

Comparing Hyperconvergence Architectures

*By Scott D. Lowe, vExpert
Co-Founder, ActualTech Media
January, 2015*

Table of Contents

- Table of Contents2
- Introduction..... 3
- Hyperconvergence Defined 3
 - Software-Defined Storage Characteristics.....4
- Benefits of Hyperconvergence 6
 - Provides a Streamlined Infrastructure Deployment and Scaling Processes6
 - Less Data Center Complexity6
 - Data Center Performance Increase..... 7
 - Ability to Leverage Commodity Hardware..... 7
 - MaxDeploy Combines Flexibility and Ease of Deployment..... 8
 - Procuring the MaxDeploy Solution..... 9
- Summary.....10
- About the Author10
- About Maxta.....10



ActualTech Media © 2015. All rights reserved.

Under no circumstances should this document be sold, copied, or reproduced in any way except with written permission.

The information contained with the document is given in good faith and is believed to be accurate, appropriate and reliable at the time it is given, but is provided without any warranty of accuracy, appropriateness or reliability.

The author does not accept any liability or responsibility for any loss suffered from the reader's use of the advice, recommendation, information, assistance or service, to the extent available by law.

Introduction

Hyperconverged infrastructure, it is safe to say, is a segment of the broader technology market enjoying significant discussion and analysis as end users and analysts alike discover the business and technical benefits that can be realized. As is the case with many technology topics that gain steam very quickly, there are a lot of moving parts around hyperconverged infrastructure and there are no one-size-fits-all solutions out there. In fact, with more vendors entering this space every day, there is consistent stream of new solutions on the market, each with its own set of strengths and challenges.

Companies of all sizes are shopping for hyperconverged infrastructure solutions and when they do, they will find that they must weigh the benefits of a software-centric approach vs. an appliance-based approach. In this paper, you will learn about the pros and cons of each of these approaches.

Hyperconvergence Defined

Before getting to the differences in solutions, it's important to understand the definition of Hyperconvergence... both what it is and what it isn't. In it's simplest form, hyperconverged infrastructure is the melding together of previously disparate data center resources, namely servers (compute and memory resources) and storage (a combination of hard disk drives and solid state disks). In such systems, the storage component takes a software-defined approach. Software-defined storage is discussed later in this paper. See Figure 1 for to see how hyperconverged infrastructure impacts the data center environment.

Because hyperconverged infrastructure is based on the concept that storage and servers are effectively recombined, some may believe that the data center market is returning to the 1990s when direct-attached storage was the de facto standard method by which storage was deployed. While this return to the past is partially true, there is a whole lot more to the hyperconvergence story.




Figure 1: Hyperconvergence eliminates an entire tier of infrastructure, vastly simplifying the environment

Unlike older architectures that relied upon server-based storage, hyperconverged infrastructure does not result in discreet islands of trapped storage. In the 1990's direct-attached storage systems resulted in significant waste and inefficiency since individual servers could not seamlessly leverage storage resources from other systems. Instead, hyperconverged infrastructure relies on an architectural concept that has become known as software-defined storage (SDS).

Software-Defined Storage Characteristics

In traditional environments, direct-attached storage is truly trapped on individual servers. For organizations that have chosen solutions based on SDS, however, new capabilities are introduced. As such, the barriers between servers are effectively shattered.

With SDS, a software-based hardware abstraction layer is introduced atop each server in the environment. In these cases, the servers basically become nodes in a larger cluster that is comprised of all of the servers included in the solution. Server-based storage is one of those technologies that just works and doesn't generally require customized drivers to software in order to operate. There are already mechanisms for making server-based storage work. As such,

hyperconverged infrastructure vendors leveraging SDS for their solution can generally leverage any kind of storage without having to rewrite software for each new hardware device. This eliminates the need for proprietary hardware to drive specific storage functions and enables the use of commodity hardware, which can save significant dollars on the data center investment.

Software-defined constructs generally follow a three-step process:

1. **Abstraction**, which happens on the individual servers, is just the first part of the process, though.
2. **"Pooling"** is the next part of the process. Pooling basically aggregates all of the previously abstracted storage into a single resource pool, which is then presented as a cohesive unit of storage to the clustered environment. This part of the process introduces a global namespace that allows the individual nodes to participate as equal members of the larger cluster.
3. **Automation** is the third and final part of the "abstract, pool, automate" software-defined manifesto and is the layer at which policies and higher order functionality is introduced into the system and, depending on the features included, determines whether the outcome is worthy of enterprise attention.

There may be concern that the use of commodity hardware with a software layer may spell the end of enterprise-class storage features that have had a positive impact on overall storage costs and data protection capabilities. However, nothing could be further from the truth. SDS-based hyperconverged infrastructure solutions carry with them several high-end capabilities, including:

- Inline deduplication
- Multilevel data protection
- Replication
- Thin provisioning
- High availability

SDS and hyperconverged infrastructure carry with them the promise for massively simplified data centers that eliminate existing points of failure and complexity, such as those routinely introduced in traditional SAN environments.

Benefits of Hyperconvergence

Hyperconverged infrastructure has a myriad of technical benefits and capabilities, but, increasingly, companies are far more interested in the financial side of the IT house and here, hyperconvergence doesn't disappoint. There is a great mix of both technical and financial benefits.

Provides a Streamlined Infrastructure Deployment and Scaling Processes

No longer do legions of IT staff members from separate functional areas need to come together to coordinate their deployment activities. A hyperconverged infrastructure solution is based on Lego-like bricks of infrastructure, with each brick carrying with it all of the resources that make the data center work. This "unit of data center infrastructure" means that, as new resources are needed, already-harried IT staff members simply add another block to the data center and add it to the existing resource pool without having to go through dozens of steps to configure individual resources.

Less Data Center Complexity

SANs emerged as the storage solution of choice because they addressed the problems that IT was facing in managing storage. Unfortunately, this centralization of the storage resource carried with it some major challenges, including performance challenges that were introduced by virtualization. The sheer variety of I/O patterns that traditional hard disk-based storage must suffer results in unpredictable performance and an inability to react to new business needs that may arise.

From an operational standpoint, the traditional SAN has become a particularly vexing object to manage, often requiring specialized – and expensive – skills. The combination of hyperconverged infrastructure and SDS brings with it significant simplification in getting work done. There are no more RAID groups to configure, LUNs to create, or monolithic SAN devices to manage. Everything is operated as a part of the cluster.

When the time comes to add resources to the data center, scaling a hyperconverged environment is the essence of simplicity. An administrator need only add a new unit of infrastructure and add it to the existing cluster. There is no need to individually configured resources.

As CIOs and other business leaders seek ways to better leverage IT, the ability to make operations simpler helps to achieve this goal. Since less effort is expended in operations, more attention can be paid to those activities that have a direct impact on the bottom line.

Data Center Performance Increase

Besides simplifying operations and scale, one of the primary reasons that organizations are turning to hyperconverged infrastructure is to improve data center performance. Although the simplicity behind the technology is compelling, modern organizations have performance demands that legacy infrastructure simply cannot meet. By leveraging a combination of hard disk drives and flash storage, hyperconverged solutions are imbued with plenty of capacity – TBs – and plenty of storage performance – IOPS – to meet even the most demanding workloads. In such environments, the solid state storage plays a vital role in acting as a performance storage tier while also accelerating any activities that take place on the much slower hard disk-based storage. Although the storage market is becoming flooded with storage arrays that have only flash disks, such solutions remain expensive. The hybrid approach taken by most hyperconverged infrastructure is more than sufficient for the majority of mainstream enterprise workload needs.

Ability to Leverage Commodity Hardware

Traditional data center environments are chock full of expensive proprietary hardware. In general, proprietary hardware costs far more than commodity hardware due to volume and due to the specific nature of specialized hardware. Hyperconverged infrastructure combined with SDS leverages commodity hardware that often costs orders of magnitudes less than specialized equipment. This lower cost hardware is then bound together by the powerful software layer, which handles everything necessary to ensure availability and performance. For the ability to leverage commodity hardware, the industry has Intel to thank. As the company keeps releasing processors packed with more and more cores, there is massive computing power that can be brought to bear to replace functions that formerly required dedicated hardware, such as comprehensive inline data deduplication.

The ability to leverage commodity hardware and to scale more linearly provides the organization with a far more predictable cost structure for the data center. At the same time, being able to scale in “blocks” enables a far more granular scaling opportunity leading to less waste.

MaxDeploy Combines Flexibility and Ease of Deployment

In general, hyperconverged infrastructure is available as either pure software or is available as a part of a bundled hardware appliance on which a hyperconvergence vendor preinstalls their tools. There are many examples of both kinds of options in the market and each option typically comes with a set of pros and cons.

For those that choose to buy software only, the primary downside is the fact that the hardware and software must be sourced separately and then integrated by either the customer or by the reseller. On the upside, though, that customer gets eminent choice and flexibility in hardware. For those that choose to buy a prebuilt appliance, the downside has generally been that the vendor chooses the hardware platform and configuration. The customer gets little to no choice and, as a result, loses a lot of flexibility. Of course, going the appliance route is often easier as the customer is buying a resource that is ready to go from Day One.

What if, however, a customer could get the best of both worlds? What if a customer could buy hyperconverged server hosts from any server vendor they wanted and size those servers to meet their specific needs while, at the same time, getting the benefits of a pre-built, pre-configured, and pre-validated appliance? In this scenario, the customer would receive their choice of servers with the hyperconverged infrastructure software already installed and ready to go.

Enter MaxDeploy.

Don't make the mistake of thinking that MaxDeploy is just a clever way for Maxta to help customers procure hardware; it's far more than that. MaxDeploy provides customers with a *single source and ordering mechanism* for a fully configured hyperconverged infrastructure solution. The customer does not buy hardware and software separately. The customer buys a bundle from a single Maxta partner. The customer gets an appliance, but they get to choose exactly what they want. Each resource – compute, RAM, storage, and networking – is sized exactly to the customer's specifications. Each appliance arrives at the customer facility configured, and ready to go.

For those customers that want to use Maxta's software defined storage product and create their own hyperconverged infrastructure appliance, Maxta sells their MxSP product. In general, MxSP is used by service providers and enterprises that have skilled IT staff and prefer to manage their own configurations and set up.

Procuring the MaxDeploy Solution

To help customers get started in their journey, Maxta makes available a configuration tool (shown in Figure 2) that completely guides them through the process, including defining their intended use case for the platform and their preferred server vendor. Maxta has partnered with Cisco, Dell, HP, Lenovo, Quanta, and Supermicro. Maxta has completed server vendor certification testing for other major server vendors as well, including Cisco and HP. The ability to choose vendor enables customers to opt for a name brand or run on commodity hardware to increase savings. By leveraging commodity hardware, MaxDeploy reference architectures are designed to save customers from 30% to 50% depending on whether they are running on branded or commodity servers. Once the server vendor selection is made, depending on the use case, the configuration tool then asks for necessary storage capacity or the number of virtual desktops that need to be supported.

This server agnostic approach provides the ultimate in resource flexibility and can often be easier to buy for organizations that are strict as to the server vendors they are allowed to use. By being able to choose the exact hardware upon which to run the data center environment, organizations gain the ability to more granularly address resource constraints as they arise. For example, for most companies, storage capacity is the resource that needs the most ongoing attention. These companies may not want to buy new nodes that have a great deal of RAM since storage is the affected resource, not RAM. Maxta's solution enables customers to grow data center resources in any necessary direction.

Maxta also takes an “any hypervisor” approach to hyperconvergence, allowing customers to select the hypervisor that best meets their needs while still providing all of the benefits of hyperconvergence.

USE CASES	VENDORS	CAPACITY (TB)	DESKTOPS
<input type="radio"/> DR Storage	<input type="radio"/> Cisco	<input type="radio"/> 10	<input type="radio"/> 100
<input checked="" type="radio"/> Primary Storage	<input type="radio"/> Dell	<input type="radio"/> 25	<input type="radio"/> 200
<input type="radio"/> ROBO	<input type="radio"/> HP	<input type="radio"/> 50	<input type="radio"/> 500
<input type="radio"/> Test and Development	<input checked="" type="radio"/> Lenovo	<input checked="" type="radio"/> 100	<input type="radio"/> 1000
<input type="radio"/> VDI	<input type="radio"/> Quanta	<input type="radio"/> 200	<input type="radio"/> 2000
	<input type="radio"/> Supermicro		

SUGGESTED SERVER CONFIGURATION			
SERVER VENDOR:	Lenovo	CPU:	8x25x12C
SERVER MODEL:	ThinkServer RD440	MEMORY IN GB:	8x128
NUMBER OF SERVERS:	8	SSD:	8x2x800GB
CAPACITY IN TB:	100	HDD:	8x14x900GB
		NETWORK:	8x2x10GbE




Figure 2: The MaxDeploy Configuration Tool

Summary

Maxta's approach to appliance-based hyperconvergence carries with it some significant benefits when compared with other appliance-based vendors.:

- **The hardware itself is defined by the customer, not the vendor.** With most appliance-based approaches, although there is often some room to tweak individual resources, there is far less flexibility in many appliance-based offerings than there is in creating a truly custom configuration. This lack of configuration may force customers to buy more resources than they need over time. Moreover,
- **Maxta is designed to support a full range of hypervisors..** Whether you are looking to run on VMware Sphere or looking to mitigate licensing fees by using KVM, Maxta provides the flexibility to run on the hypervisor the customer wishes to use and provides the flexibility to make changes when desired. Maxta is also a leader in the hyper-convergence space in providing enterprise class data services for OpenStack deployments.
- **Customers can choose to leverage existing hardware.** Maxta extends its flexibility to allow customers to use existing hardware. Many other appliances require that customers either replace existing equipment or install the solutions side-by-side.

About the Author



Scott Lowe is co-founder of ActualTech Media and the Senior Editor of EnterpriseStorageGuide.com. Scott has been in the IT field for close to twenty years and spent ten of those years in filling the CIO role for various organizations.. Scott has written thousands of articles and blog postings and regularly contributes to such sites as TechRepublic, Wikibon, and virtualizationadmin.com.

About Maxta

Maxta was founded in 2009 by industry veterans with a mission to maximize the promise of hyper-convergence. Maxta award-winning software-defined storage solutions provide organizations the choice to deploy hyper-convergence on any x86 server, use any hypervisor and any combination of storage devices. The simplicity of Maxta's VM-centric solution reduces IT management to further maximize cost savings. Hyper-scale, enterprise-level data services and capacity optimization, empower organizations to hyper-converge, eliminating the need for SAN or NAS devices. Think outside the storage box. For more information visit <http://www.maxta.com>.

