



# Hyperconvergence Facts and FAQs

*What you need to know about the next-generation data center infrastructure for optimizing VM-centric environments*

 **TechTarget Custom Media**  
A TECHTARGET WHITE PAPER

# Next-Generation Data Center Infrastructure with Hyperconvergence

Cutting-edge data centers are most efficient when built on a next-generation architectural foundation. Take virtual machines. While today VMs are often deployed on traditional architectures with servers and separate network attached storage arrays, enterprises considering broad use of VMs will see the best results by implementing a new architecture called hyperconvergence.

Hyperconvergence eliminates the silos that traditionally separate a data center's compute resources (servers) from its storage resources (such as storage area networks or network-attached storage).

In a virtualized environment, silos create operational complexity, as you need coordination between servers and storage teams. What's more, multiple storage arrays may have been purchased for specific use cases, so admins need to manage each type of storage separately. In an organization using storage-area networks (SANs), there is often a LUN-to-VM construct mismatch that can be difficult to overcome.

Silos also waste expensive storage space, underutilize processors and memory, and can make it difficult to reallocate or scale in cost-effective, agile manner.

Enter hyperconvergence, which integrates storage with compute resources to provide a more efficient data center with higher resource utilization, better operational efficiency, easier management, and far more flexible scalability.

## TOP 10 BENEFITS

- 1.** Hyperconvergence simplifies the deployment of highly available server and storage clusters by working with a data center's existing server and storage resources.
- 2.** By eliminating silos between servers and storage, hyperconvergence reduces operational complexity and simplifies data center, application and resource management.
- 3.** It is easy to add more resources to a cluster without downtime, whether it's more compute or additional capacity. What's more, storage is never orphaned in "islands of underutilization."
- 4.** Hyperconvergence is built on a distributed file system that is purpose-built and optimized for virtual machine environments.
- 5.** Data management capabilities—snapshots and clones—are built into the distributed file system, and run transparently without affecting performance.

## How it Works

Hyperconvergence creates one big pool of compute and storage resources, which is ideal for virtual machines. Clustering storage together with groups of servers simplifies deployment of virtual machines and allows for rapid allocation of storage resources. Hyperconvergence is enabled through the use of a purpose-built distributed file system that sits between the hypervisor and all of the storage resources.

Indeed, hyperconvergence reduces operational complexity and works transparently to the applications. The distributed file system is invisible, providing a pool of fast, scalable resources to virtual machines. There's no performance impact, no new tools to learn, no new APIs to manage.

Hyperconvergence is designed for flexible, agile scalability, no matter what the requirements. Data center operators can add more storage, more compute capacity, or both, wherever it is needed most. In fact, with the right solution, operators can independently scale data caching and data persistence tiers independently. The best hyperconvergence systems eliminate the wasted storage found in traditional silo-based data centers to ensure that spinning hard disk drives (HDDs) and solid state drives (SSDs) are fully utilized. What's more, deduplication and compression are designed right into the distributed file system, and are leveraged across the cluster.

Let's learn more.

## TOP 10 BENEFITS cont.

- 6.** Data optimization capabilities—global inline deduplication and inline compression—are also built in, and work automatically to improve storage efficiency
- 7.** Operators can improve performance and capacity by scaling the caching tier independently from the persistent tier, as needed.
- 8.** Hyperconvergence is self-healing; data is mirrored to multiple disks and a rebalancing engine keeps applications running while failed hardware is replaced
- 9.** Software-based hyperconvergence is easy to implement and easy to learn, and administrators can manage it at the VM level using tools like VMware's vCenter.
- 10.** The economics of software-based hyperconvergence can be transformative, with low-cost, low-risk annual licensing that lets data centers experiment with this new technology.

# FREQUENTLY ASKED QUESTIONS

## Where does the hyperconvergence management layer go?

With hyperconvergence software, a distributed file system is installed onto the vSphere ESXi server and users operate the hyperconvergence layer from vCenter. Installing and configuring hyperconvergence on VMware takes less than 30 minutes to complete. With some vendors' hardware-based solutions, there is a separate management console that users must learn in order to allocate and manage hyperconverged resources.

## Hardware or software?

Software-only solutions are the best way to go. A software-based hyperconvergence solution allows you to continue using your preferred server vendor, instead of introducing another server vendor into the data center. And because it's software, you can scale the specific elements you'd like—compute, hard drives or SSDs—with no need to add entire nodes. From an operational standpoint, purchasing the same servers that you have been buying will allow you to leverage existing server management tools, such as Cisco UCS Manager or Dell OpenManage, and keep your server management environment unified. When it's time to refresh hardware, you can simply move the hyperconvergence software licenses to the new servers—there is no need to buy expensive hyperconverged appliances.

## What's adaptive scaling?

A challenge with traditional data center architectures is that once a bottleneck appears or a storage limit is reached, it can be very hard to scale to address that specific problem. With hyperconvergence, adaptive scaling means that you can continuously scale out capacity, and specifically add only what is needed, such as more compute-only blades, more HDDs or more SSDs to independently address the tier that needs expansion.

## How does hyperconvergence help optimize data and I/O?

Hyperconvergence is implemented as a distributed file system—and that unleashes powerful, transparent capabilities across the cluster's applications. For example, it has built-in capabilities for inline deduplication, which can reduce storage utilization tremendously. Data compression is also built in, not only saving storage but also reducing storage I/O requirements.

# FREQUENTLY ASKED QUESTIONS

## What about VM snapshots and clones?

It's important to look for a hyperconvergence solution that provides fast and efficient snapshots and clones of virtual machines. When implemented correctly, clones should be taken within seconds, allowing for rapid deployment of multiple virtual machines for test/dev or production environments. In addition, look for solutions that store only the differences once the snapshots/clones are created. The outcome is very space-efficient copies, using the least amount of storage necessary.

## What about resiliency?

With hyperconvergence, the distributed file system is self-healing. It pools all the storage resources in a cluster, and automatically mirrors all data to multiple physical HDDs and SSDs. Writes are cached and persisted in multiple locations until committed to multiple storage devices. If a storage unit fails, the file system automatically serves data from the mirrors. The system also proactively monitors and signals administrators to replace the faulty devices. Once replacement storage goes online, the file system automatically redistributes the data and restores redundancy.

## What is the Springpath hyperconvergence solution?

The Springpath Data Platform is purpose-built for hyperconvergence and integrates into VMware and other infrastructure platforms to turn standard servers into a single pool of compute and storage resources. Springpath Data Platform eliminates the need for network storage and using adaptive scaling capabilities allows customers to grow compute, caching or capacity resources independently, depending on their changing business needs. Springpath proactively monitors customers' environment to ensure resilient, always-on availability. Using an advanced set of data management and optimization capabilities such as compression, deduplication, space-efficient snapshots and clones, users experience transformative levels of resource utilization at much lower TCO than traditional IT architecture.

## How can I learn more about the Springpath Data Platform?

Visit : <http://springpathinc.com/smart-hyperconvergence.php>