Is A Hardware Approach Breaking Hyperconverged Architectures?

Prepared by: George A. Crump, Senior Analyst
Prepared: July 2017

A popular hyperconverged deployment model is a turnkey solution where the hardware and software are licensed together in a pre-tested, pre-integrated solution. It brings much of the positive attributes of hyperconvergence while smoothing out the initial implementation and in theory eliminating finger pointing when support is needed. Some vendors even add customized hardware and optimize their software to take advantage of it. Is this hardware approach the right way to deliver hyperconvergence or is it breaking some of its basic tenets?

The Hyperconverged Promises

Above all else hyperconvergence is supposed to simplify the data center. It accomplishes this simplification by converging compute, storage and networking on standard servers. Instead of having to scale these different resources independently, hyperconvergence allows them to scale like Lego pieces as each brick (or node) is added to the architecture. And with that node comes an appropriate amount of additional compute, storage and networking resources.

There is also a promise of integration of compute and storage services. As additional virtual machines (VMs) are created, the storage pools for those VMs are made at the same time from the same interface. The goal is to eliminate the single biggest challenge facing a virtualized infrastructure – dealing with scale as it grows.

The Difference Between Hardware and Software Approaches to Hyperconverged

On the surface, a turnkey solution looks like the way most organizations should go. The system is delivered with software already installed. All the IT administrator has to do is apply power, plug it into the network and start creating or transferring VMs.

The alternative, hyperconverged software, means IT will need to find available server resources, make sure compute is compatible with the software, make sure the storage inside those servers is supported and appropriate for the task at hand. The advantage of a software approach is that it can leverage existing servers, can support virtually any server in the future, and can support latest generation hardware at a more rapid pace. In most cases, the software approach can also support a wider variety of hypervisors.

The Hyperconverged Problem

Ideally turnkey hyperconvergence is deployed "green field", meaning as a new project or new data center is started, it starts with hyperconvergence. This is one of the reasons hyperconvergence is so popular with virtual desktop projects since most of those started new, after the concept of hyperconvergence was understood.

The problem is that by the time most data centers are ready for hyperconvergence they already have an investment in physical servers, and in most cases already are well down the virtualization path. This reality means turnkey hyperconverged solutions have a timing concern. In theory IT needs to purchase them at about the same time old servers are coming off maintenance or have fully depreciated.

Software solutions, on the other hand, can be deployed directly onto existing hardware architectures and leverage existing hypervisors. There may be a requirement to install additional storage capacity or add flash drives in the physical servers, but the servers themselves do not need to be replaced.

Another challenge that hardware hyperconverged architectures face is how to deal with new hardware as it becomes available. A current example is new servers that are coming out with Intel's latest processors and a new protocol to interface with flash drives. The first step anytime new hardware is introduced to the market is for that hardware to be available as a standalone physical server. Once that hardware is available, turnkey solution vendors have to test this new hardware, optimize their code and then finally deliver their software based on it. Which means it could take more than a year before the turnkey hyperconverged solutions can update to the latest technology.

Another aspect to this problem is that many hardware hyperconverged solutions do not support the mixture of different types of nodes within the hypervisor cluster. The nodes must all be the same or very similar. A new generation of hardware may force the turnkey provider to require that their customers create separate clusters for each generation of hardware, increasing complexity.

Software solutions are not encumbered by these challenges. They can be deployed on new hardware as soon as the server vendors release it. In most cases customers will see an immediate performance or capabilities boost. In situations where the software itself needs updating to support a new capability of the processor or server hardware, a software update will deliver that support in-place and in most cases non-disruptively.

Related to the challenge of additional nodes is dealing with the efficient use of the nodes before any scaling is done. The hyperconverged solution should also scale-up, meaning that resources like additional compute, more memory or more storage, can be added to specific nodes as needed. The scale-up nature of the cluster, almost by default, assumes that the hyperconverged solution is able to support nodes of different sizes and types as described above.

The final consideration is vendor lock-in. Vendor lock-in is a fact of life in the data center. IT's job is to try to limit that lock-in as much as possible. A turnkey solution is essentially the black box approach. Once purchased, the organization is locked into a specific server and storage hardware vendor and, in most cases, a specific hypervisor. In addition, that hardware (in most cases) has to be provided by the turnkey hyperconverged vendor, even if the organization has a better purchasing relationship with the exact same server vendor.

Hyperconvergence software solutions eliminate much of the lock-in. Certainly the physical hardware used is up to the organization. It can leverage the contracts it has in-place with its server vendors. Some hyperconvergence software solutions are neutral as to which hypervisor the organization uses. IT can change direction in terms of the hypervisor and the hyperconvergence software solution will continue to work for them.

Overcoming the Hyperconvergence Software Problem

Hyperconvergence software has a lot of advantages vs turnkey systems. It can leverage the server hardware that is in-place and support new server hardware as it becomes available. It also enables the organization to "change its mind" in terms of server vendors, storage media vendor and even hypervisor vendor, if it so chooses. Despite these positives, hyperconvergence software solutions have one major challenge that causes a lot of turnkey solutions to be chosen, the fear of a complex startup.

Software hyperconverged solutions can overcome this challenge be developing quality partnerships with leading OEM server vendors like Lenovo, Supermicro and others. Depending on the partnership arrangement the hyperconvergence software provider should be able to provide an almost identical startup experience as the turnkey vendors but with the flexibility on future hardware and software selections.

StorageSwiss Take

Hyperconverged software solutions look the best on paper but cause too many concerns on initial implementation, which is why hardware-based solutions are so popular. Companies like Maxta are overcoming the initial bar to entry by offering solutions that are predefined and pre-validated with specific server hardware vendor solutions. These vendors combine Maxta's software along with partner platforms which removes interoperability and performance guesswork while also simplifying the ordering process. The result is the best of both worlds, rapid start-up and flexible future expansion.

Hyperconverged software solutions however are an excellent fit for an organization that recently invested in server hardware but now wants to take advantage of hyperconvergence. They can install the software directly into their existing hypervisor cluster and start seeing the benefits of a hyperconverged solution.

Finally, hyperconverged software solutions are easier to test and experiment on. For example many environments have a concern over the performance impact of adding storage processing to an environment that has been solely focused on application processing. A software solution should allow the download of software so that an initial test suite can be built without bringing in additional hardware.

Sponsored by Maxta

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplace. For more information please visit our website **storageswiss.com**