A Convergence of Pros and Cons

Committing to a hyper-converged product is a big decision. Although scaling issues could present problems, the standardization and ease of management make HCI an attractive choice.

Balancing the Pluses and Minuses

Hyper-converged infrastructure has made the full journey into the IT world. It survived as a buzzword heard often in meetings, progressed to the rank of technology tried by a few brave IT teams, and now has cemented itself as an option—most would say a viable one—for organizations that want to reap the benefits of a consolidated and well-provisioned system. But, like so many technologies, it's difficult to judge whether the hype surrounding hyper-converged infrastructure is justified.

IT teams certainly enjoy the pre-engineered nature of the hyper-converged option, but is simple really the right choice? Also, do some of the problems that it was designed to solve remain? And how big a worry should vendor lock-in be?

To answer these questions, we looked to IT analyst Clive Longbottom. In this handbook, he discusses how the benefits of hyper-convergence may outweigh the negatives, and vice versa. If an organization can handle standardization and wants direct management through built-in software intelligence, HCI may make sense, even if scaling issues are a concern.

Next, analyst George Crump narrows the lens and discusses the specific components of hyper-converged systems. Finally, tech expert Brien Posey provides a hyper-convergence checklist so that an IT team can plan for the types of integration issues it'll encounter in the move to a hyper-converged infrastructure. ■

Patrick Hammond
Associate Features Editor
Data Center and Virtualization Media Group
TechTarget

Is Hyper-Converged Over-Hyped?

Understanding why hyper-converged infrastructure is getting so much attention requires a brief jaunt through recent history. The original mainframe approach—a simple setup consisting of servers, storage and networks—was replaced by distributed computing. During this transition, distinct vendor groups grew up around each area, such as Dell and HP for servers, EMC and NetApp for storage, and Cisco and Juniper for networks.

Organizations bought racks and built their own platforms, combining the best components from each area. This approach had its problems, mainly regarding personnel. The multiple platforms required skilled employees who understood the interplay among the various components and how to provision, monitor and maintain an overall platform. Many times, however, these platforms failed in their main aim: to support the business.

A few attempts were made to streamline the creation of IT platforms. Blades and bricks made life easier, as certain platform functions were combined and a specialized chassis was created, and they reduced the skills required to put everything together in an optimized manner. Blade computing, though, didn't take hold in a significant way. Having to buy a specific chassis that required regular upgrades created the perception that vendors had much too tight a grip on their customers.

THE EMERGENCE OF THE HYPER MODEL

In hopes of solving the problems of the past, vendors such as DataCore Software, Nutanix, SimpliVity and VMware created the hyper-converged platform. This model preconfigures compute, storage and networking (along with virtualization) to provide an optimized system that can be up and running in a short period of time. Incumbent vendors, such as Dell with its

FX2, IBM with its PureFlex, and Hewlett Packard Enterprise with its ConvergedSystem offerings, have also jumped on the bandwagon.

In essence, we are seeing a distributed computing version of the mainframe—along with many of the benefits and problems that come with this type of model.

The benefits are obvious. Hyper-converged vendors have engineered the system so that it operates at optimum levels. In addition to the pre-engineered hardware, hyper-converged vendors include a pre-configured software stack, ranging from a relatively simple hypervisor, operating system and systems management environment to a full-blown elastic private cloud with intelligent workload management software. Purchasers need to have a reasonable understanding of how they will be using the system to ensure that they choose the right approach.

Whether an organization chooses a workload-specific system or a flexible cloud platform, life is a lot easier once the hardware is delivered. As everything is consolidated into one system, the buying organization's IT department has far less to do when provisioning the system. Most of the time, IT simply needs to unpack it, plug it in, input a few variables (IP address, DNS settings and so on) and then start installing applications.

Server, network and storage interactions are all managed directly by the software intelligence the vendor builds into the system. And this is where one of the battles will be fought: Which vendors can build the best intelligence into their hyper-converged systems?

With all of these benefits, and just a slight nagging worry over how intelligent the software is, it all sounds great, doesn't it?

WHAT ARE THE DRAWBACKS?

Hyper-converged infrastructure may be wonderful for some organizations, but there are downsides.

Most hyper-converged systems are built on commodity components (as in, they use standard, mainly Intel, CPUs, SATA or SAS disk drives and standard network connectors). However, a lot of the internal connectivity is proprietary, which leads back to the same issues found with blade computing and the need for specialized chassis. If the internal bus structure of a hyper-converged system requires all extra components to be purchased from the original vendor, then you are beholden to that vendor. The system may have less commodity storage (e.g., PCIe or mSATA modules) and it may include offload processors such as GPUs.

But this arrangement may not be as worrisome as it seems. As long as the hyper-converged system is standardized on the outside (that is, it speaks TCP/IP over Ethernet and has no external dependencies on that vendor's specific equipment), then using proprietary internal technology won't be an issue.

In fact, there is a hidden benefit in having vendors add their own touches to these systems. If all hyper-converged systems had to be based on standard components with standard connections and standard firmware, all systems would be the same. By allowing for innovation at the internal hardware level, Vendor A's hyper-converged system may meet your distinct requirements far better than Vendor B's.

Still, expansion issues may persist. Some hyper-converged systems are sized for a specific environment. If your organization wants to use a hyper-converged system as a complete private cloud platform, then you need to ensure that it can elastically share resources and that those resources can be easily multiplied as required.

Some systems have the requisite space to expand while using the same vendor's equipment. Others require that the buyer purchase another system and essentially cluster the two together. Some systems allow companies to easily add other vendor's equipment in areas such as NAS or SAN storage; others struggle to make use of external resources in an optimized manner. Many systems require that if one type of resource is added, more of the other resources must be added at the same time. It might make sense to look for systems that have built-in systems management and external third-party equipment management capabilities.

Overall, it makes sense to look at hyper-converged systems, as they allow an IT department to react faster to their organization's needs. Buyers need to make sure that any system is fit for purpose and that it has the flexibility to grow as needs change. —*Clive Longbottom*

The Three Tiers of Hyper-Converged Architectures

The goal of hyper-convergence is to drive out complexity and to drive down costs by consolidating compute, network and storage layers into a single tier. While it may seem that the turnkey nature of hyper-converged architectures absolves an organization of the responsibility of caring how the storage infrastructure works, the opposite is true. Understanding how the storage software works, what storage and media were selected, and why these choices were made helps IT planners make good a purchasing decision.

To better understand how hyper-converged architectures work, it's important to drill down into the three tiers.

THE COMPUTE TIER

This tier is made up of two parts: the physical server hardware and the hypervisor software. Both elements are critical to the overall success of the hyper-converged project as well as the suitability of a particular hyper-converged architecture for the organization.

There are three compute configurations available:

- A bundled hyper-converged system, in which the software and hardware come pre-integrated, and all upgrades must come from that vendor. This configuration sacrifices long-term flexibility to gain the simplicity of a turnkey system.

- Software-only configurations, where the organization provides its servers. This approach sacrifices simplicity of setup for long-term flexibility.

- A "software-mostly" model, in which the software can be installed on almost any server, but it comes with some proprietary hardware such as a PCI Express board to run specific functions.

The hypervisor software will also vary. It can be from a brand name or one of the Linux hypervisors. Some vendors have taken the Linux hypervisors and created their own, customized version of the hypervisor.

THE NETWORK

The least radically affected of the three elements being converged is the network. In most configurations, there is no longer a dedicated storage network because there is no longer a dedicated shared storage system. But there is a significant amount of networking involved, especially at scale—nodes have to communicate and storage needs to be allocated.

Many hyper-converged architectures end up dedicating network cards and switch ports to this traffic, essentially "unconverging" the network element. Network management functions are sometimes integrated so that configurations can be set from within the converged management interface.

THE STORAGE INFRASTRUCTURE

The storage tier is greatly changed by hyper-converged architectures. Typically, the shared storage system is eliminated, and internal storage within the nodes is aggregated into a virtual volume. The storage part of hyper-converged architectures is a software-defined, scale-out storage system that can run as a component of the hypervisor architecture.

The storage hardware is still important in the hyper-converged architecture design, and the parts are similar to a hardware-based storage system. First, the organization needs to make sure the shared resources (CPU and network) have enough bandwidth to drive storage functions as well as all the virtual machines it may handle. Second, the media that will be installed inside the servers is important; the right balance of flash and hard disk drives is critical so that cost and performance are in sync.

While hyper-converged architecture is often delivered as a single component, there are layers beneath the onion. The more that IT professionals understand, the more they will be able to select the right product for their organization. —*George Crump*

Hyper-Converged Deployment Issues

Hyper-converged systems have gained a reputation as a plug-and-play virtualization platform. They are designed to make the deployment process almost effortless, but every vendor operates a little differently.

While there is no such thing as a planning checklist that applies to every hyper-converged architecture system, there are certain aspects of the process that are applicable to most systems.

NETWORK CABLING

Most hyper-converged architecture vendors have requirements (or recommendations) for the requisite number of free switch ports. For example, VMware requires eight 10-gigabit Ethernet ports on the top-of-rack switch for each EVO:RAIL appliance. VMware also requires link aggregation to be disabled for these switch ports.

Keep in mind that as important as top-of-rack switch ports are, you need to ensure there is sufficient bandwidth to connect the hyper-converged system to the rest of your network.

IP ADDRESSES

One of the most critical aspects of the planning process involves establishing a plan for IP address provisioning. While your DHCP server might provide each node with one or more network addresses, there is typically more planning that must be done.

Vendors will have specific requirements for their hyper-converged architectures, but at the very least you will need to provision blocks of IP addresses for management traffic, monitoring, out-of-band management, clustering (VM-level migration) and, perhaps, storage. VMware, for example, has established guidelines for its EVO:RAIL systems:

- Reserve one IP address on the management virtual LAN (VLAN) for EVO:RAIL/vCenter Server.
- Reserve one IP address on the management VLAN for vRealize Log Insight.
- Reserve four contiguous IP addresses on the management VLAN for ESXi hosts for each appliance.
- Reserve four contiguous IP addresses for virtual SAN for each appliance.
- Reserve four contiguous IP addresses for vSphere vMotion for each appliance.
- Reserve four IP addresses for out-of-band management for each appliance.

While these requirements are specific to VMware, hyper-converged systems based on competing hypervisors, such as Microsoft's Hyper-V, tend to have similar requirements.

VLANS

Not every hyper-converged architecture makes use of a VLAN, but many of them use VLANs to isolate various types of network traffic. It is important to check your vendor's deployment guide to see if VLANs are required.

If your hyper-converged architecture supports out-of-band management, and you wish to take advantage of this feature, you will either need to have a separate switch to handle out-of-band management traffic or enough extra ports on your top-of-rack switch to provide network connectivity to each node's out-of-band management adapter.

Some hyper-converged architectures use VLANs to isolate various types of network traffic. It is important to check your vendor's deployment guide to know for certain.

Planning a hyper-converged architecture deployment tends to focus on network-related issues, such as cabling, VLAN usage and IP address provisioning. But don't forget to check your system's deployment guide to ensure you've met the power and cooling requirements. —*Brien Posey*

ABOUT THE AUTHORS

HOME

EDITOR'S NOTE

OVER-HYPED?

THREE TIERS

DEPLOYMENT

GEORGE CRUMP *is president of Storage Switzerland, an IT analyst firm focused on the storage and virtualization segments. Prior to founding Storage Switzerland, Crump was CTO at one of the nation's largest storage integrators, where he was in charge of technology testing, integration and product selection.*

CLIVE LONGBOTTOM *is the co-founder and service director at Quocirca and has been an ITC industry analyst for more than 15 years. Trained as a chemical engineer, he worked on anti-cancer drugs, car catalysts and fuel cells before moving into IT.*

BRIEN POSEY *is a seven-time Microsoft MVP with two decades of IT experience. He has published thousands of articles and has written or contributed to dozens of IT books. Prior to becoming a freelance writer, Posey served as CIO for a national chain of hospitals and healthcare facilities.*

STAY CONNECTED!

Follow @ServerVirt_TT today.

TechTarget

A Convergence of Pros and Cons is a SearchServerVirtualization.com e-publication.

Margie Semilof | *Editorial Director*

Phil Sweeney | *Senior Managing Editor*

Patrick Hammond | *Associate Features Editor*

Linda Koury | *Director of Online Design*

Rebecca Kitchens | *Publisher*
rkitchens@techtarget.com

TechTarget
275 Grove Street, Newton, MA 02466
www.techtarget.com

COVER ART: FOTOLIA