



Outlook Report

Hyper-converged Infrastructure: An SSG-NOW Outlook Report

February 27, 2015

By Deni Connor and Earl Follis
Client Relations: Phylis Bockelman

SSG-NOW
8815 Mountain Path Circle
Austin, Texas 78759
(512) 345-3850

Note: The information and recommendations made by SSG-NOW are based upon public information and sources and may also include personal opinions both of SSG-NOW and others, all of which we believe are accurate and reliable. As market conditions change however 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. SSG-NOW 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 which may appear in this document.

This report is purchased by SSG-NOW, who understands and agrees that the report is furnished solely for its distribution to customers and prospects, without the prior written consent of SSG-NOW.



Storage/Systems Strategies NOW | ssg-now.com
512.345.3850 | info@ssg-now.com



Table of Contents

Table of Contents	2
Editor's Note.....	5
Executive Summary	6
Driving Issues, Trends, History.....	7
The history of Hyper-Converged Systems (HCS).....	7
How HCS differs from converged systems	7
Benefits of converged compute and storage workloads on single hardware platform	7
Private cloud and hybrid cloud use.....	8
Remote and branch office integration	8
Data Center Consolidation	8
Enablement of analytics	8
Hyper-converged Infrastructure Features.....	9
Hyper-Convergence on Industry-Standard x86 Platforms	9
Software-defined.....	9
Hypervisor Support.....	10
Automation and Analytics	11
Maintaining State Information	11
Private and Public Cloud Integration	12
Resilience and Supportability	12
Centralized systems and network management.....	12
Support of Dissimilar Hardware Configurations	13
Distributed Systems Integration	13
Use Cases for Hyper-converged Systems.....	14
Postponing Capital Investment.....	14
Incremental growth	14
Public and Private Cloud	14
Specific Use Cases	15
Best Practices for the Implementation of Hyper-converged Systems.....	16
Vendor Name: Gridstore.....	17
All-Flash HyperConverged Appliance	17
Storage Nodes	17



Outlook Report

Windows Servers:.....	17
Server-Side Virtual Controller Technology.....	17
Gridstore Software	18
Gridstore, the company	19
Use Case Profile	20
SSG-NOW's Take on Hyper-Converged Systems	22



Outlook Report

Sponsor





Outlook Report

Editor's Note

This report on Hyper-converged Infrastructure was undertaken because there was so much noise in the industry as to what hyper-converged infrastructures are. Industry analysts, vendors and end-user IT professionals all have their own definitions. For some, Hyper-converged Infrastructures are just an extension of Integrated Infrastructure Systems. For others, it has everything to do with Web-scale computing practiced by the Amazon's and Google's of this world. Still for others, it is a type of computing system envisioned by VMware in its recently announced EVO: RAIL architecture. Those are a bunch of words that say a lot, but which provide little guidance to an IT professional or CIO that is faced with crushing system and storage challenges in implementing a hyper-converged infrastructure.

SSG-NOW is pleased to present this report on our definition of hyper-converged infrastructures. Read it with interest and if you have any questions, please contact us at the addresses below.

Deni Connor
Founding Analyst
dconnor@ssg-now.com

Earl Follis
Senior Analyst
efollis@ssg-now.com



Outlook Report

Executive Summary

Hyper-converged systems first appeared in the data centers of the world's largest Web-scale companies because organizations such as Facebook and Google found that traditional storage architectures could not scale efficiently and required high, recurring capital investments and increasing operating expenses. In addition, in order to support the rapid expansion of infrastructure associated with Web-scale services, it was not efficient to relegate servers to dedicated roles, but far better to have the same physical server capable of running both compute and storage workloads. Further, a clustered architecture was deployed in order to maximize availability and resiliency.

The Web-scale companies recognized the importance of these features and spent the resources on the required development of proprietary software as a core competency. Mimicking these innovations, a new crop of commercial hyper-converged systems have begun to enter the market with the promise of providing mainstream IT the same advantages of scalability and low incremental growth costs, with total cost of ownership less than more rigid, conventional architectures of dedicated compute farms and proprietary shared storage systems.

SSG-NOW was an early researcher in this new field and has initiated coverage of converged compute and storage functions as a high-growth product category. By using industry-standard x86 hardware platforms that include direct-attached storage rather than purpose-built and proprietary storage systems, capital expenditures and operating expenses can be minimized. The requirements for the software to operate compute workloads and storage on the same hardware include functions that are integrated with virtualization hypervisors and operating systems, while supporting both solid state and rotating storage media.

Organizations of all sizes have benefited from hyper-converged systems as a solution to sprawling server farms, expensive and poor scalability of shared storage systems and a seamless method to integrate multiple sites and hybrid cloud uses.

Driving Issues, Trends, History

The hyper-converged systems of today are a combination of years of development and refinement of virtualization software and similar advances in hardware. The popularity of hyper-converged systems can also be attributed to advances in scalable, modular computing hardware and software concepts pioneered by Google, Facebook and other Web-scale companies. Similar advancements in software-defined computing also contribute to the striking potential of hyper-converged systems to fundamentally change the way we build-out, operate and maintain modern enterprise data centers.

The history of Hyper-Converged Systems (HCS)

The roots of hyper-converged computing are based in advances in virtualization and the physical shrinking of data center hardware. Converged systems, such as the HP ConvergedSystem and IBM PureSystem, for high-density data centers was the first step in the HCS odyssey. Web-scale companies then took that concept several steps further by designing, building and deploying hyper-converged systems in very dense deployments. Google was a pioneer in the area of their modular data centers built from shipping containers containing racks of hyper-converged systems. This idea carries the modular characteristics of HCS to the next level, as those shipping containers can also be stacked and quickly connected to a data center network, creating large-scale modular computing environments based on small-scale modular computing systems.

How HCS differs from converged systems

Hyper-converged systems grew out of the market movement for converged systems, which features individual components specifically designed to interoperate seamlessly. HCS takes that idea to the next level by combining those individual components into a single chassis, allowing for even tighter integration between the storage, compute, network and management software. The results is a modular data center building block, with each hyper-converged system able to be deployed and configured very quickly. Some HCS vendors offer to load disk images and pre-configure the hyper-converged system so that it can literally be slid into a rack, plugged into power and network outlets, turned on and be operating very quickly. These are the primary advantages that HCS has over converged systems: speed of deployment, ease of installation and rapid utility.

Benefits of converged compute and storage workloads on single hardware platform

The primary benefits of HCS reside in the ease and speed of deployment, the modular nature of HCS and the tight integration of converged components offered by HCS right out of the box. This concept reduces the need to spend time designing effective, compatible, yet separate data center systems. If you need more compute or storage capacity, you simply deploy additional HCS compute or storage nodes as needed. Because all of the hardware and software has been designed from the ground up to work in unison and all come from a single vendor, support costs are notably lower, as are the opportunities for vendors to point the finger at each other when issues arise. With HCS, you have a single vendor neck to wring if problems crop up.



Outlook Report

Private cloud and hybrid cloud use

A driver for the wide-spread acceptance of the hyper-converged architecture is the ability for companies to easily and quickly utilize HCS for private cloud and hybrid cloud deployments. As many enterprise companies are just beginning to experiment with private cloud and hybrid cloud data centers, HCS is a natural fit in those environments. Considering that most public cloud companies are using HCS as their building blocks, enterprises looking to replicate that public cloud functionality in a private or hybrid makes perfect sense.

Remote and branch office integration

Another potential sweet-spot for HCS is in remote and branch office installations. Many companies find it advantageous to deploy certain server components to physically remote locations to reduce wide-area network expenses while also offering better application performance to end-users at remote offices. Email server, database servers and highly-utilized applications can all benefit from locating those computing resource as close as possible to end-users. Hyper-converged systems fit the bill perfectly due to the ease at which those systems can be deployed and supported.

Data Center Consolidation

Because of the high density of large HCS deployments, special attention must be given to power, cooling and networking in data centers utilizing HCS. The miniaturization of HCS hardware components means that racks of hyper-converged systems are extremely dense and may require increased cooling, high-amperage electrical circuits and high-speed networking within and between racks. Also, remember that increasing electrical circuit density will have an impact on backup generators and fail-over cooling strategies employed by that data center.

Enablement of analytics

One of the challenges in systems management and monitoring has always been the aggregation and correlation of key metrics collected from a multitude of underlying devices, network links and storage systems. HCS offers a plethora of tightly integrated data from the various components contained in the HCS chassis. Also, note that management software is a key component of HCS, giving HCS vendors the ability to deliver comprehensive management and monitoring capabilities that can contribute valuable data to be used in analytics and predictive analysis.

Hyper-converged Infrastructure Features

There are many features that are requisite of hyper-converged systems. Among them are being based on industry-standard x86 platforms, hypervisor-dependence or hypervisor-agnostic, distributed architecture, centralized management, both storage and compute capabilities in a single system, and enhanced scalability.

Hyper-Convergence on Industry-Standard x86 Platforms

A key concept in delivering high-performance data centers is the best practice of locating storage in close proximity to actual application workloads. Hyper-converged systems achieve that performance efficiency by integrating scalable compute and storage capabilities in the same physical chassis. (Note that there are industry-standard x86 platforms on the market that do not qualify as hyper-converged systems because they do not contain storage and compute capabilities in a single chassis.) In the example of clustered, hyper-converged servers, all direct attached storage nodes within a server cluster are aggregated together into a shared pool of storage that can be accessed by any server within that cluster. This design keeps latency and complexity to a minimum, compared to traditional hardware-based shared storage schemes, while still supporting virtualization features that require shared storage. The hyper-converged architecture seamlessly supports both hard disk drives and Flash drives that can be managed and monitored via virtualized storage controller. This controller directly ties shared storage resources to I/O requests generated by the hypervisor, thus eliminating the mismatch between compute and storage (often called 'the I/O blender effect'). By building-in robust error detection and automatic recovery resources, hyper-converged computing provides a high availability, fault-tolerant computing platform. These techniques are critical to the performance and resiliency of the virtualized storage subsystem that support highly available, hyper-converged computing.

Software-defined

The concept of software-defined computing -- including networking, storage, and compute -- is key to effective management of large-scale hyper-converged data centers. Software-defined computing utilizes inexpensive, industry-standard hardware -- servers, network devices and storage systems -- rather than dedicated, proprietary hardware on which servers, networking and storage technology has always relied. By logically detaching the software from underlying hardware, hyper-converged systems are able to evolve and support new hardware components far faster and cheaper than traditional proprietary computing hardware. Software-defined computing gives vendors and customers the ability to quickly take advantage of ongoing developments in memory, storage and processor upgrades. Because of software-defined computing, customers and vendors no longer have to wait for an OEM firmware update or an update to device operating systems in order to take advantage of new hardware capabilities. Traditional hardware-defined solutions rely on dedicated



Outlook Report

firmware and software that must be thoroughly engineered and tested prior to releasing new features and capabilities, a process that can take months or even years in some cases.

Hypervisor Support

There are two types of hyper-converged systems: those that are hypervisor-agnostic, i.e., they can operate on two or more different hypervisor platforms, and those that are hypervisor-specific, i.e., they run on only one or two of the major hypervisor platforms. At first glance, it might appear that a hypervisor-agnostic hyper-converged system has an advantage because it runs on multiple of the popular hypervisors platforms. And, certainly, if you are implementing a hyper-converged system in an environment that runs two or more hypervisors, there may be a justification for choosing a hypervisor-agnostic system. However, there are advantages to choosing a hypervisor-specific hyper-converged system, including tighter integration with the underlying hypervisor, the possibility of increased performance -- the extra code required to run on multiple hypervisor platforms may adversely affect performance -- plus a tighter integration with and increased functionality from within the underlying hypervisor management and administrative tools. There is no one correct answer for choosing between hypervisor-agnostic and hypervisor-specific hyper-converged systems. There are advantages and (minor) disadvantages to each approach and the final decision should be weighed carefully prior to purchasing and implementing a hyper-converged solution.

Data Placement, Latency and Intelligent Tiering

As mentioned above, keeping your application data in close physical proximity to the application server itself is a best practice for hyper-converged environments. This approach minimizes latency and keeps network traffic to a minimum compared to traditional compute and storage architectures. Gone are days of hugely expensive dedicated tiers of servers and physical storage arrays. However, the concept of keeping data as close as possible to the application is complicated by the fact that virtual machines (VMs) can and will migrate from one physical hypervisor host to another. When VMs move, the underlying application data must also move to the new hypervisor host. Failure to do so can have a significant impact on storage -- and thus application -- performance as data must traverse more of the data center network as needed by the application. The type of storage media used can also improve latency. Flash storage devices offer much better performance than do traditional hard disk drives. In high-performance architectures, Flash storage should be reserved for the most frequently used data, while less time-critical data can be relegated to slower storage media.

Linear Scalability

Scalability is also a critical consideration for hyper-converged environments, as application data utilization tends to increase over time. Traditional hardware-based storage systems require ample budget to procure and sufficient time to deploy well in advance of the application storage use curve. Building out your hardware-based storage today so that you'll be ready to handle increased storage utilization in the future is an expensive and wasteful approach, as you are frequently

provisioning unused storage capacity months or even years before that capacity will be put to use. Not only is this approach expensive and wasteful, but it also limits data elasticity as application and business requirements change over time. The ability to scale storage capacity up or down at the hyper-converged server level provides a lower-cost, elastic model for rapidly adjusting storage capacity as needed. The expansion or retraction of storage limits over-provisioning or under-provisioning storage while reducing storage costs. With this fact in mind, customers are encouraged to carefully select a hyper-converged platform that does not limit storage scalability by server or storage nodes.

Automation and Analytics

The old saying that you can't manage what you can't measure is especially true in the world of hyper-convergence. As the number of nodes in a physical cluster increase, there is a corresponding increase in the complexity of managing and monitoring storage nodes. Increasing complexity implies a higher level of effort to effectively manage that physical storage. The hyper-converged, software-defined data center leverages a high level of automation to avoid the error-fraught manual activities related to administering storage systems. With the automation offered by hyper-converged systems, administrative tasks can be reduced to troubleshooting tactical issues and deploying new workloads. Automation reduces operating costs by requiring less human intervention. Automation capabilities can also leverage predictive analytics to proactively identify trends and demand fluctuations that might lead to resource contention or wasted computing resources. Predictive analysis supports the ability to address and correct impending issues before they can affect application or infrastructure performance and availability, thus reducing downtime and saving money.

Native Data Management

Traditional physical storage systems include sophisticated data management algorithms that ensure data protection and recovery capabilities. Typical data management features include taking data snapshots, cloning workloads and servers and transporting data to remote storage environments to ensure disaster recovery requirements. These data management techniques can be generated via automation policies or via a manual process. The downside of these data management techniques is that they can negatively impact network bandwidth utilization, while also creating duplicate copies of data sets or large blocks of unused storage space. Hyper-converged, virtualized systems can automate these day-to-day data management chores while presenting admins with a unified view of storage, whether local, remote, or virtualized. Hyper-converged systems can also automate the backup, archiving and disaster recovery process, so be sure that the hyper-converged platform you choose supports automation of data management.

Maintaining State Information

The design of massively scalable storage requires each node to manage a subset of the storage state information, no matter how many nodes reside in the storage cluster. This design goal requires storage clusters to distribute state



Outlook Report

information among all nodes, not storing it in a single, centralized state repository. This de-centralized state information is a key requirement of scalable hyper-converged systems, yet centralized state information is almost impossible to retrofit into legacy storage architectures. Centralized state information eventually creates so much inter-nodal traffic that it can saturate network links between nodes. Hyper-converged infrastructure solves this problem by continuously synchronizing state information among the nodes, so that the failure of one node has no lasting effect on the redundant state information residing on the other nodes.

Private and Public Cloud Integration

SSG-NOW conducts annual surveys to gather real-world data on utilization of cloud storage and hyper-converged storage systems. One survey reinforces the view that public cloud utilization is in high demand as an extension of data center infrastructure for many enterprises. Fortunately, most hyper-converged solutions support autonomous participation with public, private, and hybrid cloud scenarios. Seamless integration with cloud resources and the ability to manage workloads from a central location, regardless of where those workloads run, is a key competitive advantage of hyper-converged systems. The most frequently cited example is the use of public clouds for data backup, archiving and disaster recovery. Policy-based automation can take incremental snapshots at given intervals in order to support organizational recovery time and recovery point objectives.

Resilience and Supportability

Autonomous self-repair, fault tolerance, and supportability have been design goals of infrastructure components since the first cave-man suffered a hard drive crash with no viable recovery plan. Hardware and software vendors have worked for years to bring this idea to fruition, with varying results. Fault tolerance and resilience is critical because no hardware lasts forever. Hardware-based outages occur with regularity in the modern data center so having a plan in place to address critical failures is of paramount importance. For instance, the ongoing update process for firmware and software revisions must be optimized to minimize downtime. Ideally, a hyper-converged system should be able to download and install software updates with no impact or interruption to production applications. Compute and storage nodes should be able to be added or removed from the network and workloads redistributed without disruption. Be sure to consider resilience and supportability when evaluating hyper-converged systems.

Centralized systems and network management

Effective, intuitive systems, storage and network management capabilities can make the difference between smooth operations and a hobbled hyper-converged system implementation that stymies administrator's efforts to keep everything running as expected. Do not underestimate the need for effective, intuitive management capabilities in your hyper-converged environment. With that in mind, the hyper-converged management console is a key factor to evaluate when considering implementing a hyper-converged environment. The evolution of the skill sets required for the admins of hyper-converged systems dictates that ease-of-use is of utmost concern in hyper-converged management capabilities.



Outlook Report

Using examples of the iPhone interface or game console interfaces, that same idea of a clear-cut, concise management user interface is the goal. In contrast to management interfaces for legacy components, the hyper-converged management console should cater to the fact that the admin using that console may have a strong background in servers, or networking, or storage, but probably not in all three. Hyper-converged systems that are hypervisor-specific offer tighter integration with the underlying hypervisor management tools, such as the System Center Virtual Machine Manager management console for Microsoft Hyper-V environments and the vRealize Suite console for VMware vSphere environments.

Support of Dissimilar Hardware Configurations

In traditional data centers, there is typically a variety of hardware and software from multiple vendors that may or may not interoperate seamlessly. Examples include direct-attached storage (DAS), storage area networks (SANs) and network-attached storage (NAS). Your company may use, for example, a high-end SAN for business-critical workloads, direct-attached storage for test or dev environments and NAS storage for internal use. These disparate devices and systems often rely on different vendors whose products are customized to specific workloads and compute tasks. Of course, each vendor -- and sometimes even different hardware platforms from the same vendor -- require different management and support contracts, resulting in higher operational and capital costs. The flexibility of hyper-converged systems allows customers to replace some or all of those disparate storage systems and uses with industry-standard, software-defined storage regardless of the underlying hardware configurations, saving money and time.

Distributed Systems Integration

The advantages of distributed systems are numerous, including load balancing capabilities across clustered nodes, scalability and management features and the elimination of a single point of failure in the hyper-converged system architecture. This distributed architecture has been embraced by Web-scale providers such as Google, Facebook and Amazon that minimizes management complexity, while providing extensive architectural flexibility in the hyper-converged data center. Although it has taken years for this approach to become commonplace, now that the kinks have been worked out from this design, substantial benefits include automation, self-healing, dynamic data placement and decentralized storage state repositories. Hyper-converged systems offer automagic failover policies and controls that underpin system resiliency requirements, automated data placement practices and data management policies keep critical data as close as possible to the virtual workload that access that data. This design ensures peak performance while keeping hardware, software, support and administrative costs low.



Use Cases for Hyper-converged Systems

Hyper-converged systems are typically much cheaper to acquire and deploy than traditional centralized system designs--or even converged systems. Use cases and justifications for using hyper-converged systems abound as companies realize that they no longer have to spend hundreds of thousands or even millions of dollars up-front as a capital expense just because they've always done it that way. HCS is changing the way that IT budgets and implements the hardware and software required for large-scale application projects.

Postponing Capital Investment

If you are planning to deploy a business-critical application that requires highly-available clustered servers and a high-performance SAN for storage, your start-up costs for that project can be extremely high because it's usually more cost-effective to buy the servers and SAN fully configured for the potential maximum user load, rather than buying now and upgrading later. As a result, companies may have high upfront costs that must be budgeted for even though it may be months or years before all of that hardware is actually required or is put into use. HCS allows you to add capacity as you grow, avoiding large upfront capital expenses in favor of smaller, incremental purchases that can be treated as an operating expense (OpEx) rather than as a capital expense (CapEx).

Incremental growth

Think of hyper-converged systems as compute/storage/network/management modular building blocks that can be stacked and racked as needed to quickly create sufficient computing power. If IT needs more processing power or additional storage for a specific application, additional hyper-converged systems can be easily deployed right where it's needed: in close proximity to the existing application infrastructure.

Public and Private Cloud

Many cloud providers use hyper-converged systems to power their public cloud and provide on-demand resources to customers on an as-needed basis. Knowing that public cloud providers are using HCS in their customer-facing infrastructure gives companies an obvious clue as to how they should build their own private or hybrid cloud environments. Google, Amazon and Facebook all utilize hyper-converged systems in their cloud architectures.



Outlook Report

Specific Use Cases

There are some specific use cases where HCS is well-suited to the task at hand:

- **Dev/Test environments:** The ease of deployment of HCS makes them perfect for development, test and sandbox environments. As computational requirements go up or down, the HCS building blocks can be right-sized as needed for situations that seem to always be in flux.
- **Remote office/branch office (ROBO):** The modularity of HCS and the ability to pre-configure your HCS building blocks make them an ideal solution for ROBO operations. Need more power? Simply deploy additional hyper-converged systems as needed. Problem solved.
- **Cloud-in-a-box:** As companies begin to evaluate the concept of a private cloud for their production environments, the ability to quickly, inexpensively, and easily deploy HCS in a data center makes HCS a solid foundation for cloud-in-a-box installations.
- **SQL Server Consolidation:** Most enterprise companies have ongoing problems with SQL server sprawl and the expense of managing all of those disparate SQL Servers. HCS can be used to consolidate SQL Servers into a SQL private cloud that can be right-sized as needed to provide sufficient resources without breaking the bank.



Outlook Report

Best Practices for the Implementation of Hyper-converged Systems

Best practices for hyper-converged systems vary little from traditional system and converged system best practices, except that HCS best practice take into account the integrated nature of hyper-converged systems and tight integration of those components. As always, performing a thorough analysis and benchmarking of your HCS building blocks is crucial to long-term success of your HCS-based project. Don't trust HCS vendors--or any other vendor for that matter--to provide accurate performance numbers because your specific application and environment may differ wildly from whatever benchmarking scenarios vendors use to provide performance specifications to potential customers. In addition to considering the cost, performance, and task suitability of any hyper-converged systems you are evaluating, be sure to also consider ongoing maintenance and support costs, hard disk drives versus flash drives, and the usability of the management software that ties the HCS components together. Another important consideration for evaluating and implementing a HCS-based solution is the underlying hypervisor requirements for whichever HCS platform you choose. Some hyper-converged systems are hypervisor-specific and some are hypervisor-agnostic. There are advantages and disadvantages to each approach, so be sure to factor-in hypervisor support and hypervisor licensing costs during your evaluation of HCS platforms. Lastly, considering that software-defined everything is a big media buzzword for IT pros, be sure to carefully evaluate the ability to leverage software-defined, policy-based storage, networking, and management software. Particularly, if your company has already started down the path to a software-defined data center, be sure that your HCS purchase will integrate seamlessly into that environment.

Vendor Name: Gridstore

Product Name: Gridstore HyperConverged Appliance

Link to website: www.gridstore.com

Link to Data Sheet: <http://www.gridstore.com/products/gridstore-3/#hyperconverged-appliance>

Gridstore makes hyper-converged infrastructure for the Microsoft Cloud-Inspired Data Center. Consisting of up to four compute/storage nodes per system, the Gridstore HyperConverged Appliance can be expanded to a maximum of 256 nodes and 3PB capacity by adding additional HyperConverged Appliances or Storage Nodes that contain all-Flash memory or Hybrid Nodes with flash and traditional magnetic disk drives.

Gridstore's HyperConverged Infrastructure includes:

All-Flash HyperConverged Appliance: The All-Flash HyperConverged Appliance provides up to 24TB of Flash memory capacity. It is designed for applications that require low-latency, such as virtual desktop infrastructure (VDI), SQL Server and business-critical workloads. Each 2U appliance contains up to four compute/storage nodes consisting of 24 hot-swappable SSDs and two shared power supplies; each compute/storage node has dual Intel Xeon E5-2690 processors with 24 cores and 50MB cache, up to 256GB RAM and two 10GbE connections to the network.. Each node runs Windows Server 2012 R2 Data Center plus the Gridstore software.

Storage Nodes: Storage Nodes are available in all-Flash, hybrid or capacity models ranging in scale from 4TB to as much as 12TB. Storage Nodes can be used to expand flash storage in a HyperConverged Appliance without the need for more CPU-intensive Compute nodes, or as part of the HyperConverged Infrastructure as standalone units that allow you to scale hybrid storage capacity from 36TB to petabytes in size, all managed from the same user interface for files that have less stringent latency requirements.

Windows Servers: Any Microsoft Windows Servers with or without Hyper-V can become part of the infrastructure by simply installing the Gridstore vController, a software driver that interfaces with Gridstore Storage. This allows the HyperConverged Infrastructure to work with existing servers rather than being an island.

Gridstore uses a patented Server-Side Virtual Controller Technology (SVCT) to solve the mismatch between compute and storage (often called 'the I/O blender effect'). SVCT establishes a one-on-one relationship between the virtual server and its underlying storage and manages storage functionality on a per VM rather than LUN basis.

Server-Side Virtual Controller Technology: Gridstore uses the patented Server-Side Virtual Controller Technology (SVCT), which allows the creation of Virtualized Storage Resource Pools, called vPools, that aggregate storage resources such as IOPS, bandwidth and capacity, and expands the pool size as nodes are added.

The Gridstore Virtual Storage Controller (vController), which is installed as a Windows driver presents the virtual LUNS as a local SCSI device to the hypervisor. The vController operates on IO from each VM in a direct path, improving overall performance and reducing latency compared to offerings that run inside a VM. By having the direct access to the IO, the vController can isolate I/O from each of the VMs, optimize and prioritize it for precise control of the IO performance on a per VM basis. This control is delivered via QoS, called TrueQoS by Gridstore. TrueQoS prioritizes I/O per VM and can be managed at an IOPS level by setting minimum and maximum ranges, or by policy allowing VMs to be placed in a performance category.

SVCT also offers GridScale, a technology that enables Direct I/O from the vControllers to the storage resources in the grid. Direct I/O allows massively parallel performance scaling of the grid and eliminates the need for clustering. Most importantly for a HyperConverged environment, it protects data before it leaves the server and writes parallel stripes of erasure encoded data directly across a number of nodes. Data is protected to the N+2 level and is written in parallel to N nodes. By protecting data in this manner, Gridstore uses half of the raw storage as those products that mirror data while providing two (2) levels of redundancy. This provides not just a savings of storage, but also the stack of software and hardware necessary for each Compute/Storage node. Direct I/O also increases performance through parallel network I/O while 100% of a nodes IOPS are used for primary I/O rather than making copies, resulting in more bandwidth and IOPS available as more nodes are added to the grid.

Gridstore Software: Available with the Gridstore infrastructure is software meant to improve the performance of the grid, manage it, expand it, protect it and replicate it. GridCache is a distributed write-back cache that combines server- and storage-side cache into a global cache that eliminates 80% of the network IO and IOPS from shared storage. The Gridstore HyperConverged Infrastructure provides compute level availability via clustering, and storage fault-tolerance that protects against any combination of failed network connections, entire node failures, disk failures and silent bit rot. It uses erasure coding to protect data.

Installation is easy. Additional nodes, either compute/storage or storage only can be added on the fly. With the vController being a driver, installation takes minutes rather than hours, and with Instant vPool technology additional storage capacity is automatically added to the pool as the nodes are attached to the network.

GridControl is the management software for the GridStore grid. It is a Microsoft Management Console snap-in that integrates with Microsoft Server Manager and Hyper-V Manager and PowerShell. Gridstore also provides a native interface via Microsoft System Center - Virtual Machine Manager to allow the admin to have a single interface and policy management of the entire HyperConverged Appliance.

Finally, Gridstore leverages much of the Windows storage functionality such as snapshots, deduplication and encryption so that all functions can be implement on a per VM basis. As well, Windows Hyper-V Replica as well as other third-party products such as Double-Take to provide disaster recovery and business continuity via either local or remote asynchronous failover of Gridstore nodes.

Gridstore, the company: Kelly Murphy, Tomasz Nowak, and Antoni Sawicki founded Gridstore in 2009. The company has raised \$23.5 million in venture funding from Acero Capital, GCV Capital, Investec Ventures Ireland Limited and Onset Ventures participating. George Symons is the present CEO. Murphy, Nowak and Sawicki all have technical and engineering roles in the company. The company sells through the channel.

Use Case Profile



Company Name: Interact for Health

Company Type: Non-profit organization that improves the health of people in the Cincinnati region by being a catalyst for health and wellness.

Location: Cincinnati, Ohio

Products/Services: Gridstore HyperConverged Infrastructure

Contact name: Cliff Hastings, Director of IT

Previous Environment: 35 employees and contractors; Dell iSCSI SAN coming off warranty; 1.5TB of storage capacity

Challenges

- Wanted a solution optimized for Microsoft Hyper-V
- Cost-sensitive as a non-profit --Windows Storage Spaces not mature enough for production use; VMware vSAN was expensive even with academic pricing
- Need to support SQL Server, Remote Desktop Services, Veeam One, OpenManage and Web servers

Solution

- Installed the Gridstore HyperConverged Appliance in December 2014
- 3 nodes, dual 12-core processors with 256GB of RAM and four 10GbE ports in a failover cluster
- 10TB usable, 18TB raw – Each node has three (6) 996GB drives

Benefits/ROI

- "Gridstore was engineered with Hyper-V in mind," says Hastings, Director of IT for Interact for Health. "I liked the concept of the grid, being able to sustain multiple failures and still have your data alive and running. Its deep integration with Hyper-V, down to the VM level and being able to set your quality of service, being able to guarantee that the I/Os will be serviced in the order that you want -- that's not something at the Hyper-V level any other storage product has to my knowledge."
- "The relationship with Gridstore has been extremely productive. They have been very open to my feedback and gone above and beyond to make me a happy customer."

- “In the event of a failure the failover cluster moves the VMs to an active node and things just keep running. I am set up with a GridProtect level of one. What that means is that I can tolerate one failure of any component in the Grid without loss of data. At the disk level, I can lose one disk in all three nodes and the data would still be available. At the node level, I can lose one node and the data would still be available. As you add nodes and disks, you can raise your GridProtect level to tolerate a higher number of failures.”
- “There are notable performance increases over the iSCSI SAN.”
- “I am already seeing a tremendous increase in the runtime of my UPS system. It’s doubled.”
- “Just with the nine SSD drives I have, I can saturate a 10GbE link.”
- “I don’t have to have any special tools to administer the Gridstore Grid. It’s really point-and-click. The GUI is very intuitive. When we brought the appliance in, I had it configured and was provisioning storage within an hour.”
- “I took all the physical equipment in my infrastructure and added up the cost, the number of cores, the storage capacity and the amount of RAM. The amount I paid for the Gridstore hardware was the same as I paid for the infrastructure over the past five years.”

Future

- All workloads will be moved to the Gridstore HyperConverged Infrastructure
- Possible move to cloud for disaster recovery; move to Office 365 for Exchange Server



Outlook Report

SSG-NOW's Take on Hyper-Converged Systems

For the last several years, the adoption of hyper-converged systems has been slow due to a lack of market knowledge and education about exactly how HCS can be leveraged to save companies time and money in their application infrastructure. Now that Web-scale public cloud providers and social media companies are releasing their data center architecture and designs to the public, enterprise IT customers are realizing that their traditional data center design practices may be in need of an update. Though hyper-converged systems are not the perfect fit in every situation, there are certainly enough customer implementation data points and highly-capable HCS products now on the market that missing out on HCS may be missing out on an opportunity to radically change the way your IT works -- and how much it costs. The continuing advance of software-defined storage and networking, combined with the tight integration and hardware compatibility between HCS components, makes low-cost hyper-converged systems a compelling way to reduce data center complexity while simultaneously reducing costs. The ability to dynamically move virtual workloads, as well as the ability to add or remove application capacity and performance by deploying or re-deploying hyper-converged systems as needed creates a flexible, scalable, robust computing environment for enterprises looking to model their operations on the Web-scale cloud providers.

Paperless Productivity
1402 Third Ave.
Suite 812
Seattle, WA 98122
Tel 1.888.838.0042
info@PaperlessProductivity.com

www.PaperlessProductivity.com



Copyright 2015. All Rights Reserved. Storage Strategies NOW, Inc.