere Web Client capacity view

To provide the best application performance across a SimpliVity Data Center, SimpliVity has created the OmniStack Intelligent Workload Optimizer, which includes the intelligent placement across nodes of the replica set when a virtual machine is created and an integration with vSphere Distributed Resources Scheduler (DRS). This integration with DRS brings awareness of the location of the data within the SimpliVity Data Center to maximize the performance of storage access. Intelligent Workload Optimizer creates and maintains DRS Host and VM Groups that are combined into "Suggested Affinity" DRS Rules so that DRS will prefer that virtual machines 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 anytime a new VM is created or deleted, as well as on a regular cadence to correct any manual changes that were inadvertently made by an administrator.

This integration into DRS ensures that virtual machines and their data to remain colocated. This delivers maximum storage performance relative to the approach of spreading the data throughout the datacenter and either requiring network-based storage access or having the data follow the virtual machine every time it is migrated by DRS. This "follow the VM" approach results in more CPU usage, more network usage, and more storage capacity usage.

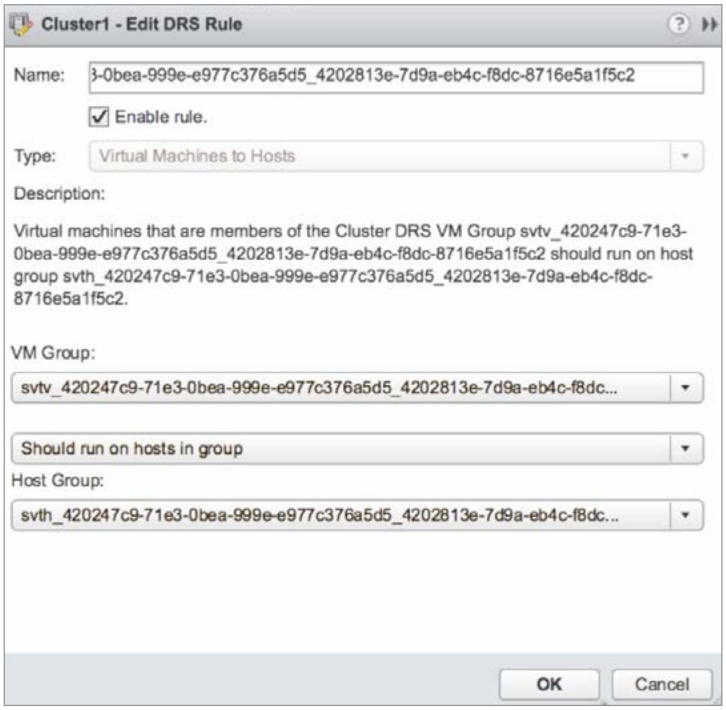


Figure 12: DRS Rule created by SimpliVity's Intelligent Workload Optimizer

## 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 virtual machines and to manually backup and restore VMs (see section 4.1.1 for supported functions).

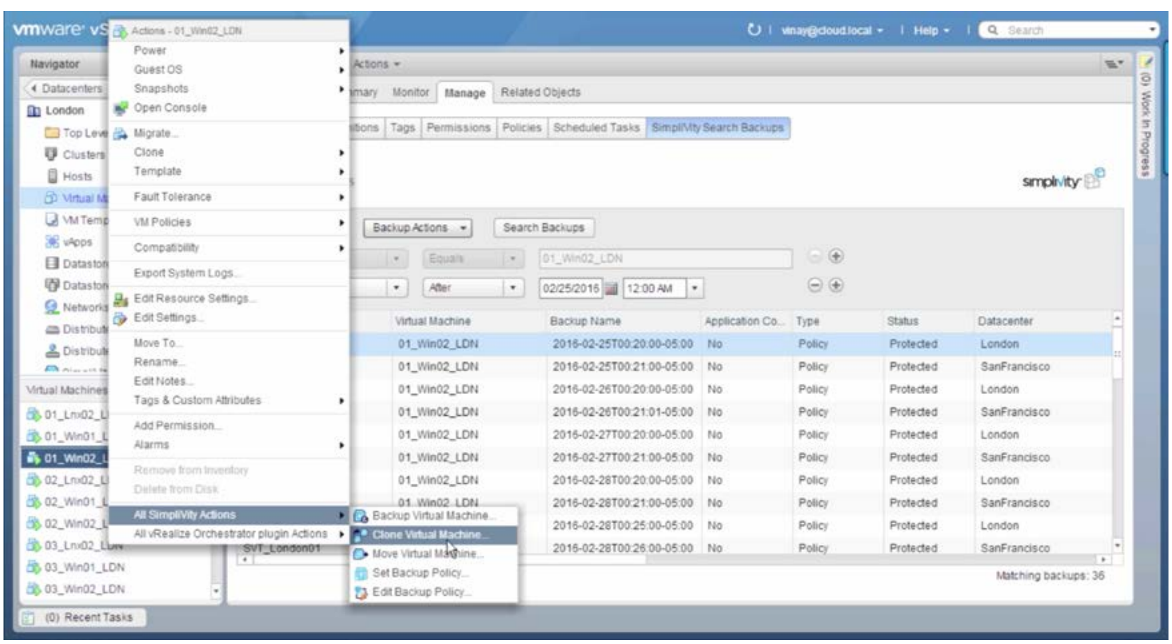


Figure 13: SimpliVity Data Protection example

## 4.4. Plug-Ins for Third-Party Service Orchestration Applications

SimpliVity offers plug-ins for third-party service orchestration applications such as VMware vRealize Automation and Cisco UCS Director. 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.

SimpliVity's Data Virtualization Platform simplifies the configuration of automated workflows by reducing the amount of disparate infrastructure technologies from as many as twelve to one. With SimpliVity's integrated solution, users can also backup, restore, clone, and move virtual machines 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.

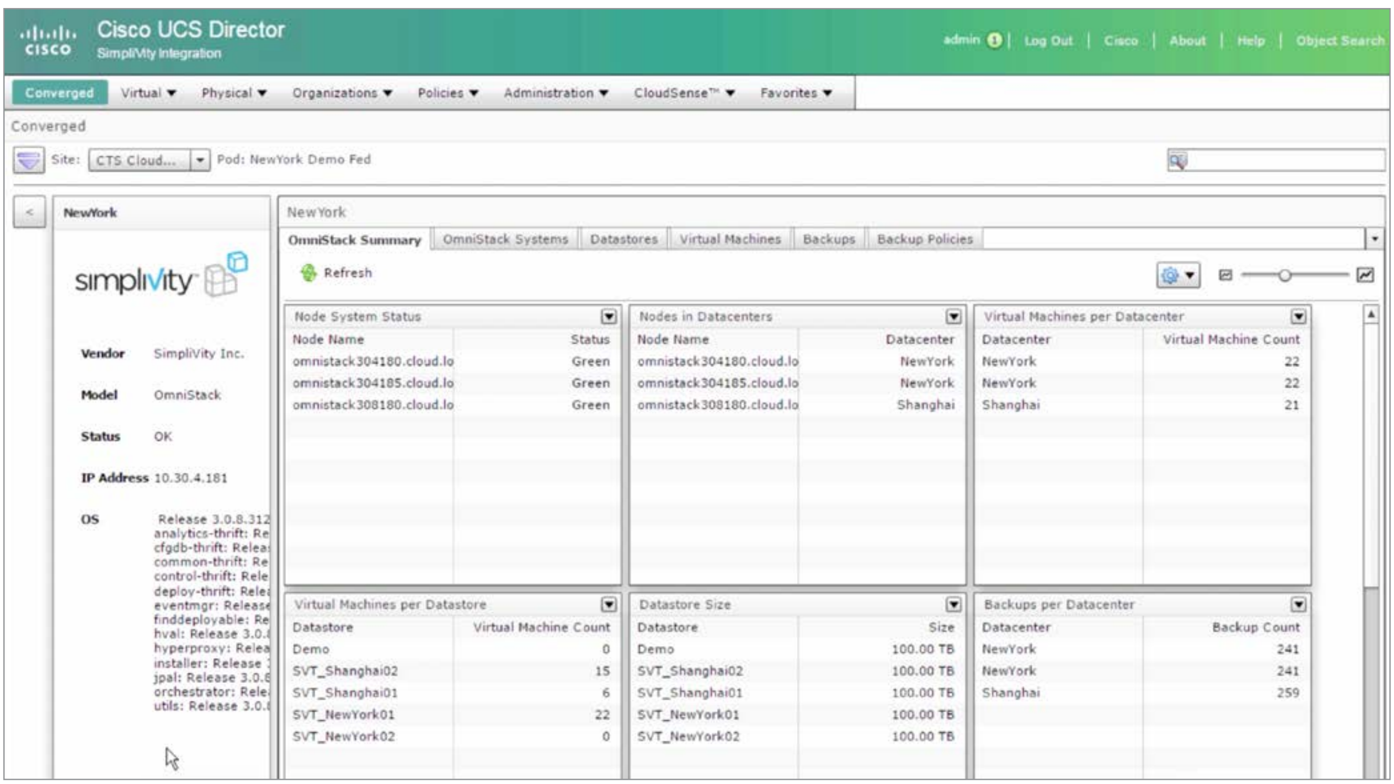


Figure 14: SimpliVity integration with Cisco UCS Director

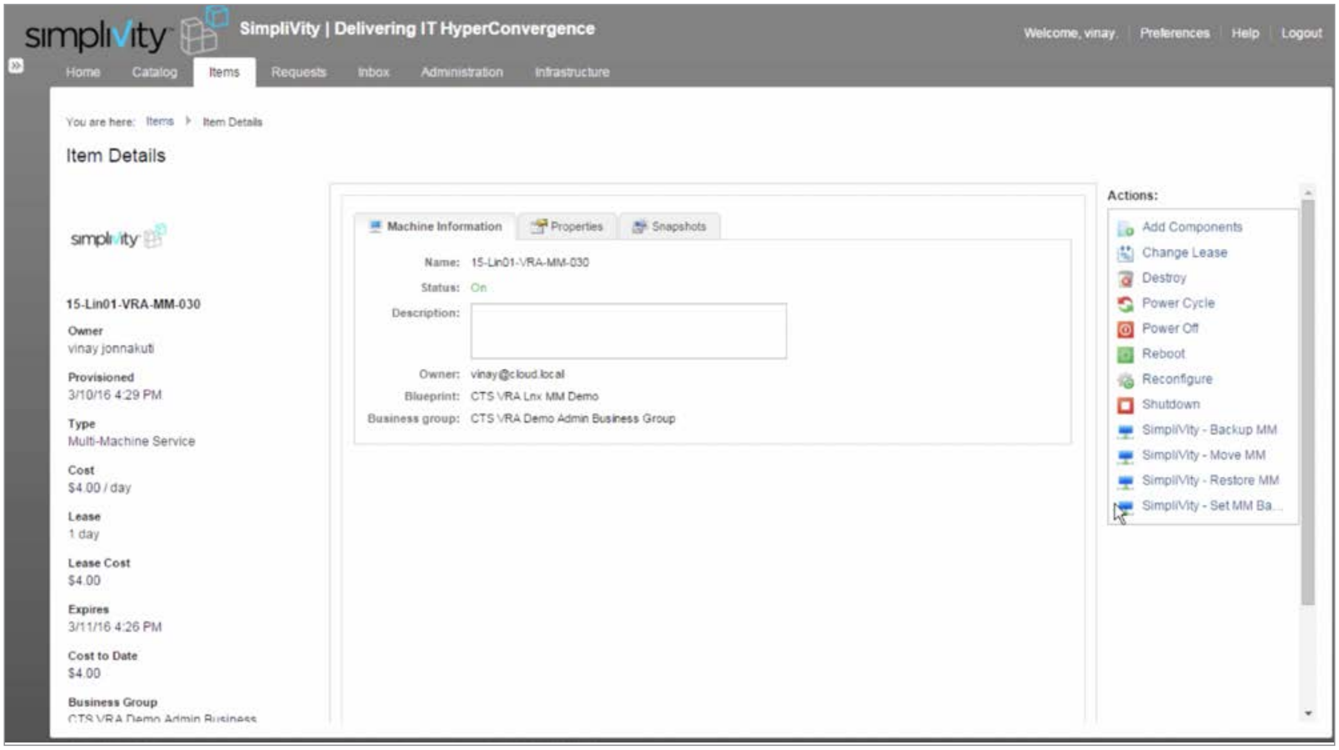


Figure 15: SimpliVity integration with VMware vRealize Automation

## 4.5. Native Management Interfaces and APIs

### 4.5.1. Command Line Interface

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 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 OmniStack Virtual Controller and is accessed via a Secure Shell (SSH) session from a terminal emulator. A single CLI session (to one OVC) provides full visibility and management across an entire Federation.

Examples of CLI command sets include:

- **Federation and SimpliVity Data Center commands** – organize and manage nodes within a Federation or Data Center
- **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 backup, clone, restore, and move VMs; apply backup policies to VMs
- **Software management commands** – manage SimpliVity software releases (upgrade, rollback, commit, etc.)
- **Support commands** – configure Phone Home settings and other maintenance options

### 4.5.2. Programmatic APIs

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

Please refer to official SimpliVity documentation for details on utilizing these APIs. Interactive documentation for the REST API is available on every SimpliVity OmniStack system at https://OMNISTACK_IP_ADDRESS/api/index.html.

## 4.6. OmniView

To improve visibility of metrics within a SimpliVity-based vSphere environment, SimpliVity created a tool called OmniView. This tool was originally developed for use within SimpliVity Customer Service 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 SimpliVity and vSphere data will be collected by the OmniStack Virtual Controller and forwarded to a SimpliVity Support Cloud database, where the data can be accessed through a secure web interface on the SimpliVity Online Support Center.

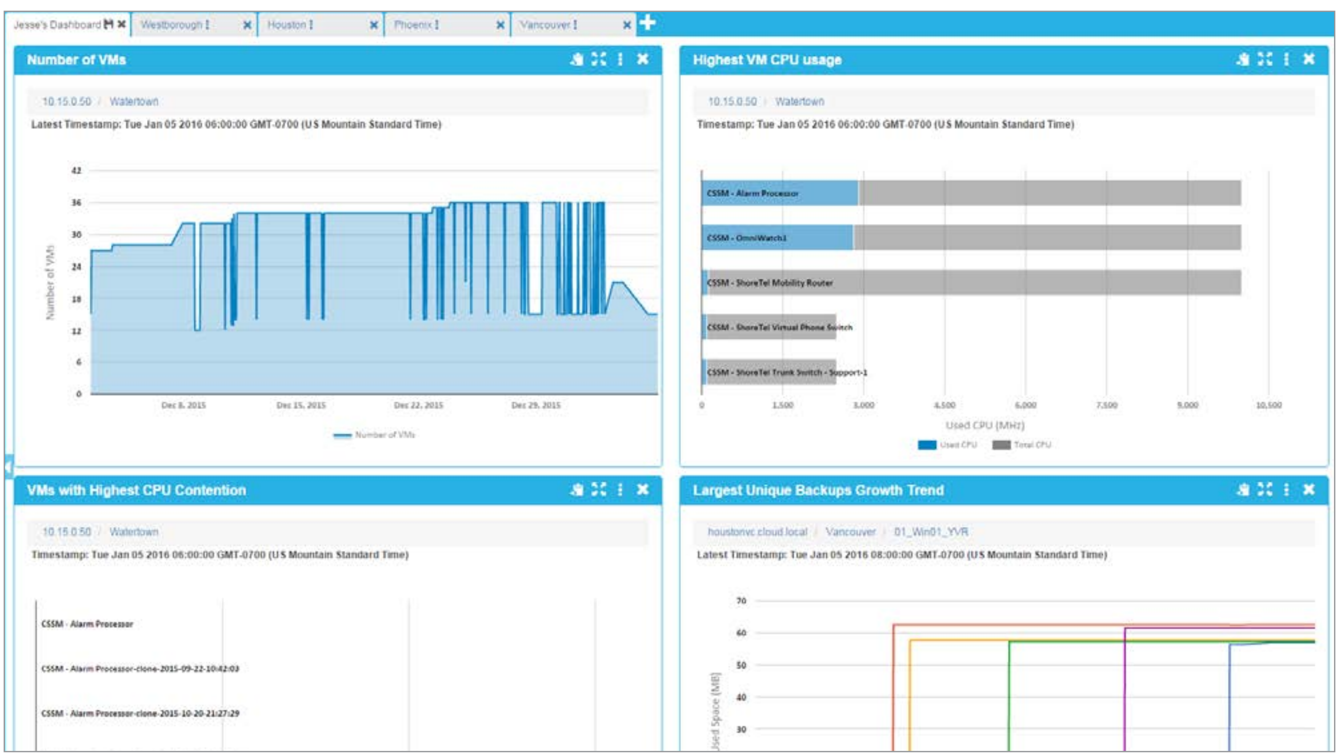


Figure 16: The OmniView dashboard

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 SimpliVity and vSphere metrics collected across multiple data centers and 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.

## 5. SimpliVity Hyperconverged Infrastructure Benefits

SimpliVity OmniStack systems are built to provide the Best of Both Worlds – Enterprise-grade capabilities and Cloud-like economics and flexibility. These benefits include the following:

### 5.1. Enterprise Capabilities

- **Improved business continuity and disaster recovery** – SimpliVity's highly resilient architecture and rapid data protection features reduce risks and uncertainty and improves availability for business-critical applications.
- **Improved service velocity** – SimpliVity's integration with system management and service orchestration tools improves time-to-value and time-to-market.
- **Better application performance** – SimpliVity's unique data management approach 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** – SimpliVity unifies global management and integrates into established administrative systems, thereby streamlining operations.

### 5.2. Cloud Economics

- **Lower CapEx** – 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** – SimpliVity's 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** – SimpliVity designed management around the VM, utilizing policy-based controls and abstraction from the underlying infrastructure to allow simple VM management and easy VM mobility.
- **Better business innovation** – SimpliVity allows IT organizations to spend less time managing underlying IT infrastructure, and more time focusing on strategic business initiatives.
- **Greater flexibility** – SimpliVity provides choice in hardware platforms, hypervisors, management systems, and deployment options.
- **Invest-as-you-grow economics** – SimpliVity's scale-out architecture minimizes upfront capital investments by "right-sizing" the initial investment and it avoids lopsided business models with long payback periods.

## 6. SimpliVity Hyperconverged Infrastructure Use Cases

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, SimpliVity customers drive 40:1 data efficiency while fully protecting their data.
- **Data migration** – SimpliVity's VM-centric approach to data management and efficient data architecture enable 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 SimpliVity's unique data migration techniques, 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.

## 7. Additional Information

For more information about SimpliVity hyperconverged infrastructure, see the following resources:

SimpliVity OmniStack Integrated Solution with Cisco UCS Reference Architecture

SimpliVity OmniStack Solution with Lenovo System x3650 M5 Solution Brief

SimpliVity OmniStack and VMware vRealize Automation Reference Architecture

The Technology behind 100:1 Data Efficiency White Paper

Native Data Protection with SimpliVity Solution Brief

Hyperconvergence Delivers Data Efficiency Solution Brief

ROBO Solution Brief

SimpliVity Corporate Blog

## 8. 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 per Second

L2 – Layer 2

LUN – Logical Unit Number

NFS – Network File System

NIC – Network Interface Card

OAC – OmniStack Accelerator Card

OVC – 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

SNMP – Simple Network Management Protocol

SSD – Solid State Drive

SSH – Secure Shell

VAAI – VMware APIs for Array Integration

VDI – Virtual Desktop Infrastructure

VM – Virtual Machine

WAN – Wide-Area Network

### For more information, visit:

www.simplivity.com

© 2016, SimpliVity Corporation. All rights reserved. Information described herein is furnished for informational use only, is subject to change without notice, and should not be construed as a guarantee, commitment, condition or offer by SimpliVity. SimpliVity, the SimpliVity logo, OmniCube, OmniStack, CN-5000 and Data Virtualization Platform are trademarks or registered trademarks of SimpliVity Corporation in the United States and certain other countries. All other trademarks are the property of their respective owners.

J352-DeepDive-WP-EN-062016