

## **SIMPLIVITY: TRANSFORMING THE DATA CENTER WITH VIRTUALIZATION AND STORAGE CONVERGENCE**

**AUGUST 2012**



The data center has fast been transformed by the emergence and mainstream adoption of virtualization over merely the past decade. Virtualization has changed the ability of IT to deploy and manage workloads, and lent tremendous power to the administrator for manipulating those workloads in clever ways. At onset, virtualization first became appealing to most users because it promised to homogenize a rather difficult part of the physical data center through a layer of software abstraction. Such abstraction would make configuration and deployment difficulties associated with physical server hardware and fat operating systems melt away. The rewards from this undertaking were tremendous – untold buckets of operational dollars were saved by avoiding time and effort intensive rack, power, install, and configure cycles that would occur repeatedly to facilitate application development, testing, production deployment, cycle replacements, break/fix, and more. The changes have infused the business with a new ability to leverage IT, and to do so at lower cost and with less risk of disruption.

The transformation is not yet complete, as the full promise of virtualization is inhibited by the underlying physical infrastructure of the data center. Realization of the need to address the infrastructure complexity problem is driving a flurry of innovation in the market. With an eye on what we call hyperconvergence, we'll briefly survey innovations emerging in response to on-going virtualization challenges, evaluate how these technologies will impact the datacenter over the next few years, and highlight one vendor as an early innovator bringing these changes to the market: SimpliVity. Unlike early players who have converged a few aspects of the IT Infrastructure stack – rudimentary storage and server functionality – SimpliVity has assimilated all the functionality of the IT Infrastructure on to a single platform. Each such unit – OmniCube, as SimpliVity calls it – offers a complete set of data center infrastructure functionality at a fraction of the acquisition and operating costs of separate IT systems. But hyperconvergence will bring about change far bigger than costs, and may well transform how IT is done. Let's take a look.

### **RIPPLES FROM THE FIRST WAVE**

Today, hypervisors running on server hardware have cast their influence over nearly every IT domain. But they have remained only an influence, and have not fully consumed all of those other valuable services that come from physical systems in the data center – services like storage, data protection, DR, and communications and storage networking.

The data center has long been physical with high degrees of sophistication and specialization in separate physical systems. As virtualization – a layer of abstraction and simplification – sought to interact, it has remained focused solely upon supporting demands inside of the virtual infrastructure. This has restricted virtualization's ability to simplify other external data center systems, and over time, the virtual infrastructure began to duplicate functionality from those systems to better support the needs of the virtual workload.

The impact of this development is two-fold. First, the unit of consumption for IT services is never fully optimized, and the infrastructure remains less than fully efficient. Storage, networks, and other resources are often over-provisioned because they must be manually touched and are less able to automatically adapt compared to virtual workloads. Consumption of resources is less than “ideal” because of physical, separate systems. Second, the combination of physical systems plus duplicated virtual functionality means the data center has become more complex versus ten years ago. This complexity often limits the success and scale of customer virtualization initiatives. Ask any DBA to implement a new database solution on a virtual infrastructure in an enterprise data center, with requisite protection, DR, and associated services, and the tale that unfolds will be daunting to the most stalwart of IT practitioner.

The next wave of virtualization will be different, and it is starting now. Technology is entering the market built upon a foundation that delivers the key storage and networking services that come from traditionally physical systems in the data center, and on top of that foundation runs the very same hypervisors and virtual infrastructure that beats within the heart of nearly any data center today.

## THE NEXT WAVE - HYPERCONVERGENCE

We call this transformational wave of technology “hyperconvergence”. For the first time, an approach has emerged that combines all of the functionality of a data center in an appliance-like form factor that can be connected together to build an *entire* infrastructure in building block fashion.

Hyperconvergence stands in stark contrast to the latest iterations of convergence – an approach that packages existing sets of technology together into denser, more tightly integrated form factors that ease consumption, but still leaves them to be managed as separate units of functionality. Hyperconvergence is a seamlessly integrated whole, built upon *homogeneous* building block appliances that deliver all infrastructure functionality – compute, storage, and networking.

Starting with a layer of highly virtualized internal storage, hyperconvergence solutions cluster together a practically unlimited set of appliances, and can scale through the continued addition of more appliances. Joined together, those appliances will serve up a huge, seamless resource pool that can support virtual machines as well as provide some services like storage to legacy physical systems still sitting on the data center floor. With internal storage virtualized across the cluster, any building block can access any stored data, while the cluster’s network makes access to any network port or amount of bandwidth instantly possible. This makes it possible to deploy and run any workload anywhere without worrying about storage configurations or network configurations. The traditional multi-system physical infrastructure nearly vanishes because all resources are highly virtualized and seamlessly connected into a single pool. This will radically simplify provisioning, scaling, and failure avoidance, and redefine utilization patterns in the data center – much overprovisioning will vanish, and utilization will match what is actually consumed. Of course, by way of management tools, virtual administrators will be able to impose restrictions, pools, or barriers to facilitate organization and separation of workloads for multi-tenancy or security.

One vendor to most recently join the ranks of hyperconvergence pioneers is SimpliVity. SimpliVity advances the genre by building a storage foundation on top of some of the most innovative storage technologies around, while seamlessly integrating data management, networking and hypervisor functionality that are each holistically connected in a way that further extends all of their capabilities so that dependencies on external infrastructure are all but eliminated. A closer look will shed light on just what this next wave of virtualization may do for the infrastructure.

## SIMPLIVITY AND THE DATA CENTER TRANSFORMATION

SimpliVity is delivering their storage+compute+networking hyperconvergence offering in the form of a product called OmniCube. Using a resource pooling layer, many OmniCubes can be woven together into a single highly available, hyperconverged infrastructure – called the OmniCube Global Federation – that runs VMware’s ESXi across all nodes, and can serve up applications from any federated OmniCube. The OmniCube’s storage capabilities are a showcase of the best innovations around, and include a full bevy of storage functions (snapshots, replication, etc.), NVRAM and Flash-based IO acceleration, pervasive granular deduplication, and WAN optimization. Moreover, the OmniCube Federation is designed to manage systems distributed across the globe, and seamlessly integrates with compute in the Cloud. OmniCube does this with a network virtualization technology that masks and automates away the underlying IP addresses, controllers, LUNs, etc. Customers can in turn manage, protect, and share VMs using a unified global management interface that can view globally distributed systems (or an administratively restricted subset of systems) and manage the total Federation as a single entity, no matter the number or physical location - even including virtual systems running on the Amazon EC2 Cloud. With the initial release of OmniCube, all such management happens through a VMware vSphere management plugin.

Key to SimpliVity OmniCube’s capabilities is an underlying data architecture called OmniStack. OmniStack deduplicates and compresses all data at inception, and data remains that way throughout its lifecycle, no matter how it moves around within an OmniCube Global Federation. This efficient, underlying data layer makes possible all of OmniCube’s customer facing functionality such as replication, resource sharing, efficient cache accelerated performance, and Cloud integration. Since data is highly efficient, snapshots can be kept forever.

The OmniCube hyperconvergence solution is the foundation for a set of highly innovative virtual infrastructure capabilities:

- **Seamless resource access, without boundaries.** OmniCube creates a clustered storage pool accessible by any node, allowing any VM to be immediately deployed without capacity concerns or configuration of external systems. Expanding this pool of resources requires little more than plugging in an additional OmniCube. In our brief examination, OmniCube looks poised to change a recurring multiple hour configuration process into rack, plug, and click-to-go.
- **Pervasive data efficiency** – OmniCube pairs up granular solid-state-based storage tiering and capacity optimization (deduplication and compression) to optimize both performance and capacity. In turn, a single OmniCube should deliver superior VM density per unit of storage, and virtualize more infrastructure per unit of floorspace. This data efficiency runs throughout system capabilities, including OmniCube to OmniCube data movement, and in turn serves up WAN optimization as well as efficient space efficient golden-master snap-clones and long-term retention of snapshots.
- **Integrated, enhanced data protection** – Pervasive deduplication and high availability provide a foundation for in-place protection, and OmniCube in turn does away with the need for external backup. Each VM is offered its specific protection scheme in terms of RPO (frequency of snaps) as well as replication of such snaps to any OmniCube globally, including the Public Cloud. Any such snap/copy is available for restoration. Mobility of snaps to designated destinations, is WAN optimized (deduplicated and compressed). OmniCube snapshots – deduplicated, compressed backup copies – can be retained on disk inside of a highly available OmniCube without practical limit to number of snaps or period of retention. The media server and dedicated backup network may well vanish.
- **Distributed data availability and access** – Finally, OmniCube Global Federations built on top of highly efficient data storage and WAN optimization enable a unique degree of mobility and access, which also becomes a foundation for availability protection. Since VM data is highly deduped at rest as well as in-flight, VM mobility is greatly enhanced and workloads can rapidly move between distributed OmniCubes. This in turn enables unique disaster protection and scaling

alternatives. As data and VMs are constantly visible from any node, all nodes are globally protected, and any data can be rapidly streamed to a site on-demand if it is not already available, applications can easily be run anywhere, and single site failures can be easily avoided. Standard vSphere mechanisms like VM affinity rules can keep application components together, while distributed nodes still appear as a single unified cluster. When it comes to scaling, customers with restricted floor space can easily join the resources of multiple locations together into one OmniCube Global Federation infrastructure to grow their infrastructure beyond their footprint. Such locations can include the public cloud, or other data center locations.

## TANEJA GROUP OPINION

SimpliVity's hyperconvergence stands poised to tackle the most complex and costly problems in the data center. Today, scaling, visibility, and integration and orchestration of multiple systems across multiple IT domains drive infrastructure OPEX costs far beyond what most customers expect at the onset of any virtualization initiative. Such costs spiral further out of control over time, and can thwart customer attempts to scale the virtual infrastructure, or simply mean the infrastructure is vastly more expensive than ever anticipated. Raw complexity many times exacerbates OPEX costs, while in other cases the storage foundation itself is the biggest obstacle, and necessitates sprawl, increasingly scattered data, artificial limits between systems in a single infrastructure, and consequent increases in management overhead.

In our opinion, tackling these challenges requires an approach from a much different direction. We fundamentally believe that tackling these challenges with external systems will take an approach built on storage, which is the domain from which the most problems are spawned. This is the common approach we see from emerging hyperconvergence vendors.

Hyperconvergence itself is not really so radical – it is simply a better-realized iteration of the vision that every major vendor is pursuing. Those major vendors hope to extend their management approach to control and automate everything in the data center, under a banner of “software defined” IT. They frankly have a long way to go, and we have seen the industry pursue this goal many times before. Hyperconvergence turns this model on its head, and instead starts by building its foundation on a scalable, storage-layer-glue integrated with compute that in turn makes all of the physically imposed boundaries and complexity disappear.

SimpliVity is showing just what storage-driven hyperconvergence innovation looks like, and the stage is set for a demonstration of just how far SimpliVity can go in dispensing with complexity and bottlenecks, while also leveraging storage innovation to take IO and VM density to the next level. Based on the rapid progression of virtualization today, and the remaining challenges, we are convinced that hyperconvergence solutions will resonate with customers. The next wave is here.

---

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.