

TRANSFORMING THE DATA CENTER

SIMPLIVITY DELIVERS HYPERCONVERGED PLATFORM WITH NATIVE DATA PROTECTION

JANUARY 2016



Hyperconvergence has come a long way in the past five years. Growth rates are astronomical and customers are replacing traditional three-layer configurations with hyperconverged solutions at record numbers. But not all hyperconverged solutions in the market are alike. As the market matures, this fact is coming to light. Of course, all hyperconverged solutions tightly integrate compute and storage (that is par for the course) but beyond that similarities end quickly.

One of the striking differences between SimpliVity's hyperconverged infrastructure architecture and others is the tight integration of data protection functionality. The DNA for that is built in from the very start: SimpliVity hyperconverged infrastructure systems perform inline deduplication and compression of data at the time of data creation. Thereafter, data is kept in the "reduced" state throughout its lifecycle. This has serious positive implications on latency, performance, and bandwidth but equally importantly, it transforms data protection and other secondary uses of data.

At Taneja Group, we have been very aware of this differentiating feature of SimpliVity's solution. So when we were asked to interview five SimpliVity customers to determine if they were getting tangible benefits (or not), we jumped at the opportunity.

This Field Report is about their experiences. We must state at the beginning that we focused primarily on their data protection experiences in this report. Hyperconvergence is all about simplicity and cost reduction. But SimpliVity's hyperconverged infrastructure also eliminated another big headache: data protection. These customers may not have bought SimpliVity for data protection purposes, but the fact that they were essentially able to get rid of all their other data protection products was a very pleasant surprise for them. That was a big plus for these customers. To be sure, data protection is not simply backup and restore but also includes a number of other functions such as replication, DR, WAN optimization, and more.

For a broader understanding of SimpliVity's product capabilities, other Taneja Group write-ups are available. This one focuses on data protection. Read on for these five customers' experiences.

Evolving Hyperconvergence

Hyperconverged infrastructure was engineered for running virtual workloads. For example, the storage controller, which would traditionally be part of the SAN hardware, is now run as a software service at the hypervisor level on every node in the cluster. This means that all storage in the appliance or cluster of appliances can be dynamically managed among all of the VMs. With

hyperconverged infrastructure, enough of the IT stack is now abstracted and integrated below the hypervisor that the VMs can be effectively decoupled from the underlying hardware, allowing the management focus to move from the IT components to the VMs.

Data protection was a missing element of original converged offerings--users had to rely on third-party packages. Some of the early hyperconverged infrastructure (HCI) vendors layered some data protection capabilities on to their platforms, but only as add-ons, not tightly integrated.

Data deduplication is another technology that has largely been missing from hyperconverged offerings. Data protection and data movement across the WAN depend on deduplication to minimize the amount of data being transferred and stored, to improve performance and use resources more efficiently. Deduplication should be global in a hyperconverged system, and be performed for all data.

Enter SimpliVity

SimpliVity understood that the key problem addressed by previous HCI vendors--virtualizing the storage aspect of the infrastructure and getting the core storage functionality up into a VM--was not sufficient. To deliver true hyperconvergence, SimpliVity went further and made data optimization, not just hardware integration, a fundamental element of their architecture. Doing that also enabled them to pull built-in data protection capabilities into their hyperconverged stack and provided intrinsic WAN optimization.

The core of SimpliVity's offering is the OmniStack Data Virtualization Platform, which is the operating software that runs on their OmniCube hardware appliance, as well as on a variety of third-party hardware. A collection of SimpliVity systems is called a federation (within a single site or across multiple sites). The Data Virtualization Platform exists as one entity globally across all nodes and sites in a federation, enabling efficient management and data movement between sites. When users add another system to the federation, it's automatically recognized and auto-configured; no manual intervention is required.

Logically, the Data Virtualization Platform is composed of three major elements:

1. **Accelerated Data Efficiency.** This means that the Data Virtualization Platform performs global inline deduplication and compression at ingestion before anything is written to disk, and maintains it that way throughout all operations. A PCI-e card, the OmniStack Accelerator, is used to assist this process. Not only does this approach deliver capacity savings, it also improves application performance by eliminating redundant reads and writes to and from disk. As a result, the solution delivers consistent performance to applications
2. **Global Unified Management.** The interconnected federation structure is what enables global unified management, the second element of the Data Virtualization Platform. That means that resources in the federation are easily managed at the VM level. 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 administrative actions are abstracted away from the underlying infrastructure, and, instead, are based on actions and policies applied to the VM.
3. **Built-in Data Protection.** The combination of accelerated data efficiency and global unified management enables the third element of the Data Virtualization Platform, built-in data protection. A VM is represented by metadata that points to the deduplicated data blocks that make up the VM.

Effectively, the metadata contains the formula for re-creating the VM, so when a user creates a backup they simply make a copy of the metadata.

How it Works: Architectural Overview

Looking more closely at SimpliVity's physical architecture, we see that the Data Virtualization Platform is an abstraction layer that resides between the hypervisor and the hardware and is used to decouple the VMs from the underlying hardware infrastructure. It comprises two sub-layers, the Presentation Layer and the Data Management Layer. The Presentation Layer abstracts the storage beneath it and gives each VM its own view into the datastore. When a VM writes data, it's registered in the presentation layer, then that data is deduplicated and compressed. When the VM reads it back, it wants the original, complete version, not the deduplicated, compressed version. The Presentation Layer manages that.

The Data Management Layer is where the actual deduplication and compression are done and where all data blocks are tracked. This layer also assembles the compressed blocks into full stripe writes to the drives. Whenever a write of a unique data block is done, it's also written to another node for redundancy. Then, if one node fails, there's still a copy of the data that was on that node available elsewhere in the federation. The Data Management Layer of each SimpliVity hyperconverged infrastructure node in a federation communicates with other layers so that all data is shared.

The OmniStack Accelerator card is an important component of the Data Virtualization Platform. Deduplication and compression, in general, place a heavy load on processor resources. This can directly impact production workloads and cause significant variations in performance. The OmniStack Accelerator offloads the main processing by intercepting VM writes and performing the deduplication and compression before data blocks are written to disk. By reserving the main processor cycles for the production workloads, this reduces system latency and decreases the amount of I/O being sent to the disks. This not only increases the overall level of performance, it decreases the variability in performance.

WHAT BUILT-IN DATA PROTECTION DOES

Due to the OmniStack architecture, integrated, simple and efficient data protection is a natural extension of the platform. VMs are decoupled from the underlying hardware via the hypervisor and the abstraction mechanisms of the DVP. This standalone nature makes them easy to manipulate and move around. Since VMs are completely described by metadata, backups or clones are easy to create; you simply make a copy of that metadata within the Data Management Layer (where all backups are maintained) and your backup is done. This means that there's no need for incremental backups, because the full backups have such a small footprint. Restore is point-and-click globally for all VMs.

Replication and DR are similarly uncomplicated. As mentioned above, to replicate a backup to another site, you merely copy over the VM metadata. The target site determines what, if any, additional data blocks need to be copied over to reconstitute the VM. If the original site goes down, the VM copy can be spun up at the new site to bring the workload back on line. Because data within the system is always in a deduplicated and compressed state, there's much less data flowing over the LAN/WAN whenever data blocks or VMs are moved around. This adds to the ease with which replication can be done.

The VM mobility, data efficiency and high level of integration also enabled SimpliVity to abstract their backup policies from the underlying infrastructure. To specify a complete backup policy for an individual VM, you need only define three things: how often do you want to copy it; where do you want to keep the copies; and how long do you want to keep them. Because of the global nature of the

SimpliVity federation, you can open a policy for a given VM, edit the rules and save it and it's automatically copied to all other sites. Policies also automatically follow VMs as they're moved around the federation. There's no requirement to manually edit them to match up with the new location; the system does it transparently in the background.

SimpliVity also allows users to send backups to the Amazon cloud for disaster recovery along with a variety of regional service providers.

SimpliVity Customers in the Real World

We spoke with five SimpliVity customers to get their detailed feedback on how and why they use SimpliVity in their enterprise data centers.

DAIRYLEA MEETS RECOVERY OBJECTIVES

DairyLea Cooperative is the largest milk-marketing organization in the northeastern U.S., with more than 2000-member dairy farms and two main data centers. Jeremy Wheeler is DairyLea's Innovation Architect.

The production environment is 100% virtualized. Before switching to SimpliVity, the company's storage and data protection consisted of three HP EVA storage arrays and Symantec BackupExec. The

"Disaster recovery was a 100% important driver for us. We also needed dedupe and compression, and a WAN optimization tool that allows us to send that data across the wire."

IT team replicated from the main data center to an iSCSI SAN at the DR site. It took 48 hours to back up DairyLea's largest VM and backups frequently overran allotted windows.

The company looked for DR failover, RPOs under 24 hours and RTOs under four hours, high performance replication, and simple manageability. They found all of this and more in SimpliVity. DairyLea bought four OmniCube CN-3000s for its Syracuse data center and two for its backup site. SimpliVity accelerated backup and restore, deeply reduced data sizes with dedupe and compression, and simplified administration.

Major Benefits

- **Improved data protection and DR.** Data protection is much stronger than it was with the legacy SAN and far less expensive.
- **Simplified backup management.** IT backup management time went from 20-30 hours a week to 5. It's so simple that SQL and Oracle admins run their own backups.
- **Lower cost.** Big savings on yearly storage and data protection maintenance fees.

HIGH COTTON PROTECTS MISSION-CRITICAL APPLICATIONS

High Cotton is a data processing and compliance firm that owns two 100% virtualized data centers in Alabama and Texas. IT Director Jeff Byers' top priority was to expand the data centers while protecting the data and controlling costs.

High Cotton's initial deployment was two OmniCube 3000s in Alabama and one OmniCube 2000 in Texas. As the company started growing, they bought additional 3000s for each site. Backups travel at all hours between the SimpliVity nodes.

Today all of their critical servers run out of the production data center. The production environment stores 3.7PB and backs up 14TB of data to Texas, which has 1.8PB capacity. SQL Server sends data remotely every 15 minutes as do the file servers that store their client applications. Other servers back up daily or weekly.

Major Benefits

- **Data protection for mission-critical applications.** “If we lost them, basically the company would go under,” Byers said. SimpliVity enables performance and protection for these crucial applications.
- **Simplified storage management.** A single IT administrator now manages the entire platform.
- **Ease of scalability.** SimpliVity makes it simple to add nodes when needed.
- **Support.** “With SimpliVity, it’s not just about sales. These guys are knowledgeable and they are nice. Their engineers will jump in and help on a moment’s notice, even weekends and nights, and they don’t complain about it,” said Byers.

“Our number one benefit is just having backup that works. Before this, a lot of our backups were based on prayer! Now I know that they’re working and I can also back up to three different places. I can failover my virtual machines and in five minutes I have them back. I’m more comfortable than I’ve ever been.”

ROMAC INDUSTRIES GETS HIGH AVAILABILITY

Based in Bothell, Washington, Romac manufactures pipe products and tools for the water and waste water industry. Russell Turley is Romac’s IT Manager.

“Frankly, I think SimpliVity is a very elegant solution. In a single point of failure environment, with one failure you’re down. SimpliVity’s high availability fixes that problem.”

The company’s aging legacy environment consisted of XenServer virtualization, five-year-old Dell servers, and a four-year-old EMC CX 4 SAN supporting XenServer and Oracle. The three-person IT team backed up some critical systems during the day, and backed up nightly incrementals to tape. Over the weekends, they quiesced systems and did a full backup to tape. But within a four-month period, the SAN experienced frequent disk failures, two failed processors, and two major crashes.

They replaced their old environment with SimpliVity and Cisco UCS servers. There is no single point of failure throughout the virtualized environment and performance is faster. Value-to-cost was excellent. “We wanted one solution that would fix it all instead of buying a SAN, multiple servers, and VMware, and then getting it all configured and working. We wanted one solution that fixed everything for us. That turned out to be SimpliVity,” Turley said.

Major Benefits

- **High availability.** Disk-level RAID and failover all contribute to high availability in the data centers. “In a single point of failure environment, with one failure you’re down. SimpliVity’s high availability fixes that problem,” Turley said.
- **Performance.** Romac’s SolidWorks is a mission-critical application with high performance requirements. It now runs about 4 times faster than on the original servers.

- **Cost and space savings.** The hyperconverged architecture saved hundreds of thousands of dollars over buying traditional solutions, and operated in half the rack space.

KHD SIMPLIFIES DISASTER RECOVERY

Joerg Ludwig is the CIO and Global Head of IT for Germany-based KHD. Ludwig is responsible for multiple company data centers around the world.

KHD's legacy environment contained HP servers, NetApp filers, and VMware farms. Key issues included DR complexity and staffing challenges. The legacy system ran a mix of Symantec and Commvault Simpana with local tape backups. The reliance on tape backup concerned Ludwig. He considered using NetApp replication and VMware SRM, but found the complex management too time-consuming.

"It is crucial to have a good DR site since we track inventory for our customers. This is a very responsible position and needs lots of customer trust. SimpliVity lets us do this huge job well."

KHD adopted SimpliVity for its efficient architecture and native data protection that replaced multi-vendor data protection products. Data replicates between primary and DR sites and is easily changed as strategy dictates. KHD's critical applications run on the OmniCubes including Exchange 2013 and NetSuite CRM. Many other applications run on OmniCubes as well including SAP, databases, and main infrastructure services.

The team prefers to run core services on SimpliVity because of its ease of replication, highly simplified management, and performance. Ludwig also

appreciates SimpliVity's single pane of glass management and the ability to use generalists instead of specialists.

Major Benefits

- **Simplified data replication.** VM-level policies and a simple management interface make it easy for IT admins to manage replication.
- **Improved reliability and performance.** The overall architecture increased performance and reliability for both primary and secondary storage.
- **Hard cost savings.** KHD used to invest heavily in tape libraries. SimpliVity lets them save on those hard costs.

SIMPLIVITY A GAME CHANGER FOR FRANCIS DRILLING FLUIDS

Francis Drilling Fluids provides drilling fluids, logistics management services, and environmental consulting to the oil and gas industries. Steve Schaaf is its CIO.

FDF's data center housed an IBM DS3500 SAN and IBM blade servers. Schaaf was interested in evolving to a "liquid computing" model where commodity servers, storage, and data protection easily move in and out. He chose SimpliVity for a pilot project. In 2013, the team bought an OmniCube with 3TB useable capacity for the rollout. Deployment was simple and straightforward and capacity was more than enough.

Schaaf referred to SimpliVity as a “game changer.” They invested in three more OmniCubes operating in a primary and a remote DR data center. SimpliVity data protection replaced Backup Exec and Veeam.

In 2014, when SimpliVity released support for OmniStack with Cisco UCS, Francis Drilling seamlessly migrated from OmniCube to four OmniStack with Cisco UCS C-series systems across their two data centers. Today the company runs a 100% virtualized environment on SimpliVity with Cisco including MS SQL 2010, Exchange 2013, SharePoint, customer call center software using Cisco Call Manager, and much more.

Major Benefits

- **Extremely reliable DR.** SimpliVity enabled FDF to protect a DR site containing critical customer-owned data.
- **Simple data protection management.** “It’s so beautiful to manage my backup and DR all at once from the same place. We just make the settings and get the reports,” Schaaf said.
- **Ease of support.** The few times Schaaf had to call, SimpliVity engineers identified the problem and fixed it – even when the problem was not SimpliVity’s.
- **Supports company growth.** SimpliVity provides the ability to seamlessly import and migrate applications and data to simplify data center consolidation with mergers and acquisitions.

“SimpliVity is working day by day and night by night, I don’t have to worry about it. You can so easily schedule replication and service level classes, and you do it with non-experts because it’s so simple. This is a global policy for my team. They only need low-level VM know-how, which frees up their time and mine.”

Taneja Group Opinion

When customers initially buy a hyperconverged system from any vendor, they are often not even thinking about data protection and other secondary uses of data. They are mostly motivated to simplify their primary storage environment, reduce cost, and improve application performance. What they often discover is that they get all these benefits and more. They get on-demand scalability. They get reduced footprint. They replace several vendors’ products with a simpler configuration. And they get VM-centricity and more.

But they often still have to buy separate backup, replication and DR software; media servers; WAN optimization and data deduplication appliances; to name a few. So while hyperconvergence simplifies their primary storage environment, the secondary storage environment remains unchanged, with all the issues we are familiar with over the years (missed backups, unreliable backups and recovery, unpredictable DR, WAN costs, and so on).

This is where SimpliVity’s OmniCube architecture comes to the rescue. Because it is architected from the outset for both primary and secondary uses of data, it simply deals with these issues inherently. The same level of simplicity that they built into the primary storage side is built into the secondary storage side. This makes its architecture unique in the industry and its products stand out from its competitors.

We knew this architectural difference but we were unsure how the customers were responding to the secondary data side, especially backup/recovery and DR. This was the purpose of this Field Report. What we discovered from interviewing five customers is that most, if not all buyers thought they got data protection and DR (and more) for free. They may not have bought SimpliVity hyperconverged infrastructure for data protection purposes but they discovered that their old nightmares simply vanished. This was an unexpected delight. This is what the interviews represent.

We don't expect a customer to buy SimpliVity hyperconverged infrastructure solely for data protection purposes but they will find that data protection simply disappears in the background so they can focus on more important IT priorities.

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.