

Comparing Hyperconverged Infrastructure Options **for Virtualized Environments**

BUSINESS COMPARISON

Hyperconverged infrastructure is radically shaking up the IT landscape, creating huge operational and economic benefits. Tier 1 applications such as Exchange, SQL Server, Oracle and others are among the many beneficiaries of this new generation of infrastructure. However, there are many vendors jumping on the market bandwagon, and not all systems that are marketed as hyperconverged really fit the criteria. IT organizations need to do their homework to ensure they are selecting true hyperconverged solutions.



For more than a decade, data center infrastructure has undergone a powerful transformation due to virtualization. Legacy infrastructure silos, typically built around dedicated servers with direct-attached storage and allocated to specific applications, proved inefficient and underutilized.

Instead, virtualization allowed organizations to create more flexible and more cost-efficient infrastructure in data centers, departmental settings, remote/branch offices and small businesses. But for many organizations that adopted virtualization, new infrastructure often amounted to little more than shared storage.

The reality is that the existing legacy infrastructure wasn't built for the virtual world. Traditional storage and data protection applications and appliances were built for the logical unit number (LUN) or the volume, necessitating an inefficient translation layer to the VMs and applications that the businesses depend upon.

In recent years, there have been important and impressive technical advances in converged infrastructure, to the point where infrastructure solutions now are more scalable, offer greater functionality, deliver greater cost efficiency and are far easier to manage. This new generation of converged solutions—called hyperconverged infrastructure—is changing the rules of the game for IT organizations.

There is a lot of hyperbole and attention around hyperconverged infrastructure, so it's important to establish common terminology and definitions. Hyperconvergence replaces the need for disparate hardware components in a legacy stack with a software-centric architecture that includes compute, storage, networking, hypervisor and unified management in a single box. Some vendors in the hyperconvergence space also integrate other functionality such as built-in data protection.

Because hyperconvergence combines the economic benefits driven by technologies such as virtualization and cloud computing with a software-centric focus for adding

functionality, adoption of the technology has surged. The global hyperconverged market will experience a compound annual growth rate of more 70% between 2013 and 2018, according to a report from Technology Business Research.¹ IDC predicts global sales of hyperconverged infrastructure will more than double in 2015, exceeding \$800 million.²

This paper will help you not only understand the benefits of hyperconverged infrastructure, but also will give you practical, actionable advice on what to look for in a hyperconverged solution. Since there are numerous suppliers targeting the market with products that offer different benefits and tradeoffs, this paper will examine three different architectural approaches to converged infrastructure. It also will look at specific products from NetApp, Nutanix and SimpliVity, and will examine how those vendors' products stack up in order to help you make smarter buying decisions.

Hyperconvergence: How We Got to Where We Are Today

Why is hyperconvergence generating so much interest and activity in IT organizations? And why are more and more organizations increasingly evaluating and purchasing hyperconverged solutions? In order to answer those questions, let's take a brief look at data center infrastructure's evolution from legacy hardware silos to converged and, ultimately, hyperconverged infrastructure.

Traditional infrastructure stacks

For decades, data centers utilized large, monolithic servers attached to expensive SAN-attached storage. They took up significant physical space, consumed large and growing amounts of energy and cooling resources, required dedicated, manual IT management and chewed up large capital equipment budgets. As data volumes grew and workloads expanded, IT organizations desperately sought newer, more economical and more efficient solutions. Traditional stacks also demanded significant

1 "Hyperconverged Platforms Market Landscape," Technology Business Research, April 2015

2 "Workload and Geographic Expansion Will Push Hyperconverged Systems Revenue Past \$800 Million," IDC, April 2015

IT commitment to specialized skill sets, such as storage and networking, in order to properly manage the environment and deal with infrastructure complexity.

First-generation converged infrastructure

The development of reference architectures (RAs) for integrating hardware components such as servers and storage was a big step forward. By using a third-party blueprint for integrating hardware components, organizations were able to follow a path that led to more efficient infrastructure. These initial RAs did not typically cover all components, weren't optimized to use software as the basis for innovation, didn't offer holistic management of the entire infrastructure, and required data protection to be bolted-on separately. However, this first iteration of converged infrastructure did help reduce deployment risk and improve implementation time.

Second-generation converged infrastructure

Moving from RAs to a more tightly integrated hardware stack that combined compute and storage components was another important improvement. It allowed organizations to begin using software building blocks to run on x86-based hardware, separating the hardware and software layers to improve performance and ease management. The downside, however, was that these solutions focused primarily on the integration of the compute and storage layer only. Notably, second-generation converged infrastructure typically included limited data protection, which meant organizations often had to purchase dedicated appliances or additional software for data protection, adding cost and management complexity.

The brave new world of hyperconvergence

Hyperconvergence, a true consolidation of all IT below the hypervisor, has evolved as the next level of convergence. Such solutions deliver optimized management and performance, and some also support built-in data protection. Simply combining the server and storage layers didn't do enough to make infrastructure more affordable from either a Capex or Opex perspective, nor did it sufficiently improve performance, resiliency and data protection.

In this new era of hyperconvergence, several essential concepts emerged as guidelines for optimized solutions, each rooted primarily in software as the foundation for dramatic, sustainable economic and operational benefits:

Data efficiency: The amount of data that needs to be captured, stored and backed up continues to be an ongoing challenge for IT organizations. Although deduplication technologies have been adopted, these introduced complexity by adding yet another infrastructure element that needs to be managed in the data center and typically handled deduplication as a post-process function. As a result, IT decision makers have looked for new ways to optimize data efficiency.

Built-in data protection: Not only have storage capacities surged, but demanding retention policies now require more vigilant and reliable data protection. Instead of boosting Capex budgets to pay for dedicated data backup and recovery products, hyperconverged solutions take a VM-centric approach to data protection and build it into the solution. This greatly simplifies replication, backup, restore and recovery, often reducing backup windows from hours or even days to minutes.

VM-centric management: Virtualization has given organizations the opportunity to view and treat their systems as a collective, integrated pool of resources, rather than as a mix of discrete components. Being able to manage in a VM-centric way relieves administrators from much of the burden of configuration complexities, monitoring system operations and optimizing performance, while enhancing VM mobility. By using a centralized management platform with a single administration console and common application programming interfaces, infrastructure administrators can spend less time and effort manually managing and tuning infrastructure, relying instead on software and automated tools to handle most of the complexities of infrastructure.

Accomplishing these capabilities requires solutions that go far beyond integration of the server and storage tiers. It necessitates a flexible, scalable, automated software platform to leverage virtualization in order to optimize all data-related functions.

The combination of tight hardware component integration and a virtualization-centric architecture makes hyperconverged solutions an ideal fit for a wide variety of use cases. These include mission-critical enterprise applications, data center consolidation initiatives, private clouds, virtual desktop infrastructure (VDI), remote office/branch office and data protection modernization projects.

Considering Your Options: NetApp FlexPod, Nutanix XCP and SimpliVity OmniStack

As converged and, increasingly, hyperconverged solutions are sought out by IT decision makers, the good news is that there are many purchasing options. These solutions come from a combination of established suppliers and agile, innovative, younger companies, some of which are pioneering new ways to design, deploy and manage IT infrastructure through hyperconvergence.

Many organizations are likely to consider solutions from three reputable suppliers: NetApp, Nutanix and SimpliVity. Each brings a different approach to converged infrastructure, providing a unique value proposition for the customer. It is essential for readers of this paper to understand the unique approaches taken by these three companies, and what those differences mean for IT buyers and their organizations.

Supplier and Solution Overviews

NetApp

Founded in 1992, NetApp is the oldest and largest supplier of the three options evaluated in this paper. NetApp's strength is storage; it pioneered an important storage solution that eventually came to be known as network-attached storage (NAS) appliances.

As a result, NetApp's biggest technical contribution to its FlexPod series of converged infrastructure solutions is storage. FlexPod is the result of collaboration between NetApp and Cisco. FlexPod components include Cisco Unified Computing System (UCS) servers, Cisco Nexus switches and NetApp

storage systems. FlexPod reflects the first-generation, reference architecture approach to hyperconvergence.

There are three different types of FlexPod solutions marketed by NetApp:

- FlexPod Datacenter, a validated solution that combines storage, networking and server hardware into a unified architecture for enterprise workloads. These solutions are based on NetApp's hard disk-based and flash-based scale-out storage, its clustered Data ONTAP and MetroCluster software and Cisco's servers and networking switches.
- FlexPod Express, a lower-cost solution designed for small or midsize businesses.
- FlexPod Select, which is optimized for specific workloads such as big data, high-performance computing, databases and data warehouses.

Due to their time in the market and the shift in buying preferences for converged systems, NetApp FlexPod has been deployed widely from small to midsize organizations for a variety of use cases, ranging from private clouds to virtualized desktops. FlexPods have often been deployed as an alternative to build-it-yourself options in order to accelerate time to running business applications and reduce integration challenges.

Nutanix

Nutanix was founded in 2009 with a goal to converge compute, storage and virtualization into a single converged solution. Nutanix's solutions are examples of second-generation hyperconvergence. These go beyond the reference architecture model, but do not fully address data efficiency and comprehensive built-in data protection.

Three versions of the Nutanix Xtreme Computing Platform are supported. The Starter edition provides basic software functionality, ideally for small-scale deployments with limited workloads. The Pro edition provides higher resilience and more extensive management features, and is aimed at enterprise customers. The Ultimate edition provides the broadest array of

functionality, typically designed for multi-site deployments and requiring higher levels of security.

Nutanix's solutions based on Supermicro servers are sold by the company's own sales force and its channel partners. Their product also runs on Dell servers and is sold through the Dell channel. Nutanix markets these solutions as the NX Series of Xtreme Computing Platforms; they range from the low-end NX-1000 series up to the high-end NX-9000 series.

Although Nutanix positions its solutions, especially in the Pro and Ultimate editions, for a wide range of workloads, the company is best known for its virtual desktop infrastructure (VDI) solutions.

SimpliVity

SimpliVity's focus in hyperconvergence is on simplifying IT through innovative, software-based solutions that run on a wide range of industry-standard hardware. Since its founding in 2009, the company has focused exclusively on bringing hyperconverged solutions to a new level in order to dramatically improve the economic and operational equation for hyperconverged infrastructure.

SimpliVity's hyperconverged infrastructure solution runs on enterprise-grade white-labeled servers under the SimpliVity OmniCube brand name, or on third-party servers under the SimpliVity OmniStack name. At the time of publication, supported third-party servers included Cisco UCS. Along with converging the server and primary storage layers, the solution also converges deduplication appliances, backup software, replication, WAN optimization technologies and cloud gateways.

What makes SimpliVity's approach unique relative to other vendors is the investment it has made in the area of data efficiency. The OmniStack Data Accelerator Card performs inline deduplication, compression and data optimization at ingest across primary and backup storage, offloading this processing so VMs suffer no performance penalty. As reported on their web

site, median data efficiency is 40:1.³ Another interesting data point is that customers achieve greater data efficiency while improving application performance.

SimpliVity's core technology, called the OmniStack Data Virtualization Platform, is the brain of SimpliVity's solution. It runs as a virtual controller on vSphere ESXi, and abstracts data from the underlying hardware. Policy management is handled at the VM level and promotes VM mobility. Unlike alternative solutions that require learning another management interface, the approach that SimpliVity has taken is to integrate with existing management consoles such as VMware vCenter, and management and orchestration software from a variety of vendors.

OmniCube and OmniStack models come in several different configurations with different levels of storage, compute and network connectivity, and are designed for a wide range of use cases. SimpliVity has had the most success in deployments supporting mission-critical applications and remote office/branch offices, along with VDI. SimpliVity solutions offer enterprises a flexible approach to deploying hyperconvergence: They can be deployed starting as a single node within a data center, with easy expansion and scale as an organization's needs warrant.

Evaluating Suppliers and Making the Best Selection

NetApp, Nutanix and SimpliVity each take very different paths to hyperconverged infrastructure. While they certainly compete against one another for business opportunities, in many ways they represent unique and almost mutually exclusive viewpoints on how to deliver convergence.

Consider **FlexPod from NetApp**. This product is certainly a solid converged solution for customers moving from traditional, disparate infrastructure stacks to a more integrated solution. NetApp is a long-established, well-resourced supplier with a reputation for good products and technology.

3 ["SimpliVity Adds Support for Cisco UCS Director, Simplifies IT and Cloud Environments," SimpliVity, June 8, 2015](#)

NetApp's traditional strength in storage has allowed it to offer a product that does a good job pre-configuring storage, and its partnership with Cisco brings computing and networking connectivity into the solution. The FlexPod solution is well regarded for some of its data protection capabilities, including FlexClones, SISclones, Snapshot and SnapMirror, and has good solutions for Microsoft workloads such as SQL Server and Exchange.

The FlexPod approach acts as a blueprint, in the form of a reference architecture, to allow customers to follow a prescribed path and deploy components faster than they could reasonably do on their own.

But the FlexPod approach doesn't go as far as other solutions in providing the full array of hyperconvergence benefits in terms of ease of use, VM-centric global unified management, data efficiency and data protection. While FlexPod does a good job with ease of installation because of its blueprint pathway, it falls short on ease of management. The compute, storage, networking and management software require assembly, and it typically takes a good amount of IT administration time and effort to tune the components for optimized performance.

This also makes scalability and upgrades more challenging because of the component assembly and integration required to add more storage, servers and/or network bandwidth. This approach also taxes staff resources when it comes to support, since "turning on" additional features while updating components may result in new configurations that fall outside of the standard support matrix.

Other challenges with FlexPod include high administrative overhead that makes data compression much more challenging than organizations will want; local deduplication rather than global deduplication; and the inability to manage multiple FlexPods as a shared resource tool. FlexPod consists of many disparate technologies that have been cabled together to provide an infrastructure for virtualization. It does not provide VM-centric management and uses legacy storage protection methodologies that bring a lot of overhead to the system.

Resilience, which is undoubtedly high on the priority list for infrastructure managers, is an area where NetApp's approach excels, but has important limitations. Leveraging mainly RAID-DP (a RAID 6-like parity solution) and a single file system structure for storing data, the solution can potentially fall prey to data loss due to file system-level corruption and triple disk failure. This exceeds the resiliency found by default on most convergence 2.0 suppliers (which lack RAID), but falls short when compared to convergence 3.0 providers. Snapshots are also leveraged to produce point-in-time recovery, but can be costly to maintain. And, if corruption occurs in a parent snap, you in essence end up saving the fruit of the poisoned tree as children snaps will likely carry forward the corrupted data.

Compared to NetApp FlexPod, **Nutanix-based converged solutions** go further in providing a software-centric model to hyperconvergence. Nutanix's solutions are most highly regarded for VDI workloads, which is where the company first gained its foothold in the market shortly after it was launched.

The Nutanix solution does a good job of creating a single pool of shared resources across compute and primary storage, and supports multiple hypervisors.

But Nutanix solutions also lack certain key capabilities typically sought after by IT organizations. In the area of data efficiency, fingerprinting of data is done inline for sequential writes of 64 kbytes or larger, but the actual deduplication processing is largely done post-process. As a result, it cannot match the data efficiency of alternative solutions.

Also, costs for enterprise-grade configurations can rise significantly when adding compression, post-process capacity tier deduplication, multi-site replication, RF3 resiliency or centralized management, since they require add-on licenses for these features.

In data protection, there are several key shortcomings that could impact operations for IT teams. For instance, the Nutanix solutions only support snapshots; backup protection requires a third-party solution. Also, multi-site replication—an important

requirement for enterprise-class deployments—requires an upgrade to the Ultimate edition. Customers evaluating Nutanix should consider whether the solution's remote replication meets their RPO/RTO requirements and understand how much network bandwidth over the WAN will be required.

It's also important to consider some of the challenges with Nutanix's philosophy for resilience. The Nutanix solution offers no RAID protection intra-node. Instead, Nutanix positions Resiliency Factor 2 (RF2) as a default; however, if you lose any two disks or a single disk while a node is offline in a cluster due to a failure or maintenance, data will be lost. For customers requiring higher resilience, Nutanix offers RF3, which requires a minimum of five nodes at each site, an additional 50% physical capacity, and moves 50% more traffic over the network to store the same amount of data.

By comparison, **solutions from SimpliVity** meet—and often exceed—requirements for the newest generation of hyperconverged solutions. This is particularly the case in the three main categories of capabilities that define and differentiate true hyperconvergence: built-in data protection, VM-centric management and data efficiency. Solutions that do not adequately support critical capabilities in each of these three segments fall short of next-generation hyperconvergence, resulting in a failure to leverage the full array of benefits of the technology.

One of the biggest advantages in the SimpliVity approach is the utilization of a familiar and widely adopted tool—VMware's vCenter management console. This obviates the need for organizations to support multiple management tools and speeds time to value by leveraging IT staffs' long-term commitment to, and familiarity with, vCenter. All operations are performed from a VM-centric mindset, including setting up local and remote backups and clones at the VM level.

SimpliVity also strives for ways to drive higher data efficiency. SimpliVity's OmniStack Data Virtualization Platform delivers inline deduplication, compression and optimization at ingest to dramatically improve data efficiency and application performance, even in highly scalable environments. The

OmniStack Accelerator Card is a PCIe card that offloads the processing of data optimization to the workloads running on system VMs with no performance impact to production workloads.

Another key capability in the SimpliVity hyperconverged solutions set is built-in data protection. This is a set of capabilities that has become increasingly important in the face of stricter compliance and governance policies; it also enables improved performance. When SimpliVity creates a new local backup, IOPS are not consumed since the creation of the backup is simply a metadata update. This is important when backing up remote VMs, as only the unique blocks are sent; this saves WAN bandwidth, remote processing and IOPS.

Built-in data protection is a critical requirement for customers evaluating hyperconverged solutions. Earlier-generation converged infrastructure designs required the addition of data protection and backup appliances, resulting in higher capital expenses and more hardware and software to manage.

When it comes to resilience, SimpliVity takes a very different approach than both NetApp and Nutanix. SimpliVity configures its nodes for high availability and commits all writes to the OmniStack Accelerator Card on two different nodes before acknowledging the write. If a node fails, all traffic is routed to the second node without disruption. Even with two drive failures on every node and a node offline, SimpliVity solutions still can serve data.

In the end, SimpliVity stands out in this comparison. Unlike alternative solutions that are based on a reference architecture approach or converge only storage plus compute, SimpliVity's approach goes further by delivering inline deduplication, compression and optimization globally across primary and backup storage tiers. Its VM-centric architecture and integration with existing management tools such as vCenter are unique in the ability to reduce complexity and ease deployment. Additionally, SimpliVity's solution is architected for enterprise-grade resilience and built-in data protection with near zero overhead.

	NetApp FlexPod	Nutanix XCP	SimpliVity OmniStack
<i>Data Efficiency</i>	Deduplication is not inline. Compression and deduplication are available but recommended to run off-peak to sustain performance.	Fingerprinting of data is done inline, for sequential writes of 64KB or larger. The actual deduplication processing is largely done post-process.	All data is deduplicated compressed and optimized inline globally across all tiers once and forever, globally in 4 to 8KB chunks. Median customer data efficiency is 40:1.
<i>VM-Centric Management</i>	Management paradigm is at the LUN level. iSCSI/ Fibre channel networking, LUN mapping and zoning are part of the standard mode of operation. Unified management requires 3rd party software solution with its own infrastructure, software and licensing cost.	VM management has an added cost due to additional licensing and requires multiple interfaces including Nutanix Prism, Prism Central and the individual hypervisor management consoles, making movement of VMs between data centers a challenge.	VM management is provided via integration with VMware vSphere and other management and orchestration software. All management is at the VM level without the complexity of LUNs and SAN concepts.
<i>Data Protection</i>	No built-in backup. Backup requires a 3rd party backup software with its own infrastructure, software and licensing cost. NetApp snaps and SnapMirror do provide local and remote data protection.	No built-in backup. Backup requires 3rd party software. Nutanix does natively offer snap shots and multi-site replication for additional license cost. File level restore requires 3rd party software.	Built-in VM backup, multi-site replication, recovery and cloning, and disaster recovery included natively. SimpliVity also includes file level restore natively.
<i>Resiliency</i>	Double-parity RAID-DP prevents data loss with double drive failure.	No RAID, resiliency is based on RAIN with Resiliency Factor (RF) 2 set by default. RF2 only protects against single drive loss or single node loss. Node loss plus an additional drive loss results in data corruption. RF3 is available, requiring significantly more infrastructure and cost investment.	Intra-node RAID6 can tolerate double HDD failure on every node. Multiple copies of data are spread evenly over several nodes for additional resiliency.

Conclusion

Data center infrastructure will never be the same—and that's a huge benefit to IT departments looking to squeeze more value, productivity, performance and reliability out of their infrastructure. As legacy hardware stacks gave way to the initial waves of converged infrastructure with integrated compute and storage, the value proposition improved significantly.

But IT decision makers, clamoring for even better efficiency with an eye toward the future, clearly are excited about the improvements brought on by the new wave of hyperconverged infrastructure. By leveraging the increased use of virtualization and software-defined architectures, hyperconverged infrastructure is delivering upgrades in performance, cost efficiency and manageability that are leaps and bounds better than had been achievable just a few years ago.

As this paper has pointed out, there are numerous viable suppliers of converged solutions. Some, such as NetApp and Nutanix, offer solid products that are focused on specific

applications or workloads, and have helped to build credibility for this category of solution. However, IT professionals are looking for even more in a hyperconverged solution.

This is where SimpliVity's solutions stand apart. SimpliVity's global inline deduplication, built-in data protection and global unified management offers better cost efficiency, increased resilience, and more simplified management than solutions from NetApp or Nutanix.

Whether you are modernizing or overhauling an enterprise data center, or implementing a greenfield deployment in a department, remote office or small business, SimpliVity's solutions deliver the enterprise-class performance, protection and resiliency that today's organizations require, with significant TCO savings compared to legacy approaches.

For more information on SimpliVity's approach to hyperconverged infrastructure solutions, go to www.simplivity.com/products.