E-Guide

Shorten Your Data Storage Problem List with Hyper-convergence

TechTarget | SearchDataCenter

In this E-Guide:

While hyper-converged systems are best known for their ability to host virtualized workloads, they can also solve almost any data storage problem plaguing IT pros in traditional environments. This guide was designed to help you shorten your data storage problem list using hyper-converged (with four particularly pertinent problems highlighted), and also includes an HCI supplier roundup from TechTarget's 2018 HCI Technology and Market Survey.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Ten reasons to consider a hyper-converged data center

Alastair Cooke

A hyper-converged data center provides many benefits over traditional setups of servers, storage arrays and networks running between them.

Businesses depend on IT innovation to meet their strategic objectives. This can lead to a steady stream of new demands that stretch the limitations of existing data center resources, and the challenges become even more apparent when we consider the dynamic nature of demands on data center managers.

While hyper-convergence, the packaging of compute and storage into networked nodes with virtualization and management software natively installed, is still far from ubiquitous, products are quickly maturing to meet the scale and manageability enterprise administrators demand.

There are a lot of different reasons why IT organizations choose a hyper-converged data center setup. Depending on your situation, one or more of these reasons may apply to you.

1. The in vogue technology

One of the top reasons IT organizations look at hyper-converged data center setups is because other customers are looking at them. Hyper-convergence is a very fashionable phrase that is getting attention. This isn't in itself a reason to check it out, but it is a sign that

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

you should look at the rest of this list to see whether hyper-convergence suits your enterprise needs.

2. Scale-out workloads

Hyper-convergence is a scale-out platform, so it suits scale-out workloads. Virtual desktop infrastructure (VDI) is the classic scale-out workload, with hundreds or thousands of desktop VMs. Each virtual machine is a fairly small load, but the total VDI load is large. Web server farms and some database workloads can also benefit from a hyper-converged data center. The scale-out platform delivers consistent performance for applications designed to grow in this manner.

3. Growing workloads

Hyper-convergence makes it easy to expand the capacity of your data center by adding new nodes over time. If you have a phased rollout or simply experience long-term IT resource demand growth, hyper-convergence allows you to purchase the needed capacity as the workload grows, aligning infrastructure spending with the value that the workload delivers. It is easy to increase node purchases if growth occurs faster than planned and decrease them if growth lags behind your forecast.

4. Simplification

A significant feature of a hyper-converged data center is simplified management. Most hyper-converged platforms include policy-based virtualization management, which reduces the effort required to manage a collection of VMs. This simplification allows IT teams to

spend less time adjusting and optimizing the platform and focus more on the VMs and applications that deliver value to the business.

5. Eliminate technology silos

A hyper-converged data center eliminates the storage silo and its dedicated network. This relieves businesses from the need to maintain expertise in the storage array and network. In addition, hyper-convergence usually removes the need to manage the logical unit numbers, presentation and storage paths that come with a storage area network (SAN). This removes the need for dedicated expertise and allows IT teams to focus on other parts of data center and application stacks.

6. Infrastructure resources

Small projects that require modest resources may not suit deployment onto a large converged infrastructure platform or dedicated fleet of servers and storage arrays. Hyper-convergence enables data center staff to deploy a much smaller unit of infrastructure for smaller requirements. Adding a few nodes to your hyper-converged data center setup can serve the needs of a small project without making a large purchase. If a series of projects each requires a few nodes, the IT team can combine the nodes into a single, larger hyper-converged cluster. Combining these smaller projects into a larger infrastructure keeps the operational cost of running the projects from increasing as each project grows.

7. ROBO support and management

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Some organizations have a vast number of remote offices or branch offices (ROBO). Many of these locations need enough IT that they will benefit from a robust virtualization platform. It gets pricey deploying SAN and servers to each site, and the complexity of operating these systems makes the cost of supporting ROBO high.

Space is at a premium for some branch offices. Combining compute and storage into one box can make hyper-convergence much more compact than conventional infrastructure. If the branch office is a vehicle, ship, truck, aircraft or even submarine, then the space and power savings may prove critical. If there are many ROBO locations, then managing the platform at these sites becomes a challenge. Hyper-convergence simplifies management, keeping operational costs from spiraling out of control.

8. Hardware refresh

IT organizations often purchase, and replace, entire virtualization platforms together. This is an ideal time to evaluate hyper-converged products for your next infrastructure. If you deploy a hyper-converged data center, it is likely that the future process of refreshing the hardware will be far simpler than replacing virtualization platforms en masse. Hyper-convergence generally allows node-by-node replacement to update a cluster, meaning there are no large outages or long VM migrations to get a new infrastructure in place.

9. Staffing

Finding competent, specialized staff is a challenge for any IT department. The simplicity of a hyper-converged data center reduces the variety of skill sets you will need on the team.

Managing individual VMs eats up a lot of time. Using policies to manage groups of VMs and only having to think about noncompliant VMs reduces this workload. The same number of staff can do a better job of managing even more VMs than they currently handle when traditional processes are replaced by good policy-based management.

10. No public cloud

The ultimate simplification of IT is to not run IT, which is what makes cloud services so attractive. However, many organizations have legal or compliance requirements, or simply business rules, that prevent them from using cloud services. A hyper-converged data center allows for reduced infrastructure management of cloud services while keeping all IT on premises. Hyper-converged vendors design the platforms for programmability, making it easier to build private or hybrid cloud platforms on top.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Shorten your data storage problem list with hyper-convergence

Brien Posey, Microsoft MVP

Hyper-converged systems consist of a variety of resources -- not just storage. And while these systems are best known for their ability to host virtualized workloads, they can also solve almost any data storage problem plaguing IT pros in traditional environments. Here are four data storage problems hyper-convergence can solve.

Data storage problem #1: Storage silos

One of the primary challenges a hyper-converged infrastructure addresses is that of storage silos. Silos evolve as a result of business growth or changing technology. There are several definitions for storage silos, but the term often refers to a situation in which storage has been dedicated to an application, workload or group of systems.

Storage silos can present many different challenges, but most of these boil down to a lack of flexibility. Dedicated, siloed storage tends to be very rigid with regard to what it can be used for. It is usually far better to treat storage as a pool of resources that can be allocated on an as-needed basis rather than to dedicate storage hardware to specific applications.

One of the benefits of using a hyper-converged system is that it has its own storage hardware. The management layer makes it easy to treat that storage as a pooled resource that can be dynamically allocated as a workload's storage requirements change over time.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Data storage problem #2: Branch office storage

Branch offices and remote offices have always presented a special set of problems for enterprise IT. The main problem is that branch offices tend to be relatively small and often lack dedicated IT support. As such, providing a branch office with dedicated enterprise storage hardware might not be practical. Not only are there costs to weigh, but the IT staff must also consider supportability.

In a branch office, storage I/O must traverse a WAN link between that office and the organization's data center. This architecture can result in high latency and a poor overall experience for employees in the branch office. Caching can help to a point, but it isn't usually practical to cache everything.

A hyper-converged system tends to be a very good fit for branch offices. These systems are compact, self-contained and much simpler to deploy and maintain than a la carte hardware. Branch office employees can benefit from local access to virtual machines and to data storage. The IT department benefits from easy remote manageability. Furthermore, many hyper-converged systems support storage replication, which allows secondary data copies to be stored in the primary data center.

Data storage problem #3: Storage expertise

Most IT pros will probably agree that storage is more complex than it was even just a few years ago. But it isn't just storage complexity that can prove problematic -- problems can stem from the use of dissimilar hardware, as well.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Over time, an organization will inevitably need to add storage hardware or to replace existing storage hardware. Because technology continues to improve, newly acquired hardware will most likely differ from what the organization already has in place. It can be very challenging for an IT shop to support a variety of dissimilar storage hardware options.

Hyper-converged systems solve this data storage problem. Hyper-converged resources are designed to be modular, and the systems are also designed to hide most of the storage complexity. Storage is treated as a readily available, pooled resource. As such, IT pros are freed from the complex configuration and maintenance tasks inherent to using dedicated storage hardware.

Data storage problem #4: Storage capacity planning

Hyper-converged systems tend to eliminate, or at least reduce, the need for capacity planning. Hyper-converged systems are designed to be modular and to use a pay-as-you-grow architecture. This allows IT pros to purchase the exact amount of storage they need and to add additional modules should more resources be required. This eliminates the risk of over-purchasing storage resources.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Hyper-converged infrastructure 2018 technology and market survey

Chris Evans, Contributor

Hyper-converged infrastructure (HCI) has been around for a number of years. HCI systems consolidate the traditionally separate functions of compute (server) and storage into a single scale-out hardware platform.

In this article, we review what hyper-converged infrastructure means today, the suppliers that sell HCI and where the technology is headed.

HCI systems are predicated on the concept of merging the separate physical components of server and storage into a single hardware appliance. Suppliers sell the whole thing as an appliance or users can choose to build their own using software and hardware components readily available in the market.

The benefits of implementing hyper-converged infrastructure are in the cost savings that derive from a simpler operational infrastructure.

The integration of storage features into the server platform, typically through scale-out file systems, allows the management of LUNs and volumes to be eliminated, or at least hidden from the administrator. As a result, HCI can be operated by IT generalists, rather than needing the separate teams traditionally found in many IT organisations.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

HCI implementations are typically scale-out, based on deployment of multiple servers or nodes in a cluster. Storage resources are distributed across the nodes to provide resilience against the failure of any component or node.

Distributing storage provides other advantages. Data can be closer to compute than with a storage area network, so it is possible to gain benefit from faster storage technology such as NVMe and NVDIMM.

The scale-out nature of HCI also provides financial advantages, as clusters can generally be built out in increments of a single node at a time. IT departments can buy nearer to the time the hardware is needed, rather than buying up-front and under-utilising equipment. As a new node is added to a cluster, resources are automatically rebalanced, so little additional work is needed other than rack, stack and connect to the network.

Shared core

Most HCI implementations have what is known as a “shared core” design. This means storage and compute (virtual machines) compete for the same processors and memory. In general, this could be seen as a benefit because it reduces wasted resources.

However, in the light of the recent Spectre/Meltdown vulnerabilities, I/O intensive applications (such as storage) will see a significant upswing in processor utilisation once patched. This could mean users having to buy more equipment simply to run the same workloads. Appliance suppliers claim that “closed arrays” don’t need patching and so won’t suffer the performance degradation.

But running servers and storage separately still has advantages for some customers. Storage resources can be shared with non-HCI platforms. And traditional processor-

intensive functions such as data deduplication and compression can be offloaded to dedicated equipment, rather than being handled by the hypervisor.

Read more on hyper-converged infrastructure

- The rise of hyper-converged infrastructure – with compute, storage and networks in one box – seems ideal for SMEs, but is it always a better idea than traditional IT architecture?
- Hyper-converged infrastructure and containers are key deployments planned by UK customers in 2018, with flash storage, virtual machine storage and disk backup also prominent.

Unfortunately, with the introduction of NVMe-based flash storage, the latency of the storage and storage networking software stack is starting to become more of an issue. But startups are beginning to develop solutions that could be classed as HCI 2.0 that disaggregate the capacity and performance aspects of storage, while continuing to exploit scale-out features. This allows these systems to gain full use of the throughput and latency capabilities of NVMe.

NetApp has introduced an HCI platform based on SolidFire and an architecture that reverts to separating storage and compute, scaling each separately in a generic server platform. Other suppliers have started to introduce either software or appliances that deliver the benefits of NVMe performance in a scalable architecture that can be used as HCI.

HCI supplier roundup

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Cisco Systems acquired Springpath in August 2017 and has used its technology in the HyperFlex series of hyper-converged platforms. HyperFlex is based on Cisco UCS and comes in three families: hybrid nodes, all-flash nodes and ROBO/edge nodes. Fifth generation platforms offer up to 3TB of DRAM and dual Intel Xeon processors per node. HX220c M5 systems deliver 9.6TB SAS HDD (hybrid), 30.4TB SSD (all-flash) while the HX240c M5 provides 27.6TB HDD and 1.6TB SSD cache (hybrid) or 87.4TB SSD (all-flash). ROBO/edge models use local network port speeds, whereas the hybrid and all-flash models are configured for 40Gb Ethernet. All systems support vSphere 6.0 and 6.5.

Dell EMC and **VMware** offer a range of technology based on VMware Virtual SAN. These are offered in five product families: G Series (general purpose), E Series (entry level/ROBO), V Series (VDI optimised), P Series (performance optimised) and S Series (Storage dense systems). Appliances are based on Dell's 14th generation PowerEdge servers, with E Series based on 1U hardware, while V, P and S systems use 2U servers. Systems scale from single-node, four-core processors with 96GB of DRAM to 56 cores (dual CPU) and 1536GB DRAM. Storage capacities scale from 400GB to 1,600GB SSD cache and either 1.2TB to 48TB HDD or 1.92TB to 76.8TB SSD. All models start at a minimum of three nodes and scale to a maximum of 64 nodes based on the requirements and limitations of Virtual SAN and vSphere.

NetApp has designed an HCI platform that allows storage and compute to be scaled separately, although each node type sits within the same chassis. A minimum configuration consists of two 2U chassis, with two compute and four storage nodes. This leaves two expansion slots. The four-node storage configuration is based on SolidFire scale-out all-flash storage and is available in three configurations. The H300S (small) deploys 6x 480GB SSDs for an effective capacity of 5.5TB to 11TB. The H500S (medium) has 6x 960GB drives (11TB to 22TB effective) and the H700S (large) uses 6x 1.92TB SSDs (22TB to 44TB

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

effective). There are three compute module types: H300E (small) with 2x Intel E5-2620v4 and 384GB DRAM, H500E (2x Intel E5-2650v4, 512GB DRAM) and H700E (large) with 2x Intel E5-2695v4, 768GB DRAM. Currently the platform only supports VMware vSphere, but other hypervisors could be offered in the future.

Nutanix is seen as the leader in HCI, bringing its first products to market in 2011. The company floated on the Nasdaq in September 2016 and continues to evolve its offerings into a platform for private cloud. The Nutanix hardware products span four families (NX-1000, NX-3000, NX-6000, NX-8000) that start at the entry-level NX-1155-G5 with Dual Intel Broadwell E5-2620-v4 processors, 64GB DRAM and a hybrid (1.92TB SSD, up to 60TB HDD) or all-flash (23TB SSD) storage configuration. At the high end, the NX-8150-G5 has a highest specification Dual Intel Broadwell E5-2699-v4, 1.5TB DRAM and hybrid (7.68GB SSD, 40TB HDD) or all-flash (46TB SSD) configurations. In fact, customers can select from such a large range of configuration options that almost any node specification is possible. Nutanix has developed a proprietary hypervisor called AHV, based on Linux KVM. This allows customers to implement systems and choose either AHV or VMware vSphere as the hypervisor.

Pivot3 was an earlier market entrant than even Nutanix, but had a different focus at that time (video surveillance). Today, Pivot3 offers a hardware platform (Acuity) and software solution (vSTAC). Acuity X-Series is offered in four node configurations, from the entry level X5-2000 (Dual Intel E5-2695-v4 up to 768GB of DRAM, 48TB HDD) to the X5-6500 (Dual Intel E5-2695-v4 up to 768GB of DRAM, 1.6TB NVMe SSD, 30.7TB SSD). Models X5-2500 and X5-6500 are “flash accelerated” as both a tier of storage and as a cache. Acuity supports the VMware vSphere hypervisor.

Ten reasons to consider a hyper-converged data center

Shorten your data storage problem list with hyper-convergence

Hyper-converged infrastructure 2018 technology and market survey

Scale Computing has had steady growth in the industry, initially focusing on SMB and gradually moving the value proposition of its HC3 platform higher by introducing all-flash and larger-capacity nodes. The HC3 series now has four product families (HC1000, HC2000, HC4000 and HC5000). These scale from the base model HC1100 (Single Intel E5-2603v4, 64GB DRAM, 4TB HDD) to the HC5150D (Dual Intel E5-2620v4, 128GB DRAM, 36TB HDD, 2.88TB SSD). There is also an all-flash model (HC1150DF) with Dual Intel E5-2620v4, 128GB DRAM, 36TB HDD and 38.4TB SSD. HC3 systems run the HyperCore hypervisor (based on KVM) for virtualisation and a proprietary file system called Scribe. This allowed Scale to offer more competitive entry-level models for SMB customers.

Simplivity was acquired by HPE in January 2017. The platform has since been added to HPE's integrated systems portfolio. The Omnistack software that drives the Simplivity platform is essentially a distributed file system that integrates with the vSphere hypervisor. An accelerator card with dedicated FPGA is used to provide hardware-speed deduplication of new data into the platform. The HPE Simplivity 380 has three configuration options: Small Enterprise all-flash (Dual Intel Xeon Broadwell E-2600 v4 series, up to 1467GB DRAM and 12TB SSD); Medium Enterprise all-flash (Dual Intel Xeon Broadwell E2600-v4 series, up to 1428GB DRAM and 17.1TB SSD); and Large Enterprise all-flash (Dual Intel Xeon Broadwell E5-2600v4 series, up to 1422GB DRAM and 23TB SSD). Systems are scale-out and nodes can be mixed in a single configuration or spread over geographic locations.