

# Hyperconvergence Fundamentals & the Maxta Advantage

Authored by  
**David M. Davis,**  
Global Trainer & vExpert

Edited by  
**Scott D. Lowe,**  
Consultant & Industry Veteran

Your Guide to the Innovative  
World of Hyperconvergence  
& Maxta's Choice-Driven  
Hyperconvergence Solution

Brought to you by

**maxta**<sup>®</sup>



# Hyperconvergence Fundamentals & the Maxta Advantage

Authored by  
**David M. Davis,**  
Global Trainer & vExpert



Edited by  
**Scott D. Lowe,**  
Consultant & Industry Veteran

**Data Center Innovations Learning Series Hyperconvergence  
Fundamentals & the Maxta Advantage**

Author: David M. Davis, ActualTech Media  
Editors: Scott D. Lowe, ActualTech Media  
Hilary Kerchner, Dream Write Creative  
Kiran Sreenivasamurthy, Maxta

Book Design: Braeden Black, Avalon Media Productions  
Geordie Carswell, ActualTech Media

Layout: Braeden Black, Avalon Media Productions

Copyright © 2015 by ActualTech Media

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review.

Printed in the United States of America

First Printing, 2015

ISBN 978-1-943952-01-4

ActualTech Media  
Okatie Village Ste 103-157  
Bluffton, SC 29909  
*www.actualtechmedia.com*

# Foreword

Hyperconverged infrastructure is a data center architecture that has all of the makings of a revolutionary advancement in how IT provides this business-critical environment. From the promises of lower costs, easier growth and scalability, massively simplified administration, and even the complete elimination of the expensive, complex and monolithic SAN, hyperconverged infrastructure appears to touch on almost all of modern IT's pain points.

However, as in all things that seem too good to be true, the devil is in the details. Different hyperconverged infrastructure vendors provide vastly dissimilar solutions that don't always meet customer expectations. It takes some effort and open, honest, and transparent communication from vendors and industry analysts to help you understand the reality behind solutions.

ActualTech Media's mission is to help consumers understand this reality. We shine the light on today's data center innovations in a way that is positive and accessible. To assist in this learning process, we have developed a brand new program to shine light on the underpinnings beneath these new technology approaches: Data Center Innovations.

In this book, the very first in the series, we are thrilled to have Maxta as the inaugural solution. You will learn about the real reasons that hyperconverged infrastructure is taking the data center market by storm. You will also learn about how Maxta's innovations specifically address the most serious data center pain points and the unique benefits that Maxta brings to the hyperconverged infrastructure space.

Scott D. Lowe  
Co-founder  
ActualTech Media

# Table of Contents

## Section 1: Hyperconvergence Fundamentals

<b>Chapter 1:</b> Introduction to Hyperconvergence	<b>3</b>
<b>Chapter 2:</b> Solving the IT Lifecycle Problem with Hyperconvergence	<b>7</b>
<b>Chapter 3:</b> Simplifying the Data Center with Hyperconverged Infrastructure	<b>13</b>
<b>Chapter 4:</b> Data Protection in the World of Hyperconvergence	<b>21</b>
<b>Chapter 5:</b> Data Reduction and Hyperconvergence	<b>29</b>
<b>Chapter 6:</b> Agility, Availability, & Hyperconvergence	<b>31</b>

## Section 2: The Maxta Advantage

<b>Chapter 7:</b> Who Is Maxta?	<b>39</b>
<b>Chapter 8:</b> Delivering Day One Results	<b>45</b>
<b>Chapter 9:</b> Independent Resource Scaling	<b>49</b>
<b>Chapter 10:</b> A Question of Form (Factor)	<b>53</b>
<b>Chapter 11:</b> Hyperconvergence Pricing and Budgeting	<b>59</b>

## About the Author



**David M. Davis,**  
Global Trainer & vExpert

David Davis is a partner and co-founder of ActualTech Media. With over 20 years in enterprise technology, he has served as an IT Manager and has authored hundreds of papers, ebooks, and video training courses. He's a 6 x vExpert, VCP, VCAP, & CCIE# 9369. You'll find his vSphere video training at [www.Pluralsight.com](http://www.Pluralsight.com) and he blogs at [www.VirtualizationSoftware.com](http://www.VirtualizationSoftware.com) and [www.ActualTech.io](http://www.ActualTech.io).

## About the Editor

**Scott D. Lowe,**  
Consultant & Industry Veteran



Scott Lowe is a vExpert and partner and Co-Founder of ActualTech Media. 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 [www.EnterpriseStorageGuide.com](http://www.EnterpriseStorageGuide.com) & [www.ActualTech.io](http://www.ActualTech.io).

# About ActualTech Media

ActualTech Media provides enterprise IT decision makers with the information they need to make informed, strategic decisions as they modernize and optimize their IT operations.

Leading 3rd party IT industry influencers Scott D. Lowe, David M. Davis and special technical partners cover hot topics from the software-defined data center to hyperconvergence and virtualization.

Cutting through the hype, noise and claims around new data center technologies isn't easy, but ActualTech Media helps find the signal in the noise. Analysis, authorship and events produced by ActualTech Media provide an essential piece of the technology evaluation puzzle.

More information available at [www.actualtechmedia.com](http://www.actualtechmedia.com)





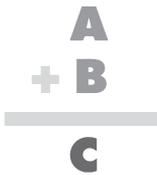
## Maximize the Promise of Hyper-convergence

Maxta's hyperconverged solutions dramatically simplify IT, deliver enterprise-class data services and hyper-scale while enabling significant cost savings.



### Maximize Flexibility

- Any server
- Any hypervisor
- Any storage device
- Any workload



### Maximize Simplicity

- VM centric
- Data Agility
- MaxDeploy™ = MxSP™ + Any Server



### Maximize Savings

- Industry Standard Components
- Simplified Management
- Storage Efficiency

# Maximize

Visit [www.Maxta.com](http://www.Maxta.com)

# Book Features

---



## Real World

How a technical point can impact the real world of IT.

---



## In Depth

Takes readers into the deep, dark depths of a particular topic.

---

## Introduction

It may come as no surprise that technology in the data center is advancing at break-neck pace. However, what you may not realize is that there is a relatively new option for the data center called hyperconvergence. Whether your data center is a few servers under the desk or a few thousand servers in a massive enterprise data center, hyperconvergence has the potential to make massive improvements in efficiency, simplicity, agility, availability, and data protection — all while saving your company real money. It accomplishes this by using virtualization to integrate compute and storage resources into an innovative, complete solution.

Hyperconvergence is unique. Unlike other new enterprise technology that improves performance or capacity in a single area, hyperconvergence offers a wide breadth of benefits in multiple infrastructure and business areas, such as the extending the IT lifecycle.

The term hyperconvergence can be confusing because it is thrown around by different vendors doing very different things. Thus, it's important to not only understand the multiple ways that hyperconvergence helps the infrastructure group and the business, but also how it really does what it does. In other words, you need to know how hyperconvergence truly works.

For that reason, we'll cover the fundamentals of hyperconvergence technology in the first half of this book. In the fundamentals section of the book, we'll explore how hyperconvergence helps the IT lifecycle, simplifies the data center, provides data protection improvements, reduces data storage demands, and improves application availability and business agility.

In the second section of the book, we'll cover a specific hyperconvergence technology – the Maxta storage platform, or MxSP. We'll find out how companies are using Maxta to implement hyperconvergence, scale resources independently, better protect their data, and quickly see improvements in efficiency.

As mentioned previously, don't look at hyperconvergence as just offering improvements in a single area of the data center or the business — it is so much more. Read on to learn how hyperconvergence provides a whole new way to think about your data center technology.

# **Section 1**

## **Hyperconvergence Fundamentals**

# 1

# Introduction to Hyperconvergence

---

It's been called Server SAN, Hyperconvergence, Hyper-Convergence, Hyperconverged Infrastructure (HCI), and software-defined storage... but what truly is hyperconvergence? How did it develop?

## Hyperconvergence History

In the past, applications ran best when they were given their own physical server and own LUN on the SAN for storage. We IT professionals wanted to make the business and application owners happy by providing applications with everything they needed *and* simplify administration.

Then along came server virtualization, and we consolidated as many of our physical servers as we could justify. Server virtualization provided amazing efficiency across so many areas that it actually ended up kicking off today's hyperconvergence revolution. This is because, as the number of virtualized servers grew, it became apparent that the storage infrastructure did not offer the performance required, the utilization of that storage infrastructure was inefficient, and the management of the storage infrastructure was much more painful than it needed to be.

Since that point, many vendors stepped forward to offer storage solutions for the problems triggered by server virtualization. The most comprehensive and most efficient of those solutions to-date hyperconvergence.

## Hyperconvergence Defined

It's obvious that the *-convergence* part of the word *hyperconvergence* means to “bring together.” However, the word part *hyper-* either means “energetic,” or “hypervisor,” or both.

Convergence in the data center first became popular for combining servers and storage into a single offering when VMware Cisco and EMC formed VCE and began offering the converged Vblock solution. However, the vast majority of companies require a much more affordable, flexible, and granular convergence solution. This is where *hyperconvergence* comes in.



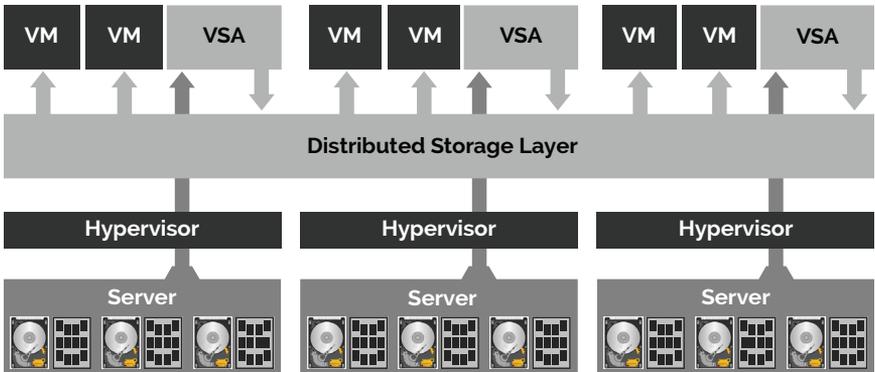
Hyperconvergence leverages the increasing computing power of x86 servers, layers both server and storage virtualization by integrating compute and storage, and offers greater efficiency than was available from siloed infrastructure resources in the past.

The term *hyperconvergence* is attributed to blogger and technologist, Steve Chambers, who in 2012 said that hyperconvergence was an extension of the overall convergence trend, collapsing the data center into an appliance form factor.

Some companies define *hyperconvergence* simply as “a solution that, using virtualization, brings the computer layer and storage layer

together.” You should be aware that hyperconvergence today is much more than that. Yes, hyperconvergence uses virtualization and brings together compute and storage resources in the data center; however, hyperconvergence also should integrate those resources into a complete solution that is much easier to manage, support, and scale. Additionally, true hyperconvergence should offer enterprises additional functionality from the integrated solution over what they had before, such as data reduction, data protection, increased availability, and greater business agility.

## The Hyperconverged Infrastructure



**Figure 1-1:** Hyperconverged Infrastructure

As you see in **Figure 1-1**, a hyperconverged infrastructure, or HCI, is typically comprised of multiple servers running a virtualization hypervisor and distributed storage layer, offering high-availability for compute and storage, ease of scale for both types of resources, and a simplified management console.

Keep in mind that not all hyperconverged infrastructures are created equally. You will want to make sure that your hyperconvergence

solution doesn't just join the resources together, but also does so in a way that provides real business value and additional functionality.

## The Future of Hyperconvergence

Hyperconvergence is a fast-moving target. It's an area of enterprise technology that is developing rapidly. The hyperconvergence solutions of today are quickly changing the way enterprise data centers function, the way they are managed, and how they are planned and paid for. Plus, if the current trends are any indication, hyperconvergence will bring greater and greater value to enterprises by offering more integrated functionality and advanced features. For example, when your company's data is stored in an HCI, that data could be easily deduplicated to reduce storage capacity demands and also easily replicated to a public cloud or another data center for protection. Additionally, future hyperconvergence solutions will bring together not only compute and storage resources but also network resources to create the ultimate hyperconverged data center with the most flexibility and efficiency.

## Up Next

Are you overburdened with with the continuous refresh and complexity of your IT infrastructure? There is an answer. Read on to find out how hyperconvergence can help.

# 2

## Solving the IT Lifecycle Problem with Hyperconvergence

---

Hyperconvergence benefits companies in many ways, one of which is by improving the IT lifecycle. What is the “IT lifecycle”? And why is it so important?

The IT lifecycle, or ITLC, dictates the life of an IT service, from planning and optimizing the service to aligning the service with the business strategy. The ITLC outlines the service design and delivery as well as its ongoing operation and support. It’s important because management of the ITLC is a constant in the enterprise, requiring a dedicated investment of time and money.

While the ITLC should always be aligned with the needs of the business and its critical applications, at too many companies, many aspects of the ITLC have become huge expenses that are disproportional to the demands of the business.

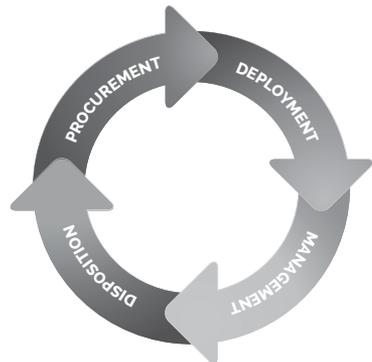


The IT lifecycle, or ITLC, manages the life of an IT service. Examples of IT services where the ITLC applies are desktop computers, applications, or virtual machines. In the case of hyperconvergence, we are discussing the ITLC related to datacenter servers and storage.

The ITLC is broken down into the procurement, deployment, operation, and scale or disposal of hardware and software. Let's examine the ITLC in the enterprise data center specifically and how hyperconvergence can make huge improvements, across each of these areas.

## Improving Procurement

The typical IT department is in a state of continual hardware and software procurement (see **Figure 2-1**). Initially, IT pros enjoy “spec'ing out” new servers and storage; however, as the number of servers increases, so does the number of storage arrays and the complexity. The enjoyable part of procuring new technology quickly disappears when the procurement process consumes many weeks and even months, taking away from other valuable IT activities and causing the business to question the effectiveness of the IT group.



**Figure 2-1:** Procurement cycle

With so many servers and storage arrays (each with different lease schedules) in the data center, the procurement process becomes an ongoing, daily process. In fact, many large enterprises have technology

procurement groups whose sole job is to purchase new data center technology. This type of constant technology refresh is a wasteful use of time and money. The company's only true need is for their applications to run reliably and within a reasonable associated cost.

Hyperconvergence changes the procurement cycle in a few different ways:

- **Elimination of the legacy storage array** — Costly storage arrays are typically procured for a 3- to 5-year cycle because of their enormous cost and complexity. Since the single most expensive resource in the data center is the storage array, hyperconvergence can in many cases, be justified simply because it eliminates the legacy storage array and its costly maintenance contract.
- **Homogenized compute nodes** — As companies purchase more and more servers on different lease cycles, the types of servers in data center will vary. This creates unnecessary complexity and incompatibility for advanced features. With hyperconvergence (using virtualization) in place, compute and storage nodes become a simplified building block for the data center.
- **Build as you go** — Traditional data center procurement usually includes buying a large rack of servers and storage for a single application and for a long period of time. With hyperconvergence in place, you only purchase the servers (compute and storage) that you need. As the data center grows, you add the new capacity you need, thus simplifying the procurement cycle.

## Improving Deployment

Once procurement is completed, the deployment process is the next important phase of the ITLC. In deployment, the servers and storage that were procured are installed in the data center. In most cases, de-

ployment of new server and storage infrastructure can take from a week to a month, depending on the scope of the deployment. The reason that traditional infrastructure deployments consume so much time (and money) is because multiple teams are involved (server and storage) due to the challenging configuration complexities and compatibilities that must be resolved between server and storage infrastructure.

Consider for a moment how much this deployment process could be improved if, instead of using complex server and storage infrastructure for each application, you had a flexible hyperconverged infrastructure that could share the compute and storage load for all applications running in the data center. So exactly how could the deployment of new infrastructure be simplified?

In a hyperconverged infrastructure (HCI) deployment, new nodes are simply brought online to increase data center compute and storage capacity as new applications are added or as application usage increases. With hyperconvergence, data center deployments (meaning, replacing hardware or adding additional capacity) can be performed quickly, easily, and non-disruptively.

## Improving Operations

For the admins and teams operating and administering the data center on a daily basis, hyperconvergence will change their life the most. By being able to reduce management interfaces down to just one interface for management of storage, servers (virtual infrastructure), and perhaps other applications (like backup and disaster recovery), administrators will be able to reduce the amount of time that it takes them to perform common administrative tasks. Fewer management interfaces will also allow the IT infrastructure team to be more proactive by identifying

problems before they happen, and to solve problems in the virtual infrastructure faster by reducing the time it takes to remediate issues.

## Improving Scale and Disposal

The last phase of IT lifecycle management is to either scale up the infrastructure that you have or to reduce/dispose of unneeded infrastructure. Also, part of this phase includes the continuous refresh of the infrastructure we discussed earlier.

Capacity monitoring of any data center resources is also a critical function that must be done to ensure that your company's applications get the resources that they need. Today, legacy storage and server resources are a struggle to monitor, because they have separate monitoring interfaces and use different “measuring sticks.” It's challenging to monitor capacity in the storage infrastructure because storage is typically broken out into LUNs with each LUN having different performance and redundancy characteristics.

With hyperconvergence in place, the scale, disposal, and technology refresh of the storage and server infrastructure can be much more efficient. Hyperconvergence eliminates the LUN and RAID striping paradigm and makes all storage resources a single shared pool. Thanks to hyperconvergence eliminating the legacy storage array, storage just becomes a building block of the server infrastructure. As the compute pool grows, so does the storage pool (and there are ways to grow storage or compute at different rates, if needed). By virtualizing the servers and the storage, hyperconvergence makes the refresh of those resources easier to do since applications are no longer tied to hardware.

# The Future of the IT Lifecycle

Due to legacy hardware and software designs, lifecycle management of IT infrastructure has grown out of control at most companies, causing great inefficiency, unchecked costs, and application delays. Hyperconvergence creates vast improvements in efficiency for IT infrastructure groups because it simplifies the IT lifecycle across all phases, including procurement, deployment, operations, and scale/disposal. It does so by eliminating legacy storage resources, making storage easier to scale, capacity easier to manage, and simplifying infrastructure management. With hyperconvergence solutions becoming more efficient to deploy, manage, and scale every day — consider how hyperconvergence could change the future at your company.

## Up Next

Why is storage so complex to administer? Why are enterprises forced to buy more storage and compute than they really need? And why does everything in the data center have to cost so much? Hyperconvergence can help to solve all of these problems. Read on to find out how.

# 3

## Simplifying the Data Center with Hyperconverged Infrastructure

---

If you think about it, most of the technology you select you do so because it's the most efficient, easiest to use, and simplest available to accomplish the tasks that you require. For example, consider your thought process while shopping for a new smart phone: You would choose the device that has the best user interface, is the easiest to use, allows you to do everything you need to do, and has an affordable price tag. However, in the enterprise data center, technology decisions are tremendously more complex. We have complex server, storage, network, and virtualization layers that require experienced administrators. There are interoperability issues between the silos, the software constantly needs updating, the infrastructure constantly needs refreshing, and we've purchased copious vendor solutions in place to try to solve an overabundance of issues.

So why does enterprise technology have to be polar opposite in complexity to consumer technology? Today, the answer is that enterprise technology is becoming more and more simplified, thanks to hyperconvergence.

Let's examine 4 ways that hyperconvergence will simplify your data center.

## #1 Right-Sized: Scalability in the Scale You Need

Today's data centers scale in multiple ways, but rarely do the scalability options meet the needs of the business. Resources often scale in chunks that are too large, and different resources scale in different ways. For example, servers and storage don't scale at the same rate. Storage arrays scale in huge chunks. Servers need to scale evenly across a cluster to keep the load equal. One storage LUN may be out of capacity or I/O where others have room and IOPS to spare. Or perhaps you have to purchase a whole new SAN just to get the I/O throughput you need.

Your company has applications and data to run and store in the data center. Shouldn't you be able to simply add those applications and increase that data without it feeling like you're fighting an evil game of Tetris or balancing a house of cards?



Scale Out is when you add more nodes to a cluster. Scale Up is when you add more resource capacity to existing nodes.

Hyperconvergence allows you to create a single pool of server and storage resources that can be consumed by applications and data. When that pool is running low on any of the server or storage resources (CPU, memory, storage capacity, storage throughput/IOPS) simply add **another node** to the cluster to increase resources, or “scale out.” Let’s say, for example, that you are running out of storage throughput (I/Os per second); you would add another node to the hyperconverged infrastructure (HCI) cluster. In the hyperconverged world, that additional node would provide more CPU, memory, storage I/O, and storage capacity. In this way, hyperconvergence nodes can be easily scaled up, similar to LEGO® bricks, such that you can simply add more nodes to the data center as you need more capacity (shown in **Figure 3-1**).

Some hyperconvergence solutions also give you the option to “scale up” by adding more storage / disks to the existing servers. This is a good option to consider if your resource constraints are just at the disk capacity or I/O level.

With hyperconvergence, you’ll finally be able to scale your data center in “bite-sized pieces” that make sense for you and your company.



**Figure 3-1:** Simple Scaling of Hyperconverged Nodes

## #2 Flexibility: Not Tied to a Single Storage Vendor — Choose Your Own Hardware

Traditional legacy storage arrays are costly and complex, but they are also proprietary and inflexible. With a legacy storage arrays, when you need to add more disks, disk drawers, or controllers, you must purchase hardware from your storage vendor, typically at a very high cost. Some hyperconvergence solutions sell only hardware appliances. What if you could eliminate both of these problems?



Hyperconvergence allows you to pick the lowest cost hardware vendor, this month, and add new server/storage infrastructure. This means no more lock-in and finally gives IT the upper hand in negotiations!

With the right hyperconvergence solution, you'll be able to bring your own hardware and scale the solution to fit your needs. This “BYO-hardware” flexibility offers:

- **No vendor lock in** — Use any server (with built-in storage) that your applications need.
- **Use a mix of servers** — A heterogeneous cluster eliminates the requirement to make all servers the same in the HCI cluster.
- **Lower cost** — This is due to the option to use the lowest cost server hardware available (and eliminate your costly SAN/NAS vendor in the process).

## #3 Simplicity: Make the Data Center Easier to Manage and Prevent Downtime in the Process

As discussed earlier, when you choose consumer technology, you search for the solution that's easiest to use. Now more than ever, simpler solutions are available in enterprise technology. Hyperconvergence makes your life easier and prevents downtime in the process by:

- Eliminating multiple management interfaces and moving toward a “single pane of glass.”
- Eliminating the storage management burden by eliminating complex storage infrastructure.
- Eliminating data center complexity by folding storage into the compute layer and providing the ability to manage a VM.
- Reducing the Mean Time To Repair (MTTR) your data center infrastructure, when it breaks. As the infrastructure has been simplified, management is easier, infrastructure troubles are prevented, and problems, when they do occur, can be solved much quicker.

## #4 Cost: Lowering the Data Center Cost of Ownership

There's no doubt that you'll lower the total cost of ownership (TCO) in your data center by having the flexibility to bring your own hardware (and not be required to use specific hardware) but there is more to the cost savings of hyperconvergence. In fact, all of the simplifications listed above help to reduce costs.



Lowering your companies total cost of ownership doesn't just help the business. Lower TCOs mean that IT may have more money to spend on other infrastructure or on technology that delivers greater profitability for the business. In many cases, lower TCO in the datacenter could mean that there is more money available for bonuses or profit sharing for IT staff.

Costs are also lowered by:

- Right-sizing your scale so you won't over-buy or over-provision on your data center infrastructure.
- Eliminating your dedicated and proprietary storage infrastructure and its associated maintenance contract.
- Being able to use existing hardware and select future hardware for the hyperconverged infrastructure based on the lowest cost instead of being required to purchase your hardware from your hyperconvergence vendor.
- Reducing management and monitoring interfaces, meaning you'll be able to prevent infrastructure problems before they happen and troubleshoot infrastructure problems, if they should happen, faster.

## The Future of Simplicity in the Enterprise Data center

By simplifying the data center infrastructure using hyperconvergence, fewer administrative hours will be spent administering and

troubleshooting. The end result is a streamlined, simplified, and more efficient data center infrastructure, which ultimately benefits the entire enterprise, its applications, its end users, and its profitability.

## **Up Next**

Enterprises purchase and deploy numerous data protection solutions in the typical enterprise data center. Why does data protection have to cost a small fortune and require so much time and effort? Hyperconvergence simplifies data protection and lowers cost. Read on to find out how.

# 4

## Data Protection in the World of Hyperconvergence

---

Just as hyperconvergence simplifies the overall data center infrastructure, hyperconvergence also vastly improves enterprise data protection. The first step, for those who aren't using virtualization or who haven't attained 100% virtualization, is to virtualize your servers. Simply virtualizing physical servers and running them in the virtual infrastructure will improve your data protection options tremendously (including snapshots and clones) as well as offer more efficient backups and replication.

Additionally, as hyperconvergence virtualizes storage by folding it into the compute layer, you can immediately gain significant improvements by virtualizing the storage layer with intelligent software. It's this intelligent storage software that provides so many of the advanced storage features within the hyperconvergence design. This chapter will examine several of these features and benefits.

# Snapshots

Virtualization administrators love snapshots. Up until recently, enterprises have had only two snapshot options: virtualization layer snapshots and storage snapshots. Each of these snapshot options has its own set of benefits and drawbacks.

In the virtualization world, the virtualization layer snapshots allow you to “take a picture” of a virtual machine’s (VM’s) virtual disk file, including its operating system, applications, and data. Snapshots are used to preserve the state of a VM before an upgrade, installation, or configuration change. Snapshots at the virtualization layer are per-VM, but they are an inefficient use of server resources and can consume great amounts of disk space to store.



Snapshots, clones, and backups are all point in time states of your data so why do you have to use different software solutions for each resulting in tremendous data inefficiency? Look for hyperconvergence solutions to combine all of these “pictures” of your data and maximize your storage efficiency.

Storage arrays also offer their own flavor of snapshots where an entire LUN on the array is preserved at a particular point in time. When that LUN is used for virtualization it will likely contain many VMs. Thus, storage snapshots aren’t as granular or as useful as snapshots taken in the virtualization layer. However, snapshots taken at the storage layer are typically are more high-performance than virtualization layer snapshots.

With hyperconvergence, enterprises can get the best of both worlds. Hyperconvergence snapshots can be per-VM *and* high performance.

When you want to remove the snapshots, hyperconvergence snapshots can be deleted without a performance penalty. Plus, these snapshots can be utilized without impacting capacity or performance.

## Clones

Like snapshots, clones have been available both in the virtualization layer and the storage layer for some time. Clones allow you to quickly duplicate and then reuse virtual machines to create new VMs (or groups of VMs) based on existing deployed VMs.

While infrastructure admins love clones, like snapshots, the determining factor of quality always comes back to the design of the cloning mechanism. How the cloning mechanism's designed determines the performance of the cloned VMs as well as the storage capacity required to keep the clones on disk. You should always understand if full clones are being used or, if not, how clones are linked to the original and what the performance impact will be of using clones in production.

As with snapshots, clones in a hyperconverged infrastructure provide the best of both the virtualization and storage worlds, but without the performance impacts. Clones created in a hyperconverged infrastructure can be created on a per-VM basis, no actual copying has to be done (so there is no performance impact to create them), and fully-cloned VMs can be instantly used to create new VMs in the virtualization layer.

## Backups

The introduction of virtualization into the data center provided massive improvements in how critical data could be protected with virtualization backup being a major piece. Virtualization backups give

you the option to back up an entire VM (guest OS, applications, and data) and restore either the entire VM or just individual files inside. However, virtualization backups are traditionally done with additional tools that provide varying degrees of efficiency in the protection of the backup data, storage of the backup data, and how it's protected offsite.

In the world of hyperconvergence, many solutions still recommend that you use whatever virtualization-aware backup and recovery software you would have historically chosen, (if you already use virtualization, it can be the same tool you have in use). However, some hyperconvergence solutions offer integrated backup and replication (which we'll talk about in the next section). By integrating backup and replication into the hyperconvergence layer, you will gain far greater efficiencies and performance than any add-on tool ever could.

When it comes to data protection, most enterprises have virtualization layer snapshots, storage snapshots, virtualization backup data repositories, offsite replicated data, and more. Why do you have to use so many different systems to protect your data, and why do those tools have to create so many inefficient copies of your data?

Increasingly, hyperconvergence solutions are able to protect your data more efficiently. For example, what if your storage layer could provide you the ability to snapshot your VMs, and those snapshots could be used for point-in-time recovery onsite and automatically replicated offsite for disaster recovery protection? After all, it's the storage layer that has the most efficient access to your data and offers the most efficient storage of that data.

It's innovations like these that you should look for when considering hyperconvergence solutions.

# Replication

To protect your data, you need multiple copies of your data onsite in your local data center and offsite in a remote data center. Traditionally, this was done with RAID on a SAN (for local redundancy) and SAN-based replication (for offsite redundancy). Today, SAN storage accomplishes this with great complexity, inefficiency, and expense when compared to the hyperconvergence alternatives.

Whether it's a snapshot, clone, back up, or replication of a virtual machine, all of these different forms of data protection are really just creating multiple copies of your data. Consider for a moment how much easier the job of an administrator would be if all of these similar data protection functions could be done intelligently, in software, without any performance impact and without ever creating multiple copies of the data.

You should look for hyperconvergence solutions that will allow you to create per-VM policies to control both local and remote data protection, including replication.

# Data Integrity

When it comes to data protection, an important piece that is all too often left out is data integrity. Data integrity is a mechanism to ensure that the data hasn't been modified or corrupted on disk, as it moves between nodes, or as it is replicated across a network. Virtualization administrators typically assume that data integrity is taken care of by their SAN.

Hyperconvergence gives you the opportunity to do a ground-up rethinking of the data center and an important piece of that is data

integrity. What if your hyperconvergence solution included strong end-to-end data integrity for all of your data, no matter where it lives?

When considering hyperconvergence options, you need to find out what is included in terms of data integrity to protect your critical data.

## **The Future of Data Protection**

For infrastructure administrators, there may be no more critical task than protecting the company's data. Traditional storage solutions limit how you can protect your data, increase complexity, and carry a huge price tag. The virtualization layer helped improve data protection in many ways but didn't go far enough. Today and in the future, newly innovated and evolving hyperconvergence solutions finally provide the data protection options that companies need and do so without the performance impact, limitations, and high costs that they've experienced in the past.

## **Up Next**

Data growth in the enterprise is out of control and results in skyrocketing data storage costs. What can hyperconvergence do to slow or even reduce storage capacity requirements in the enterprise? Read on to find out how hyperconvergence can help.

# 5

## Data Reduction and Hyperconvergence

---

For many years, traditional storage arrays have handled the storage capacity and the storage I/O throughput demands of enterprise applications. In most cases, data reduction is an afterthought, something done once the data is already written to disk during a scheduled deduplication process. It's also common not to be able to run deduplication processing because you must purchase a separate license key at additional cost. Finally, some companies who have the option to perform in-line deduplication simply don't do it because they are afraid it will cause I/O latency and application slowdowns.

However, in addition to extra cost, one of the problems with post-processing deduplication is that the I/O writes have already occurred, and the storage capacity has already been used. Thus, it's inefficient to write data that didn't have to be written, and it's inefficient to consume disk capacity that didn't have to be consumed.

In this chapter, we'll examine how hyperconvergence can perform data reduction in different ways, providing greater efficiency, performance, and cost savings. These data reduction processes

can be built into hyperconverged storage and can be done in real-time, not after-the-fact and not at an additional cost.

## Compression

The first way hyperconvergence can help perform data reduction is through data compression. Compression identifies the redundancies within a file or block and eliminates these statistical redundancies. Unlike deduplication, compression doesn't do this on an entire data set; instead, it works by one block or a file at a time.

Compression requires “deflating” and “inflating” (compressing and uncompressing) CPU processing capacity, and this processing can introduce latency. Thus, the data reduction benefits of compression must outweigh the potential latency increase (or that increase must be mitigated). Innovative hyperconverged storage systems can perform data compression in-line, in real time, and mitigate any potential latency increase to perform the compression and decompression with caching.

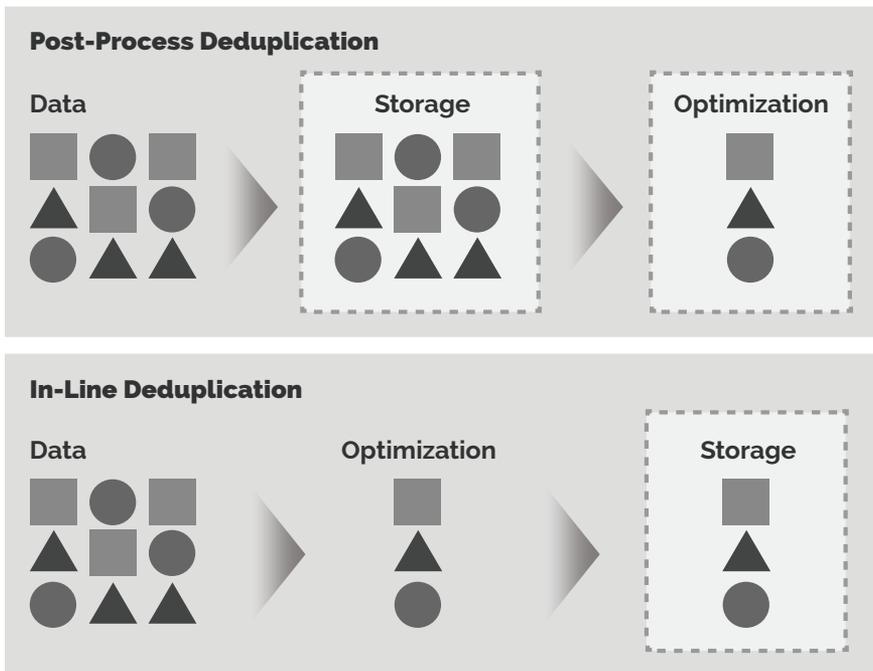
## Deduplication

The second way in which hyperconvergence helps data reduction is through deduplication. Deduplication is simply the removal of redundant data over a larger data set (as compared to compression). You only have to picture the massive duplication in OS and application data between all your virtual machines to begin to understand the types of data reduction possible with deduplication.



Compression is used to reduce the size of small amounts of data, like compressing a file. Deduplication is used to reduce the size of very larger data sets, such as across an entire storage array.

The great thing about deduplication is that, with the most common types of data, the more your data set grows, the greater your deduplication ratios (and data reduction savings). This is because of the greater chance of data redundancies in the data set.



**Figure 5-1:** In-Line Deduplication vs. Post-Process Deduplication

As shown in **Figure 5-1**, the most innovative hyperconvergence storage systems perform this redundant data deduplication in real-time, in-line,

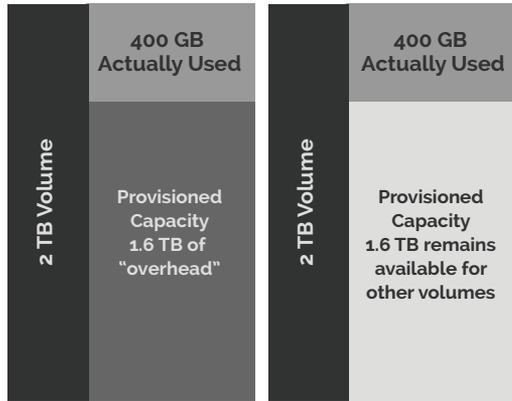
before the data unnecessarily consumes throughput in the I/O path and capacity on disk.

You might think that I/O latency would be increased, similar to compression; however, hyperconverged storage systems can mitigate any potential increase in latency from in-line duplication via caching.

## Thin Provisioning

Technically, thin provisioning is not a data reduction technique even though many people associate it with data reduction. Thin provisioning doesn't actually reduce the amount of data. Instead, it ensures that a disk (or a set of disks / LUN) isn't unnecessarily consumed by preallocating the storage device. In other words, the amount of data stored on the device consumes the same amount of disk space, but a thin-provisioned volume will only use capacity from the storage device as it's needed (**Figure 5-2**).

Innovative hyperconverged storage systems include thin provisioning to ensure that storage capacity is only consumed when needed.



**Figure 5-2:** Thin provisioning in action

# The Benefits of Data Reduction

It's easy to see how the data reduction options offered by hyperconvergence could save you storage space, but there are more benefits.

Data reduction provides you with:

- **Improve performance** — By having to write fewer I/Os, you make the storage path more efficient and, ultimately achieve higher performance. Thus, there is more I/O throughput available for production I/O, especially for bursty I/O. Further, the more you can reduce, the more that you can fit into fast storage caching layers in the hyperconverged cluster. This means that you can accelerate more workloads, leading to even better overall performance across the cluster.
- **Lower costs** — Less capacity used means less expensive storage capacity you have to purchase.
- **Data management efficiency** — With less capacity utilized in the storage, other data-related activities become more efficient such as backup and replication.

## The Future of Data Reduction

Integrated data reduction features in hyperconverged storage allows enterprises to perform data reduction as efficiently as possible in real-time, in-line, and built-in. With hyperconverged storage and data reduction in place, enterprises will be able to make most efficient use of their storage investment, now and in the future.

## Up Next

How can you meet end user demands for the “always-on data center”? How can you deliver the applications and capacity that end users expect, just like their web-apps and smart phone? Hyperconvergence helps you do this by increasing agility and availability in the enterprise data center. Read on to find out how.

# Agility, Availability, & Hyperconvergence

---

Enterprise IT groups are pushed to their limits to meet the demands of unlimited scalability and the 24/7 x 365 availability of Internet-based software as a service (SaaS) applications that their end users use daily. Enterprise data center groups are looking for ways to meet these expectations without losing control by moving to the public cloud, and do so using technology that their company can actually afford.

Traditional storage wasn't designed to have unlimited agility and affordability, it requires dedicated storage teams with years of training. To achieve availability with traditional storage, the often cost-prohibitive and complex answer is to simply buy redundant storage arrays. Traditional storage systems just can't offer the agility and availability that modern enterprises need because of their complex and monolithic designs.

Enterprises need modern, innovative storage options that can give the data center the agility and availability that it needs.

How can hyperconvergence help?

# High Availability

You may recognize you need complete data availability, but, unfortunately, may not be able to afford the complexity or cost required for high availability in the traditional storage world. With hyperconvergence, you don't have to. With hyperconvergence, storage is virtualized and distributed across the server layer. But there is more to it than that.



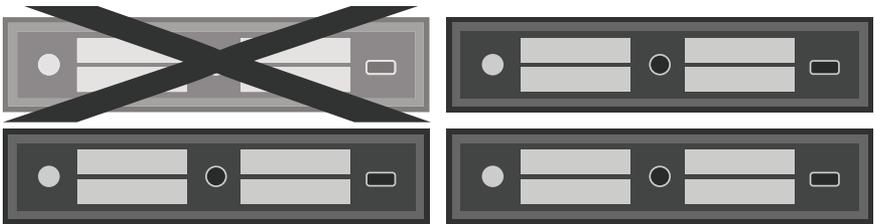
Many companies have high availability in their compute layer thanks to VMware's vSphere High Availability (HA) however most companies don't have complete storage high availability with fully redundant SAN or NAS arrays. Complete storage high availability is something that most hyperconvergence solutions include by default and without additional cost.

Hyperconvergence leverages server virtualization, which makes virtualized servers hardware-independent and includes built-in high availability for the compute/server layer. Hyperconverged storage systems include the high availability storage characteristics that traditional storage systems lack (**Figure 6-1**). For example:

- **Highly available storage is part of the design** — Hyperconverged storage systems not only utilize the local storage on the virtual hosts, but they also synchronously replicate data across multiple hosts, in real-time, to ensure that data is protected and highly available should a server (and its associated storage) fail. Because of this design, highly available storage is included when hyperconvergence is implemented. High availability is also included in the cost (unlike in many traditional storage arrays).
- **Fault domains are reduced** — If a traditional storage array fails there is a vast impact-and-rebuild time. With hyperconvergence,

if a server providing storage in the cluster fails, only a small amount of the data has to be redistributed to the other hosts.

- **Per-VM availability** — With traditional storage arrays, you specify performance and availability characteristics on a LUN basis. If you want to change a VMs storage characteristics, you have to move it to another LUN with different characteristics. On the contrary, with hyperconvergence, you have one large storage pool and you can specify performance and availability characteristics through policies. Hyperconvergence also features the ability to make dynamic changes to these per-VM storage policies and have the underlying storage layer immediately adjust.



**Figure 6-1:** Storage and Compute High Availability in Hyperconvergence

## Scalability

As data grows in the enterprise, you need to be able to add more storage capacity, more storage I/O throughput, or both. In many cases, what you really need is just to add storage I/O throughput, independent of capacity, but traditional storage design doesn't allow it. The result is that you must buy much more traditional storage than you need just to obtain the I/O throughput required. Unfortunately, the traditional storage that you are forced to purchase comes in an expensive, complicated, and over-sized “block.”

What your enterprise really needs is the ability to scale I/O throughput independent of capacity and to be able to add more storage (be it

capacity or throughput) in small blocks — blocks that are exactly the size you require.

Hyperconvergence provides granular scalability. You can add more servers (compute) to increase processing for virtual machines (VMs). You can add more storage (in servers or attached to servers) as greater capacity is required. Finally, you can add flash storage, as more storage I/O throughput required. All of these resources are independent and can be easily scaled as the business needs change.

## **VM-Centric Storage Services**

Hyperconverged storage systems create a single pool of shared and distributed storage. The easiest way to picture this is to think of the traditional Windows desktop PC that has a single local “C:” drive. With a single drive, it’s easy to monitor storage capacity and throughput. There is no complexity related to multiple storage LUNs with different RAID characteristics and different caching levels.

Hyperconvergence, integrated with a virtualization hypervisor like vSphere, allows you to create storage policies for VMs, specifying the quality, availability, and performance characteristics required. The result is that the burden of storage administration is eliminated and VMs (running business critical applications) get the storage characteristics they need.

## **The Future of Agility and Availability**

Hyperconvergence, with its software-defined storage layer, simplifies storage management, includes highly available storage, allows for granular scaling of resources, and provides VM-centric storage. By

combining powerful x86 server, virtualization, and distributed storage, hyperconvergence gives the enterprise the agility and availability needed to handle current and future demands.

## **Up Next**

We've learned a lot about what hyperconvergence is and how it can help. But what does hyperconvergence look like in practice? Now it's time to look at a specific hyperconvergence solution from Maxta. Read on to learn about Maxta's hyperconvergence solution and what makes it unique.

# **Section 2**

## **The Maxta Advantage**

You've learned what hyperconvergence is and the many ways it can benefit those who administer, design, and utilize the modern data center. Additionally, you've learned how the incredible efficiency in design and management of hyperconvergence over traditional storage options translates into real financial savings when companies adopt it. But hyperconvergence isn't just a "theoretical unicorn", hyperconvergence is being implemented at break-neck pace in data centers around the world. No doubt you'd like more details about specific hyperconvergence solutions and how they can help you. To that end, this section of the book will cover Maxta — a top provider of hyperconverged solutions and a leading innovator in the space. Read on to find out who Maxta is and what advantages they can bring to you and your company.

# 7

## Who Is Maxta?

---

Recently named one of the coolest vendors in storage by the Gartner Group, Maxta is on its way to becoming a household name in the data center. Maxta was founded in 2009 to help enterprises around the world realize the benefits of hyperconvergence. Since then, Maxta has been supported by top-tier investors who admire their innovative solution and believe that enterprises will, too.

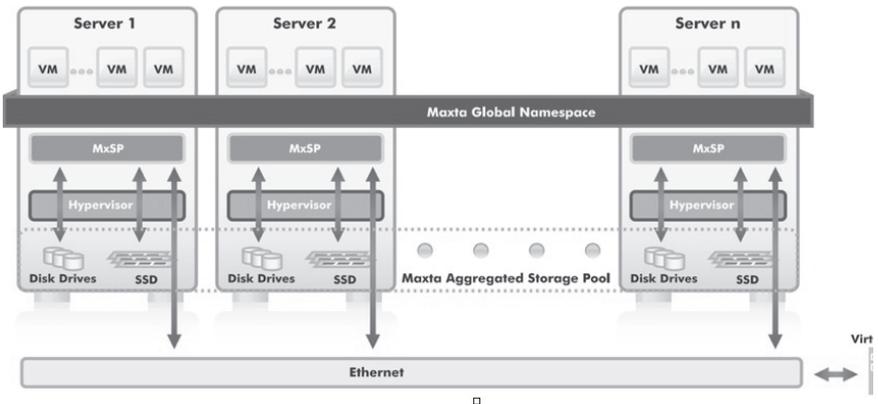
Unlike traditional storage solutions, Maxta's solution is flexible, simple and affordable, while still providing you the same advanced storage functionality. Here's what Maxta's Storage Platform, known as MxSP, will provide you with:

- Flexible virtual storage that can run on across any industry standard servers.
- Elimination of your traditional storage array and storage network.
- Dramatically simplified storage administration.
- Enterprise-class storage with per-VM management and visibility.

- Hybrid storage that can be optimized for performance using flash storage, for tier-1 applications, but also provides the capacity needed with spinning disk capacity for lower tier applications.
- Hyperconvergence of compute and storage on commodity servers to manageably achieve the software-defined data center vision.

## How Does the Maxta Storage Platform (MxSP) Work?

MxSP creates an aggregated virtual storage pool across using the local storage of your existing commodity servers (shown in **Figure 7-1** below). With availability, resiliency, end-to-end data integrity, data protection, and virtual machine (VM) awareness, MxSP provides features beyond traditional storage and allows you to eliminate your traditional storage array/silo.

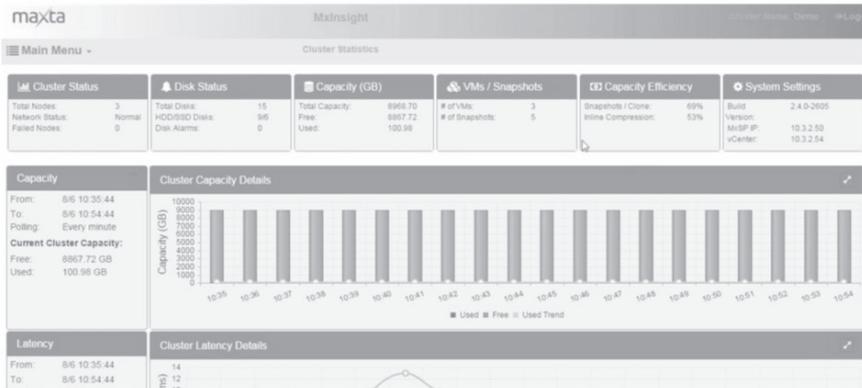


**Figure 7-1:** The Maxta Storage Platform, or MxSP

Here are the types of advanced storage features offered by the Maxta Storage Platform:

- **High Availability and Resiliency** — VMs are distributed across multiple nodes in the cluster to ensure that if any node is lost, other nodes immediately take over and redundant copies are automatically rebuilt on other nodes in the cluster. MxSP offers best-in-class resiliency and high availability with end-to-end data integrity and no single point of failure. With Maxta Metro storage cluster, MxSP can sustain even datacenter/rack failures and still serve data.
- **Data Protection** — With the ability to create an unlimited number of virtual machine snapshots via the storage platform (not hypervisor snapshots), MxSP provides an integrated backup and recovery option for virtual machines. Additionally, MxSP configured in a stretch cluster configuration with VMs replicated across a metro-area network to an alternate data center can protect against data center failures.
- **Per-VM Storage-Awareness** — With the ability to instantly clone VMs or provision new VMs, MxSP allows admins the ability to tune their storage based on the needs of their VMs and the associated applications. New VMs can be instantly provisioned or cloned (based on existing VMs) in seconds, because there is no data to copy. Additionally, MxSP allows you to perform per-VM tasks in the MxSP console or you can configure MxSP from the vSphere (or another hypervisor) centralized management console (which would be vCenter in the case of vSphere).
- **Scale out expansion capabilities** — The individual nodes in a Maxta cluster operate on the robust ZFS file system, which is not known for its scale out capabilities. MxSP imbues each node with an overlay file system that instantly transforms the individual ZFS-based nodes into scaling powerhouses that can allow organizations to expand at will and without having to worry about hitting inbuilt limits.

- **Simplified storage management** – The entire Maxta storage cluster as well as common virtual machine administration tasks can be performed from the Maxta MxInsight web-based administrative console, shown below in **Figure 7-2**. This administrative console is also accessible through the vSphere Web Client.



**Figure 7-2:** Maxta’s MxInsight Administrative Console

In summary, MxSP is a tier-I completely virtual storage system that is so fully featured that it allows companies to eliminate their existing storage array and storage area network. With the storage array eliminated, a tremendous amount of IT spending and excess resources can be eliminated. The end result with hyperconvergence in place is an easier life, both for those who administer the data center infrastructure as well as those who pay for it. You’ll learn more about the economic impact of hyperconvergence later in this book.

## Up Next

Thankfully there are a lot of innovative storage systems on the market today, compared to when most companies bought their SAN. What really makes Maxta unique in the hyperconvergence world? Let’s find out more about how Maxta differentiates itself by delivering day one results.

# Delivering Day One Results

---

Many IT innovations promise results. With enough capital investment and time, most solutions should be able to provide some kind of positive result. However, the most treasured solutions are often hard to find. These are the solutions that don't require massive time or capital investments to realize those results, and Maxta is one of them.

Maxta is one of those rare solutions that is able to provide day one results when implemented across existing hardware or when new technology is released. For example, Maxta was the first hyperconvergence vendor to support several new technology innovations.

In this chapter, we'll examine how Maxta combines with these innovations to deliver day one results.

## Intel Grantley

Intel's new server CPU family dubbed "Grantley" is the Xeon E5-2600 V3 chip. Grantley provides more processing speed, NVMe storage, DDR4 memory, and 40 Gigabit Ethernet networking support. If you

set all those technical specs aside, what Grantley provides to companies are real business benefits such as higher VM density or host-to-VM ratio.



NVM Express, NVMe or Non-Volatile Memory Host Controller Interface Specification (NVMHCI), is a specification for accessing solid-state drives (SSDs) attached through the PCI Express (PCIe) bus.

Meaning that enterprises will get a greater return on their investment (ROI) and lower their total cost of ownership (TCO) in the data center. For example, by combining Maxta and Grantley, companies can achieve 40% more Microsoft Exchange users than before.

Maxta was the first hyperconvergence vendor to not only support Intel Grantley, but release a reference architecture. The Maxta “MaxDeploy” reference architecture is the fastest path for enterprises to deploy the latest technology in a predictable, repeatable, and scalable way. The MaxDeploy reference designs work with Cisco, Dell, HP, Intel (server boards), Lenovo and Supermicro servers. They include Intel PCIe flash products and Intel’s XL710 family of 40 Gigabit Ethernet controllers as well. Maxta offers an online MaxDeploy Appliance Configuration Tool that allows you to specify your use case, preferred server vendor, and amount of storage required in order to obtain suggested server configurations for your hyperconverged infrastructure, such as the one you see in **Figure 8-1**.

## SUGGESTED SERVER CONFIGURATION

SERVER VENDOR:	Cisco	CPU:	8X2603v3 6C 1.6GHz
SERVER MODEL:	Cisco UCS C240 M4	MEMORY IN GB:	4x128
NUMBER OF SERVERS:	4	SSD:	4x2x600GB
CAPACITY IN TB:	48	HDD:	4x6x2TB
		NETWORK:	4x2x10GbE

Cisco UCS C240 M4



**Figure 8-1:** Results from the Maxta MaxDeploy Configuration Tool

## VMware vSphere 6

Another first for hyperconvergence vendors was Maxta's support of VMware's latest vSphere 6 release. On the same day that VMware® vSphere 6 announced general availability (GA), Maxta's Storage Platform (MxSP) and their reference architectures (MaxDeploy) were able to support it. Maxta is able to provide this quick support because of their design.

Maxta's hardware-agnostic and hypervisor-agnostic design means that they will continue to be the leading hyperconvergence vendor. By supporting vSphere 6, you can immediately start leveraging the latest vSphere 6 advanced features and increased scalability.



Maxta's founder Yoram Novick reiterates this commitment: "Our commitment to providing customers the choice of hardware and software they most prefer is reflected in our continued support of the latest technological advancements in the industry, such as VMware vSphere 6."

## OpenStack Kilo

When OpenStack announced their newest edition, version 11, dubbed "Kilo", Maxta announced support for it on day one of Kilo's availability. Enterprises are building more private clouds than ever,

and OpenStack is the solution of choice. Day one OpenStack support is yet another of the many examples that show how Maxta is a leading innovator in the hyperconvergence world.

In summary, when considering hyperconvergence solutions, you want a solution that will allow you to use the latest technology as soon as it becomes available. Having the freedom to use the latest enterprise technology isn't just exciting, it provides your company real business value, such as increasing your VM consolidation ratio and maximizing your data center ROI. Consider choosing a hyperconvergence vendor that is hypervisor and hardware agnostic to give you and your company the greatest ability to achieve day one results.

## Up Next

Achieving the benefits of hyperconvergence as quickly as possible is critical. But what about scale? What if your storage and compute resources need to scale at different rates? Read on to learn how Maxta's MxSP handles independent resource scaling.

# Independent Resource Scaling

---

Many hyperconvergence vendors offer appliance-based solutions that include both compute *and* storage. While this is excellent if both resources scale at the same rate, it's inefficient if, say, you have a large amount of data that consumes all the appliance storage and leaves behind unused compute. Or, vice versa, the more virtual machines (VMs) that use, by comparison, very little storage, means the more storage in appliances that is wasted.

The most common scenario is that an enterprise has a very large dataset, perhaps for a data warehouse or file server, which uses all the storage across the hyperconvergence appliances. Then they are forced to purchase more appliances, simply to add more storage, and the compute resources (CPU and memory) in the appliances is virtually unused. This is an inefficient use of resources and financial investment and it's exactly the type of situation that hyperconvergence is supposed to help companies avoid. That is, why are you using hyperconvergence if you end up with a large pool of wasted resources?

So why can't appliance-based solutions overcome this and how can Maxta help? Read on to find out.

# One Problem with Hyperconvergence — Appliance Limits

Deploying new hyperconverged nodes in the hardware-based form factor isn't a bad approach. In fact, Maxta partners with numerous hardware providers that offer Maxta in their appliances (HP, Dell, Cisco UCS, Lenovo, Quanta, and Supermicro). For the majority of enterprises, appliances are the best approach. The virtualization hypervisor distributes compute and storage workloads across the appliances. Appliance sizes can be scaled up and down to fit the resource requirements of your company's workloads.

While this works well for most companies, the issue with appliance limits is a problem for a minority of companies who typically have storage requirements beyond the scale of the appliance form factors. Unlike a Storage Area Network (SAN) or Network Attached Storage (NAS) array where you “scale up” your array by adding more disk drawers and disks, appliances have only so many disk slots and, thus, can support only so much storage before you have to add another hyperconverged node to your cluster.

Many hyperconvergence solutions also require that you purchase the appliance and associated disks from them. You may not have the option to simply replace server drives with larger “off the shelf” drives, because you must use the drives that only the hyperconvergence vendor supports and sells.

So when you have a large dataset that is beyond the limits of what appliances can offer, what are your options?

# How Maxta Addresses Resource Needs

Too many companies are “holding off” on hyperconvergence because the appliance form factors don’t address their large data set requirements. Thankfully, Maxta has an answer that fits the needs of all enterprises.

Because of Maxta’s hardware-agnostic design, MxSP offers you the following options to scale storage disproportionately to compute:

- **Add drives** — To add more storage capacity to the hyperconverged cluster, you can simply add more drives to your MxSP nodes. For example, you might add 2TB drives in addition to the existing 1TB drives to increase your storage capacity by 200%.
- **Change drives** — To add more storage capacity to the hyperconverged cluster, you can replace existing drives in your MxSP nodes. For example, you might replace 1TB drives with 2TB drives, across the cluster, to double your storage capacity.
- **Scale out** — Add more storage capacity to the hyperconverged cluster, by adding more nodes to the MxSP cluster. Those new nodes don’t have to be identical in resource configuration to the existing nodes. For example, you could add a couple of new hosts to the cluster that were both very storage-heavy, with lots of directly attached storage.

In summary, not every organization is going to be the same when it comes to application and data requirements. It’s not realistic to offer only appliances that scale uniformly in terms of compute and storage, because that approach results in wasted resources at many companies that have heavy storage requirements for their applications. Maxta’s MxSP gives you ultimate flexibility when it comes to the independent resource scaling of compute and storage.

## Up Next

We've learned how to overcome the challenge of independent resource scaling. But what form factor should your hyperconvergence solution take? And why does form factor matter? Read on to learn how Maxta's hyperconvergence solution offers the most flexibility when it comes to hyperconvergence form factor.

## A Question of Form (Factor)

---

You've learned how hyperconverged infrastructure works. You've seen the benefits it can provide to you and your company. Based on all of that, now you're interested in using hyperconverged infrastructure in your data center. So, what shapes and sizes does it come in? How do you purchase it? Read on to learn what hyperconverged solutions are available today and which type is the best.

### **Software-Based Solutions**

Hyperconverged storage and compute is possible thanks to server virtualization solutions, such as the VMware vSphere hypervisor, and software-defined storage solutions, such as the Maxta Storage Platform (MxSP). When you layer software-defined storage on top of server virtualization, you have the foundation of hyperconvergence.

However, as you have learned earlier in this book, there is more to hyperconvergence than just virtualizing the storage. Innovative hyperconvergence solutions should offer many more advanced features within their software, such as a single management point for the hyperconverged infrastructure, compression, de-duplication, and data protection options.

But if hyperconvergence is really “all in the software”, why do so many hyperconvergence companies require you to purchase new hardware to move to hyperconvergence? Good question.

The Maxta Storage Platform is one of the very few hyperconvergence solutions that can be implemented by only purchasing the MxSP software and using your existing hardware or purchasing additional hardware from the server vendor of your choice. The benefits to this are:

- **Cost savings** — By not having to purchase any new hardware, there is an immediate cost savings compared to comparable solutions. Plus, don’t forget the cost savings and complexity reduction you immediately gain by eliminating your SAN or NAS!
- **Flexibility** — Many companies don’t want to have to rip and replace their existing servers that are already on-lease, nor do they want to replace their proven and trusted server brand with a new brand from a hyperconvergence vendor.
- **Acquisition** — Enterprises have worked hard over the years to develop relationships with vendors or resellers for specific products, such as servers. Enterprises pride themselves on knowing the value of their server purchases and knowing that their relationship with the reseller (or hardware vendor) gives them bargaining leverage when it comes time for hardware acquisition. If they were to purchase hyperconverged solutions from a vendor that required them to replace their existing servers, they lose the bargaining leverage, discounts, and relationships that they have worked so hard to build.

Proponents of appliance-based (hardware-based) hyperconvergence solutions bring up the following concerns with software-only solutions:

- **Support** — How can you know that you will receive quality support from a software-only solution installed on your own hardware? There are two parts to this question 1) who will support the server hardware when it breaks? and 2) who will support the storage software? The answer is, hardware support is provided by the customer's existing server vendor, just as it would be without hyperconvergence. Software support of the hyperconvergence layer is provided by Maxta. Overall, this support model will allow each vendor to do what they do best, resulting in better quality and lower cost support for the customer.
- **Performance and Sizing** — How do you know what kind of performance you can achieve? How do you know how to size your hardware? When using your own hardware, you are also accepting that you can perform your own performance sizing and management, over time. If you aren't willing to accept that responsibility, then you should consider the customer-provided appliance approach where you select hardware based on one of the MaxDeploy reference architectures.

## Appliance-Based Solutions

While software-based hyperconvergence solutions offer the ultimate flexibility and the lowest cost by far, hardware-based appliances aren't all bad. Many companies want to be able to purchase an all-new hyperconvergence solution that includes both hardware and prebuilt software. This allows them to build an all-new hyperconverged infrastructure. With an all-new hyperconverged infrastructure in place, they can then migrate from the old infrastructure to the new.

Hyperconverged appliances come in two different forms: vendor-provided and customer-supplied.



A Hyperconverged Appliance is a commodity server, hypervisor, and hyperconverged storage, all in one. Some hyperconverged appliances are offered as a single SKU from a vendor or reseller or, by using a reference architecture, you have the option to create your own.

## Vendor-Provided HCI Appliance

Many hyperconvergence vendors will offer to sell you a “vendor-provided appliance” which is a set configuration of a hardware server (with included local storage), hypervisor, and hyperconverged storage, all in one. Typically, you are required to purchase the vendor-provided appliance from the vendor or through a reseller, and only a few configurations are typically available. Usually, you cannot change the appliance configuration once you own it. If you need more performance or storage, the only options are to purchase new vendor-provided appliances.

By consuming hyperconvergence using the vendor-provided appliance form factor, you are limiting the flexibility you have in your data center to purchase your own hardware and to configure your own hardware however you see fit. The result is that your hardware flexibility goes down and your cost of acquisition may go up.

## Customer Supplied HCI Appliance Based on a Reference Architecture

Another appliance alternative is the “partner supplied” appliance where you are essentially buy a build to order hyperconverged appliance from a partner based on a reference architecture.

A reference architecture is a pre-designed and pre-tested hardware configuration that specifies the supported hardware you can use, the exact configuration, and how many IOPS and/or virtual machines

(VMs) you may be able to run. Maxta offers numerous reference architectures with hardware partnerships from Cisco, HP, Supermicro, Dell, Lenovo, and more.



Think of a Reference Architecture as a blueprint for your hyperconverged infrastructure. You wouldn't build a new house without a blueprint, and you shouldn't build a new hyperconverged infrastructure without a blueprint.

For example, the reference architecture shown in **Figure 10-1** is Maxta's MaxDeploy Cisco UCS reference architecture. As you can see, based on the configuration given, you'll be able to achieve 35,000 – 120,000 I/Os per second using 4K block sizes. With a pre-tested configuration, you know that the HCI design you are using is supported, has been pre-tested, and what your rough performance expectations should be.

In summary, there are two different form factors for acquiring and deploying hyperconvergence in your data center. You can purchase hyperconvergence infrastructure software only and deploy it on your own hardware. Or, you can purchase hyperconverged infrastructure appliances in a few different form factors. Each solution offers its own pros and cons, with the software-based form factor being the least expensive and most flexible option for getting started with hyperconvergence.

## System Configuration Information Key Solution Benefits

Components	Description
Number of servers	4 (Cisco UCS C240 M3 Systems)
Processor per server	2 Socket 6 cores (Intel Xeon® CPU E5-2643 @ 3.30GHz)
Memory per server	128GB DDR3
HDDs per server	12x1TB 7.2K RPM SATA
SSDs per server	2x300GB SATA (Intel DC S3700)
Network configuration	Dedicated 10 GigE port
Maxta virtual machine	4 vCPU/8GB memory
Virtualization products	<ul style="list-style-type: none"><li>- VMware vSphere 5.x</li><li>- KVM in CentOS 6.5</li><li>- RedHat Enterprise Linux 6.x</li></ul>
Management	<ul style="list-style-type: none"><li>- VMware vCenter 5.x</li><li>- OpenStack Icehouse, Juno</li><li>- Maxta Native Web Interface and API</li></ul>
Performance	35,000 – 120,000 IOPS @ 4K Block Size

**Figure 10-1:** Maxta's MaxDeploy Cisco UCS Reference Architecture

## Up Next

We've learned about the three different form factors that hyperconvergence can be consumed in, but what are the pricing and budgeting models for hyperconvergence? After all, you have to pay for hyperconvergence at some point, right? Read on to learn about Maxta's unique pricing model and how to budget for hyperconvergence.

## Hyperconvergence Pricing and Budgeting

---

With hyperconvergence available in so many form factors and deployment methods, you may wonder how much hyperconvergence is going to cost and how you should budget for a hyperconvergence implementation. Read on to learn all about the economic factors related to a new hyperconvergence implementation.

### Pricing Models for Hyperconvergence

You learned that the software-based form factor of hyperconvergence consumption has, by far, the lowest cost and the lowest barrier to entry simply because you don't have to purchase new hardware. As with many software licensees, hyperconvergence software can be purchased in either a subscription or a perpetual license.

#### Subscription-Based Pricing

With the subscription-based pricing model, you pay for what you use. Much like software as a service (SaaS), the subscription model sets a unit of measure, and then you pay the software vendor as you consume that unit. Whereas other subscription-based hyperconvergence

solutions charge you based on the CPU socket, Maxta offers subscription-based pricing based on the amount of storage and the CPU cores that you consume. For example, if you have 12TB of storage and 24 cores then you pay for that amount.



More and more, enterprises are moving to subscription-based pricing to synchronize their IT expenses with their company revenue. If you think about it, subscription-based pricing makes more sense for everyone. Companies should only pay for the services that they consume and have the option to cancel when needed.

Another benefit to subscription-based pricing is that support and software maintenance are built into the price and are not a separate line item. Just as Internet-based software as a service applications have grown in popularity, thanks to their subscription-based pricing, the Maxta MxSP user base has also grown quickly because of this easy-to-consume subscription-based pricing model.

## **Perpetual License Pricing**

The alternative to subscription-based pricing is perpetual pricing. With this model, you purchase the software license, and you own it forever. Typically, however, the support and software maintenance is then purchased separately and renewed each year. The plus side to perpetual licensing is that you can typically use the software as much as you want because you will always own the license. The downside is that perpetual software licensing usually requires a large upfront capital investment.

The design of the pricing model is going to be directly related to how you go about budgeting for your hyperconverged infrastructure.

# How to Budget for Hyperconvergence

Just like any new data center project, enterprises have to figure out how to fit hyperconvergence into their budget. The main decision you have to make is which one of the form factors (discussed in Chapter 10) you will choose as it will have the greatest effect on your overall budget.

The software-only form factor is the easiest on your budget because you will only have to pay for the software. For example, with Maxta, you would only pay for what you use with their MxSP subscription-based pricing model. Additionally, you will be able to eliminate your SAN or NAS, and you won't have to purchase any new hardware. Therefore, the implementation of MxSP likely can save money in your overall budget so that you can use those leftover funds for other technology innovations that have a positive impact on your company's bottom line.

If you want to or need to purchase new hardware, you must choose from the (hyperconvergence) vendor-provided appliance or the partner-provided appliance options. The partner-provided hardware option (based on a reference architecture) will give you the greatest flexibility in the brand and model of hardware that you choose, and thus, give you the greatest flexibility in purchasing and budgeting. By being able to select the hardware from a few different vendors, you'll likely be able to purchase the hardware from your existing vendor or reseller and maintain the highest level of negotiating power (and the lowest possible price).

## Conclusion: Next Steps

### Find Your Entry Point for Hyperconvergence

For every company, there is an ideal entry point for hyperconvergence

that requires the least amount of effort but still makes the most sense. Here are some of the most common:

- **SAN or Server lease replacement** — When you have an expensive SAN, NAS, or server cluster that is up for lease replacement, it's an ideal time to consider hyperconverged solutions because the budget for the hardware refresh/lease replacement will typically provide your company much more for the same investment.
- **SAN/NAS Upgrade** — Many aging SAN/NAS infrastructures can't meet the I/O demands of current applications. When you consider a SAN/NAS replacement, it's an ideal time to consider hyperconvergence.
- **New Tier-1 Application** — Any time your company is about to implement a new application (financial, email, VDI, CRM, e-commerce, or similar) is an ideal time to maximize your infrastructure investment by considering hyperconvergence.
- **Disaster Recovery** — When you are implementing a disaster recovery solution or refreshing your current solution, it's an ideal time to consider hyperconvergence.

## Try Maxta For Yourself

You've learned what hyperconvergence is and that it can benefit you and your company in many ways. You discovered who Maxta is and what makes them unique: they are one of the few software-only hyperconvergence solutions available, and because of that, Maxta's MxSP is the easiest solution to implement, has the lowest cost, and is the easiest to try!

Try Maxta for yourself, for free, at [www.Maxta.com](http://www.Maxta.com) !

# Discover hyperconverged infrastructure and how Maxta's innovative solution can transform your data center

Hyperconvergence is one of the greatest innovations to come to the data center since virtualization. It's transforming the data center with massive improvements in efficiency and simplicity, reducing both the time and cost it takes to run a modern data center. In this book, you'll learn how hyperconvergence works, how it can help you, and how Maxta's hyperconvergence solution is unique. Data center infrastructure no longer has to be complex and expensive. Make your life easier, today and in the future, by learning about hyperconvergence!

## Learn:

- What hyperconvergence truly is and how it can help you
- How hyperconvergence modernizes, simplifies, and streamlines the IT Lifecycle
- Why hyperconvergence is the answer to so many datacenter challenges including data protection and availability
- How Maxta's unique solution can help you get started with hyperconvergence today



ActualTech Media

For bonus content & more  
Visit [www.ActualTech.io](http://www.ActualTech.io)

ISBN 978-1-943952-01-4



Not for Resale