

# Extending Hyperconvergence: Building the Complete Hyperconverged Data Center

Authored by  
**Scott D. Lowe &  
David M. Davis,**  
Partners, ActualTech Media

Your Guide to  
Understanding Stratoscale's  
Full Stack Hyperconverged  
Infrastructure Solution

Brought to you by



**STRATOSCALE**



# Extending Hyperconvergence: Building the Complete Hyperconverged Data Center

Authored by  
**Scott D. Lowe & David M. Davis**  
Partners, ActualTech Media

**Data Center Innovations Learning Series**  
**Extending Hyperconvergence: Building the Complete  
Hyperconverged Data Center**

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

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

Layout: Braeden Black, Avalon Media Productions

Project Manager: Amy Short, ActualTech Media

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

ActualTech Media  
Okatie Village Ste 103-157  
Bluffton, SC 29909  
*[www.actualtechmedia.com](http://www.actualtechmedia.com)*

# Table of Contents

<b>Introduction</b> . . . . .	<b>. xi</b>
<b>Chapter 1:</b>	
Complexity In The Traditional Data Center . . . . .	<b>13</b>
Data Center Complexity . . . . .	<b>13</b>
Vendor Finger-Pointing . . . . .	<b>14</b>
Policy Sprawl . . . . .	<b>15</b>
Data Center Costs . . . . .	<b>15</b>
Poor Acquisition Economics . . . . .	<b>15</b>
Significant Zero-Return-on-Investment Zone . . . . .	<b>17</b>
Constantly Requires Expansion/Attention . . . . .	<b>18</b>
Requires Deeply Skilled Technical Staff . . . . .	<b>18</b>
Zero Direct Business Value-Add . . . . .	<b>18</b>
Data Center Inflexibility . . . . .	<b>19</b>
Incomplete Integration . . . . .	<b>19</b>
Performance Problem Identification Challenges . . . . .	<b>20</b>
Incomplete Resource Control and Insight . . . . .	<b>20</b>
<b>Chapter 2:</b>	
Solving Complexity With the Hyperconvergence Revolution . . . . .	<b>21</b>
Shared Resource Pool . . . . .	<b>22</b>
The End of the "Hamster Wheel" . . . . .	<b>22</b>
Procurement . . . . .	<b>23</b>
Deployment . . . . .	<b>24</b>
Operations . . . . .	<b>24</b>
Scale/Disposal . . . . .	<b>25</b>
<b>Chapter 3:</b>	
The Missing Pieces in Hyperconverged Infrastructure . . . . .	<b>27</b>
The Hypervisor . . . . .	<b>27</b>
The Network . . . . .	<b>29</b>
Complete Integration . . . . .	<b>30</b>
Hardware Flexibility . . . . .	<b>31</b>

Multi-Directional Scale . . . . .	32
Single Point of Management . . . . .	33

**Chapter 4:**

Data Center Desired State . . . . .	35
Hardware Heterogeneity . . . . .	35
Resource Granularity. . . . .	36
Just-in-Time Data Center . . . . .	36
Complete Workload Support. . . . .	37
Virtual Machines. . . . .	38
Traditional Workloads . . . . .	39
Hardware Emulation . . . . .	39
Containers . . . . .	39
Shared Binaries . . . . .	39
Less Overhead = More Density. . . . .	39
OpenStack . . . . .	40
Eliminate Proprietary Components . . . . .	40
Improve Data Center Operations . . . . .	40

**Chapter 5:**

Simplifying and Securing the Heterogeneous Data Center . . . . .	41
Hardware Heterogeneity and Independent Resource Scaling . . . . .	41
Any Vendor, Any Time . . . . .	42
Hyperconverged Infrastructure Can Exacerbate the Problem . . . . .	43
Avoid the Forklift . . . . .	44
“Set It and Forget It” . . . . .	44
Granular Resource Control . . . . .	44
Heterogeneous Hardware as a Key Enabler . . . . .	45
Leverage Existing Hardware Resources. . . . .	45
Reduce Sunk Costs . . . . .	46
Adopt the “Just in time” Data Center . . . . .	46
Extreme Abstraction . . . . .	46
Data Center Hardware and Plumbing . . . . .	47

A Focus on Outcomes . . . . .	48
Abstract It All . . . . .	48
<b>Chapter 6:</b>	
Understanding the Network in the Hyperconverged World . . . . .	49
Missed Networking Needs . . . . .	50
Security . . . . .	50
Workload Isolation . . . . .	51
Service Provider Needs . . . . .	51
Ensuring Security and Performance . . . . .	52
QoS and Dynamic Workload Rebalancing . . . . .	53
Hyperconvergence and Software Defined Networking . . . . .	53
<b>Chapter 7:</b>	
Improving Operations: DevOps in the IT World . . . . .	55
<b>Chapter 8:</b>	
Building The Perfect World. . . . .	57
The Hyperconvergence Vendors Own It All . . . . .	57
Deep Integration . . . . .	58



## About the Authors



**Scott D. Lowe,**  
Partner, ActualTech Media

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).



**David M. Davis,**  
Partner, ActualTech Media

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 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)





# STRATOSCALE

*At Stratoscale, we're focused on how technology can be leveraged to help IT teams make better and more profitable usage of existing infrastructures. We know that data needs are growing at an ever-increasing pace, so we've build a hardware-agnostic and hyper-converged software solution for data centers that facilitates scale-out, simplifies operations and allows your IT infrastructure to keep up with business growth.*

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



# Introduction

Hyperconverged infrastructure has emerged as a force to be reckoned with in the modern data center. As CIOs and IT pros alike tire of what has become the status quo for the data center, they are seeking solutions that are simpler to manage, potentially less expensive to maintain, and that enable expected business outcomes.

However, many of the hyperconverged infrastructure solutions on the market today, while excellent products, lack comprehensive resource integration and don't enable companies to granularly configure individual resources. For these companies, there is a need to, for example, bring the network resource into a hyperconverged infrastructure's computer/storage sphere of integration. Moreover, there is a need for these companies to be able to specify exactly hardware requirements, from the amount of CPU they want to the RAM density to the amount of flash and spinning disk they want in the solution.

In this book, you will learn about what has brought us to this point in the modern data center. You will also learn about where current hyperconverged solutions can be improved upon and how you can fully leverage commodity hardware to achieve your goals.



# 1

## Complexity In The Traditional Data Center

---

Take a walk through a traditional data center and what you'll find is a plethora of hardware and software all deployed and maintained just to support critical business applications. Servers enable the company to run a myriad of business workloads, while storage devices enable the company to hold on to critical data. Network systems then bind it all together. Last, but not least, an abundance of integration points surround these basic resource silos and make the whole data center work the way it's supposed to.

### **Data Center Complexity**

The ideal, at least, is to fuse all of this into one technological entity that runs the company's applications quickly, reliably, and affordably. In reality of the data center, however, is far from perfect. In fact, it's simply broken in many organizations. The more complex a system, the easier it breaks and the more work it takes to keep it operational. Therefore, the modern data center is often the very definition of complexity, and complexity is the antithesis of productivity.

Over the years, well-meaning — often overworked and under-budgeted — IT staff have crafted solutions to meet the immediate demands of the business. Unfortunately, time and budget constraints often force these same people to simply “make it work.” As a result, the long-term impact of these “make-it-work” solutions may not be exactly favorable.

In fact, even solutions that were put in place with proper time and financial consideration may, in the long term, begin to actively work against the data center’s ability to remain flexible. For example, if such solutions were put into place to solve a specific point need, they can eventually hinder primary resources by introducing additional integration and policy points, and ultimately impacting the budget because of their maintenance needs.

The fact is that not every solution works perfectly with every other solution. Anything that has to be cobbled together to “make it work” will certainly require more maintenance and will be less reliable than something designed to work seamlessly, right out the box.

### **Vendor Finger-Pointing**

One result of data center expansion has been the associated expansion of the number of vendors providing data center services. As you acquire services from an increasing number of vendors, there may come a point when some of the vendors, rather than actively helping you solve problems, begin to blame other vendors whose products you’re using with the first vendor’s products. This vendor finger-pointing creates the potential for extended and expensive downtime.

## Policy Sprawl

Most companies have some level of IT policies in place. For example, you may have storage policies that address how much storage capacity a single virtual machine (VM) can consume. You may have another policy on a deduplication device that places some kind of preference on specific storage volumes. Also, you might have different replication / disaster recovery policies on across different groups of VMs, offering varying levels of protection.

Here's the problem: in many data centers, the sheer magnitude of different devices and services makes it impossible to achieve perfect alignment in policies. Administrators are forced to use different management interfaces and attempt to align policies with different resources.

In the example above, one policy targets a VM, another targets a storage volume, and yet other policies target *groups* of VMs. In a perfect world, you would be able to collapse this policy pyramid into a single management interface that works against VMs alone. After all, those are the workloads you're trying to manage, and having the ability to target policy consistently across all aspects of a VM could be powerful.

## Data Center Costs

With great complexity comes great budget need. The more complex a system is, the more expensive it is to support — and this is often already beyond the amount of some challenging initial data center investments.


## Poor Acquisition Economics

Consider for a moment the manufacturing industry. Almost all manufacturers have adopted a “just in time” manufacturing methodology. Rather than produce all of



the widgets needed to satisfy customer demand for the next three *years*, they produce enough widgets to satisfy customer demand for the next three *days*. This method carries a lot of benefits for the manufacturer:

- **The need to carry inventory is all but eliminated.** The company doesn't tie up capital investment for inventory that just sits on a shelf.
- **Return on investment increases.** Inventory on a shelf carries zero return on investment and, in many cases, actually actively costs the company money in warehousing and other costs.
- **It provides protection against market forces.** For the company that created three years' worth of widgets, what happens when the widget market implodes in the next 18 months? The company is stuck with widgets that it can't sell.



**Figure 1-1:** Actual Data Center Load vs. Purchased Capacity

Now, consider the way that companies buy data center equipment today. In many cases, companies *are* looking out 3 to 5 years for many investments. The capital budgeting process has forced these types of buying cycles, which have themselves spawned replacement cycles. Replacement cycles may vary depending on the equipment being purchased.

### **Significant Zero-Return-on-Investment Zone**

Take a look at **Figure 1-1**, particularly, the gap between the actual data center load and the purchased capacity for Year 1 to Year 3. That gap is, for all intents and purposes, a zero-return-on-investment zone. In the case of the data center infrastructure, it is comprised of resources that the company purchased which are just waiting for the business' needs to catch up to the available capacity.

That may never happen. IT decision makers generally make their best efforts when sizing new resources, but their predictions about needs 3 to 5 years in the future may not hit the mark. In fact, the way that IT departments have traditionally purchased equipment carries a lot of risk, and a lot of it mimics the manufacturing risks discussed earlier, such as:

- **Delayed reaction to change in market forces.** The more “inventory on the shelf” that is in place, the tougher it can be to pivot the business if needs change. If money is tied up in data center equipment that won't even be fully used until 2 years from now, that's money that can't be used to meet emerging needs.
- **Capital constraints.** The more money that is tied up in excess capacity means the less money that is available for other investments. There is a significant

opportunity cost with how organizations fund their data centers today.

- **Technology changes.** The pace of technological change seems to be accelerating with each passing year. Today, IT departments buy data center equipment that needs to last 3 to 5 years for a complete lifecycle. This lengthy cycle forces many companies to wait until a new refresh cycle before adopting the latest technology because they don't have the budget available.

### **Constantly Requires Expansion/Attention**

Many data centers require constant care and feeding. You have to keep adding capacity and performing software updates; you have to keep adding staff that can manage the environment. They are also sometimes unpredictable despite best efforts to size them properly.

### **Requires Deeply Skilled Technical Staff**

Data center equipment is expensive, but it's not nearly as expensive as the highly skilled people necessary to keep it all operational. Data center IT staff tend to be among the more expensive personnel that organizations hire, and they require constant training in order to allow their skills to keep pace with ongoing changes to products.

### **Zero Direct Business Value-Add**

When it comes down to the bottom line, the data center itself has no direct value-add to the business. It's a cost center for most organizations. The real business value comes from the outcomes created by the applications that run in the data center. As such, there needs to be a way to leverage data center assets at the best possible cost while enabling

the most substantial outcomes from the various business applications.

It is this very scenario that has made cloud-based systems so popular today. Organizations leveraging cloud systems are able to focus on application outcomes and not on the complexity of the various inputs, which include networking, servers, storage, and the expensive personnel needed to maintain it all.

With that said, not everyone wants to move to the cloud, but they would like similar outcomes and economics. They would like the ability to adopt a more consumption-based data center model that is far simpler to manage.

### **Data Center Inflexibility**

The typical data center has, in many cases, become the very antithesis of what IT is supposed to be about, which is helping to accelerate the business. Instead, it has often become the chain that binds IT and, by extension, the business. The complexity inherent in many data centers has led directly to situations in which it has become difficult to change direction as business needs change. It's also often expensive, when the need comes, to perform mid-cycle upgrades.

### **Incomplete Integration**

Data centers are often sprawling with different kinds of equipment from many different vendors. This has been the normal scenario for a long time, so IT pros have become used to it. However, it's a part of the reason why organizations see less flexibility from the overall conglomeration. The more equipment there is, the more difficult it is to fully integrate everything.

## **Performance Problem Identification Challenges**

Sufficient application performance is one of the most critical services provided by a data center. Poorly performing applications can directly impact the bottom line by either impacting revenue opportunities or increasing expenses. Neither situation is desirable. Over time, there are all kinds of reasons why applications can become plagued by performance challenges. Tracking down the root cause for these problems can be difficult and often exacerbated by data center complexity. In fact, these challenges have led to an entire infrastructure monitoring market with solutions intended to help people try to figure out the root cause of application performance challenges. In other words, companies have to spend yet more money on the data center because it's so complex.

## **Incomplete Resource Control and Insight**

The traditional data center carries with it a multitude of touch points, which are the points at which an administrator has to interact with the systems. Even with a lot of touch points, though, there are still challenges around holistic control and insight into what is really happening in the environment. Moreover, each resource acts against different constructs that don't always align. For example, storage is managed by created LUNs while networks are provisioned using virtual LANs. On the server side, servers are created as VMs.

These constructs are all measured differently. The LUN is managed based on capacity, the network by throughput, and the VM is, well, a virtual machine. When it comes to policy alignment across these various constructs, these varying measurement methods force administrators to “reinvent” policies at each level of the infrastructure.

# 2

## Solving Complexity With the Hyperconvergence Revolution


---

Hyperconverged infrastructure has emerged as a data center architectural opportunity with the potential to transform the way that IT does business. Hyperconvergence aims to bring simplicity and clarity to the data center in order to enable a more focused effort on delivering business results rather than technical ones.

So, what exactly is hyperconverged infrastructure?

### **Conglomeration of Compute, Storage, and Management**

At its most basic, hyperconverged infrastructure is the result of bringing together the compute and storage layers of the data center and coupling this newly merged construct with massively simplified management.



**Figure 2-2:** Hyperconvergence Illustrated


### Shared Resource Pool

Every resource in the data center currently requires individual management. Each resource is individually sliced up and configured by highly skilled IT personnel and then these resources are individually assigned to support workloads. Nowhere is this paradigm more prevalent than with the storage resource, where storage administrators toil day in and day out creating LUNs, configuring RAID sets, and monitoring available capacity.

With hyperconverged infrastructure, the management process is very different.

### The End of the “Hamster Wheel”

IT goes through many cycles, one of which has become a regular upgrade cycle dictated by capital budgets and depreciation schedules.



**Figure 2-3:** Data Center “Hamster Wheel” Buying Cycle

Shown in **Figure 2-2**, the data center buying cycle generally consists of four distinct phases, beginning with procurement.

## The IT Lifecycle

Each phase of the traditional IT buying cycle has its own set of challenges, each of which can be solved through the implementation of hyperconverged infrastructure solutions.

### Procurement

Buying data center equipment in the traditional way can be risky. First, because of the hamster wheel cycle discussed earlier, IT departments are basically forced to use a crystal ball in order to determine what they'll need for the next 3 to 5 years. Very often, these projections are simply best guesses because business needs are almost certain to change in that kind of time horizon. Because of this, IT leadership often needs to request mid-cycle funding in order to close a capacity or performance gap. This situation is part of the reason that many organizations see IT as very much a cost center that requires constant feeding.

Moreover, the current IT funding model ensures difficulty in realizing relatively short return-on-investment (ROI) cycles. Because IT is buying hardware that may not be leveraged for three to five years, maximum ROI can be quite difficult to achieve.

Hyperconverged infrastructure begins to enable IT to take a more consumption-based approach to the data center. With hyperconvergence, the ability to quickly, easily, and seamlessly scale the environment allows IT to shorten their "resource horizon". That is, they don't need to look out 3 to 5 years; instead, they can plan for the next 6 months. This means that there is potentially less direct initial capital outlay and IT can



pivot much more quickly as business needs change as there is less in the way of sunk costs.

## **Deployment**

The need to integrate equipment from different vendors can be challenging, resulting in the need to perform significant proof of concepts and testing to ensure that everything interoperates the way it's supposed to. However, proof of concepts and integration testing almost certainly extend the amount of time that it takes to deploy new infrastructure and require a significant allocation of staff time. This staff time is laser focused on bits and bytes and not on the company's bottom line. As such, the deployment process can be expensive both in terms of direct costs for staff and opportunity costs when considering what business-facing activities IT staff could have been addressing.

## **Operations**

As new hardware enters the data center, so do new operational procedures. Disparate technology throughout the infrastructure stack creates a continuous challenge for organizations. Every component of hardware, software, networking, and management solution, becomes a gap that requires new operations procedures. These procedures, in turn, have to be designed, documented, tested, and adhered to.

Inter-component communication has also been historically difficult. The lack of abstraction among storage, compute, and networking required tightly coupled systems that are rigid and difficult to manage and maintain. Each component often has its own management tool and operations view. Alerting and automation add an additional layer which is also often disconnected between components.

Hyperconverged infrastructure reduces these deployment challenges by reducing the number of integration points and

vendors that have to be involved with a solution. By combining — at a minimum — storage and compute, there is no SAN to integrate. Plus, with these resources already combined into a single entity, there is no need for integration-focused proof-of-concept testing. By shortening the deployment process, hyperconverged infrastructure allows organizations to achieve a much faster time to value for the data center investment.

Although procurement and deployment improvements are excellent outcomes, they are one-time events that pale in comparison to what can be achieved by improving the operational processes in an IT organization. In fact, this phase is where hyperconverged infrastructure demonstrates its real value to a company.

No more LUNs. No more SAN. No more RAID groups. Easier and more workload-centric policy management. Hyperconverged infrastructure takes a bold step in the direction of data center simplicity. Such simplicity can help companies begin to rethink IT staffing strategies. With the complexity of a SAN, storage engineers instead work to help support common infrastructure elements, including the network and the hyperconverged environment. Simplicity can also help IT staff focus more on outcomes rather than raw inputs all the time. In fact, overworked IT staff might even be able to start going home at a reasonable hour instead of drowning in work.

### **Scale/Disposal**

At some point in a data center's life, there will come a need to either expand what is already in place or to dispose of the equipment as the equipment replacement cycle comes full circle. For some solutions, either of these activities can be momentous events, requiring carefully planned change-management processes to force IT to integrate new equipment

with the old. This is particularly challenging in situations where the legacy equipment has been discontinued and a full-scale migration to a new platform needs to be carried out.

When it's time to grow a hyperconverged infrastructure environment, you just add more. It's that simple. Physically, someone will still have to rack and cable a new appliance, but beyond that, new nodes in a hyperconverged infrastructure are simply assimilated into the already existing resource pool. The net effect is that new resources are made immediately available.

This scalability is really important. In such systems, “scale” is a built-in feature that just works. It's not an event that needs to be carefully staged lest something terrible happen. Hyperconverged infrastructure systems are designed with the intent that new nodes will be added on an ongoing basis.

# 3

## The Missing Pieces in Hyperconverged Infrastructure

---

Even with the outcomes identified in previous chapters, many modern hyperconverged infrastructure systems have gaps in design and packaging. For example, hypervisor choice, lack of network integration, requirement to rip-and-replace existing hardware, limited scalability options, and even inefficient management.

To find out what to look out for when considering your next hyperconvergence solution, read on.

### **The Hypervisor**

The hypervisor itself has become a commodity. The actual real-world feature differences between VMware vSphere, Hyper-V, and open source hypervisors such as KVM is negligible from a practical perspective.

That said, VMware is pretty ubiquitous today. It serves a great purpose and is the current leader in the market for some very good reasons. First, VMware has had a great lead in innovation in the space. Second, the support ecosystem around vSphere is

incredible, both from outside the company and from within. In fact, in recent years, as the hypervisor has become more commoditized, VMware itself has placed more emphasis on their own products that supplement vSphere, even when doing so places them in more direct competition with their partners.

For most mainstream organizations using vSphere, they could get by with Hyper-V or even KVM if they really needed to. Although there are some advanced features in vSphere, feature for feature, Hyper-V or KVM more than hold their own.

Buying into and maintaining the VMware ecosystem can be pretty expensive; it's the price tag that's of concern to many. However, the company's constant product-name changes and licensing changes also confuse end users.

As IT departments are forced to cut costs or even just maintain the same costs year over year, something has to give. Further, many data centers sprawl over the years, resulting in an explosion in complexity and new costs.

VMware and Hyper-V, while excellent products, miss the mark in that they are always "add ons" in the data center. Individual vendors don't get to directly manage that layer. With these systems, no one "owns" the hypervisor layer except VMware and Microsoft. As such, many hyperconverged infrastructure systems based on these products must work around limitations in the product, which can include steep licensing costs. There is no direct control of the hypervisor from the hyperconvergence vendor.

That's where the open source KVM hypervisor comes into play. Although the idea of open source is great, most companies want mission-critical systems running on something fully supported. Therefore, "raw" KVM enterprise deployments

are relatively rare. However, when bundled into a complete solution and when the providing vendor has customized the hypervisor to provide a more complete — and, more importantly, a fully supported experience, it can be a compelling option.

Those that have embraced KVM as a core part of their offering also build the support infrastructure around that tool, effectively eliminating the primary concern that around the product. Further, as vendors embrace and extend KVM, more and more core services make their way into complete solutions that might be able to supplant vSphere and its ecosystem.

## The Network

The network is the tie that binds everything together in the data center. All communication between physical machines, VMs, storage, and other data center equipment relies on some kind of functional networking fabric. Even given the importance of the network, it remains the least integrated function in most hyperconverged infrastructure solutions. There are a couple of reasons for this:

- **The hypervisor.** Most of the networking functionality between VMs, hosts, hyperconverged systems, and client systems is handled by the hypervisor's networking capabilities.
- **Physical hardware.** Although there has been a marked increase in software defined networking (SDN) solutions, the network is the one IT resource that permeates the entire physical infrastructure of the business, not just the data center. With or without SDN, there are cables running through all of the buildings and connected to edge networking devices that sit in their own closets. It's a far more difficult resource to fully and completely assimilate into a combined infrastructure environment.

It's safe to say that in most hyperconverged systems, the hyperconvergence software simply consumes existing networking resources, whether those are provided by the hypervisor for inter-virtual machine communication or by the physical network for other communications needs.

This brings things back to the way that the hypervisor vendor has chosen to support networking in the hypervisor, which may also be dictated by which hypervisor edition is in use in the hyperconverged infrastructure cluster. For example, if you're using VMware vSphere Enterprise Plus, you have access to the vNetwork Distributed Switch, which is not available in other editions of vSphere. When coupled with a hyperconverged infrastructure solution, Enterprise Plus does provide enhanced networking capabilities, but this would also be true in a non-hyperconverged infrastructure scenario. The hyperconvergence hosts simply consume this resource from the hypervisor.

This also means that the hyperconverged solutions that support vSphere remain reliant on VMware for decisions around networking. These decisions can have a negative impact on the overall manageability of the solution and add an additional policy control point that needs to be considered by administrators.

## **Complete Integration**

All too often, companies are pushed into incomplete integration options for data center equipment. The need for what have become termed as “bolted on” components is substantial as companies do everything possible to run an efficient IT operation.

However, systems that sport these kinds of ad hoc integration provide a less than optimal experience, resulting in a lot of

seams between systems. Every seam is an opportunity for error and for inefficiency. Currently, these tightly coupled systems are prone to rigidity with upgrades that may fracture already sensitive systems.

### **Hardware Flexibility**

Very few datacenters today have the advantage of what would be considered a ‘greenfield’ deployment. Enterprise datacenters today are a heterogeneous patchwork of different brands, models, and hardware configurations. Additionally, IT infrastructure teams have invested a great deal of time and energy to gain experience and a degree of comfort with their chosen brand of server infrastructure. Finally, the decision-makers in IT have built established relationships and purchasing power with the reseller, VAR, or vendor from who they have been procuring server hardware with over many years.

Due to their ‘one size fits all’ business model, most hyperconvergence solutions today must ignore the hard-earned trust and relationships that have been built up with existing server hardware vendors and require and instead that enterprises ‘rip and replace’ their existing server infrastructure with entirely new converged server infrastructure. In most environments, not only is this unrealistic but it can devalue the institutional knowledge and expertise developed around existing server hardware standards, not to mention effective and established support and warranty workflows.

What many enterprises need is a solution that gives them ability to realize all the benefits of hyperconvergence while allowing them to deploy hyperconvergence on their *existing* infrastructure. A hyperconvergence solution, architected for hardware flexibility, would enable enterprises with the capability to implement hyperconvergence - and reap its benefits - *on*



*their own terms*, on their proven, preferred servers, procured through their existing vendor relationships.

### **Multi-Directional Scale**

Obviously, different enterprises have different workloads, generated by different applications. Some of those applications work best when they are allowed to scale out (horizontally, by adding more nodes) where other applications work best when they are allowed to scale up (vertically, by adding more resources to existing nodes).

Most hyperconvergence solutions today offer unidirectional scale – scale out. What that means is that when you need more resources of any kind (CPU, memory, storage capacity, storage throughput, or network), your only option is to add additional net-new nodes to the hyperconvergence cluster. While this may not be an issue in some scenarios, scaling out is not always the right solution for every enterprise and every application. There are situations where applications need a large amount of resources on a single node. There are also many scenarios where enterprises would like to save money by simply adding more resources (CPU, memory, or disk) to an existing node or nodes. Additionally, there are often other associated costs when deciding to scale out or scale up which must be taken into account. For example, hypervisor license costs must be considered, in addition to power, cooling, rack space, network connectivity, and more.

By allowing enterprises to use their own hardware, the decision to scale up or scale out can be made by the people who are most effected by the outcomes – infrastructure experts in the datacenter.

## Single Point of Management

Everyone knows that the more complex the management tools are, the more time will have to be spent in management, the greater the difficulty of troubleshooting, and the more likely it will be that an enterprise will suffer downtime simply because monitoring and management was overly complex. One of the stated benefits of hyperconvergence is to provide a single point of management for the server infrastructure, the virtualization layer, and the software-defined storage infrastructure. Thus, instead of having individual management interfaces to manage server hardware, the hypervisor, the SAN/NAS, and the third-party performance/capacity tool, there should ideally be just one, single, interface.

Unfortunately, it's common for hyperconvergence solutions not to provide a single point of management for the requirements listed above. Instead, you end up with, for instance, a new storage management interface that includes a marginally useful virtualization management functionality. For most tasks, administrators still must use their traditional virtualization management tool, server hardware management tool, a third-party performance/capacity tool, and more.

What enterprise administrators need is -truly- a single point of management for the “infrastructure”. That single point of management should, minimally, provide comprehensive virtualization management, storage management, and server hardware management. Ideally, other management tools (network, data protection, performance/capacity, etc) can be rolled into that “single pane of glass” as well.



# 4

## Data Center Desired State

---

Don't forget that there are a lot of different hyperconverged infrastructure options available on the market. With that said, there are some additional desirable outcomes that many organizations will seek as they look to adopt hyperconverged infrastructure.

### **Hardware Heterogeneity**

Appliance-based hyperconverged infrastructure systems have become a very popular option among organizations seeking to simplest way to begin to leverage this burgeoning technology space. These kinds of systems are shipped from the vendor with all of the hyperconverged vendor's software products already loaded and ready to go. As a customer, you receive your new appliance, rack it, cable it, and move on.

But what do you do if you already have existing hardware you want to use? What do you do if you want to or need to use hardware from a specific vendor? For some companies, the ability to deploy truly heterogeneous hardware is an important outcome. In other words, they need a hyperconverged infrastructure solution that is software-only.

Commodity hardware plays an increasingly key role in the modern data center. By leveraging commodity hardware, there is little to no custom engineered hardware in the infrastructure. The end result is that customers can generally have their pick of server vendor, disk (spinning disk and flash disk) vendor, RAID adapters, network adapters, and more.

### **Resource Granularity**

Another challenge is that hardware provided by an appliance-based hyperconverged infrastructure vendor doesn't always ship with the ability for you, as the customer, to fully configure individual resources. As such, you may end up buying more compute than you really need in order to get the amount of storage capacity that you need.

In any environment, hyperconverged or not, it can be useful to have the ability to carefully customize individual hardware resources so that there is minimal resource waste. There needs to be a recognition that different customers have different needs and different tolerance for certain technical activities. While some want the “data center in a box” offered by appliance-based vendors, many want the ability to tailor their needs more granularly while still obtaining the benefits offered by hyperconverged infrastructure.

### **Just-in-Time Data Center**

Earlier, you read about the concept of just-in-time manufacturing and why that approach returns the most dividends for manufacturing companies. With that in mind, how does the hyperconverged infrastructure help move you to this operational paradigm in the data center? Simply put, it's because hyperconverged infrastructure essentially helps you move away from kludging together disparate hardware in favor of a single “unit of infrastructure.” Also, because these units of infrastructure are easy to scale, it's simple to buy what you

need for the short-term and add more units of infrastructure as necessary.

This approach means that you may be able to preserve capital on the solution acquisition front while also having the ability to grow the resource environment in granular terms that make sense for your business. The key term here is your business. What you need in your business is vastly different than what might be needed in other businesses, so why should you be stuck with the same levels of scaling granularity?

Hyperconverged infrastructure allows you to begin adopting a more consumption-focused economic model for the data center. It's not quite like the cloud where you are paying *only* for what you use, but it does get much closer than traditional environments.

### **Complete Workload Support**

As time continues its inexorable march, things change. This fact is doubly true when it comes to the world of technology, which continues to see change at a dizzying rate. It seems like not all that long ago that technology pundits were debating the merits of virtualization, which today, has become a *de facto* data center standard.

Now, we see similar discussions happening around what have become known as containers.

But what exactly is a container?

To answer that, let's talk about VMs for a moment. When you create a VM, you're actually creating an entire server, complete with RAM, CPUs, network adapters, and disks. Of course, it's a software-based instantiation of a physical server. Even though it sits on top of a hypervisor, that doesn't mean that it doesn't consume resources. VMs are comfortable for IT pros

because they look and feel like a real, live server. Applications often don't run any differently inside a software-based construct rather than directly on hardware. The virtualization layer, which consists of the hypervisor, fools the operating system running inside each VM into thinking it has all of its hardware to itself.

There is no debating the fact that virtualization has become the great enabler for all kinds of new ways to manage workloads in the data center, but even this robust technology has some challenges. First, every single virtual server needs its own virtual hardware, which is carved out from physical hardware. So, if you're running 150 virtual machines and they're all running Windows Server 2012 R2, you have hardware overhead for all 150 of those VMs as well as major overhead from 150 running Windows Server 2012 R2 instances. In many cases, IT pros still take the "single application per server" deployment approach to reduce compatibility issues and make troubleshooting easier, so there are a lot of identical operating system instances in play, each of which requires its own resource in order to function.

The reason: virtualization via a hypervisor slices up a physical server and provides hardware-level virtualization in which hardware is abstracted and workloads run on the resulting software layer. Containers take that abstraction to the next level by slicing up the operating system rather than the hardware.

## **Virtual Machines**

Every hyperconverged infrastructure offering supports VMs, although some have different levels of support for different hypervisors. VMs have been absolutely transformational for IT and are the current de facto standard method by which new enterprise applications are deployed.

## **Traditional Workloads**

Traditional data centers carry more traditional workloads in which applications are very much server-based. These are the full virtual servers that are effectively one generation removed from true legacy servers which were once physical.

## **Hardware Emulation**

The driving force to workload portability was the introduction of virtual hardware emulation. This allowed VMs to appear as if they were all running in the same environment, regardless of underlying physical topology. This, in effect, made the physical layers more agnostic to the workload. In essence, the workload just saw the underlying virtual machine as just another server upon which to operate.

## **Containers**

VMs have a lot of overhead associated with them. Each operating system requires resources that are duplicated for each new operating system instance. Recently, containers have begun to appear on the market as an answer to eliminating the duplicative nature of VMs.

## **Shared Binaries**

One layer further up the virtualization stack, container infrastructure allows the further abstraction to include pre-packaged binaries and supporting libraries to launch and run applications more rapidly. In plain English, this means that containers, rather than relying on hardware abstraction, basically slice up a single operating system into multiple chunks and run applications inside each slice.

## **Less Overhead = More Density**

Because the containers require less overhead to deal with the virtualization stack and virtual hardware, they can be



much lighter in nature and allow for much heavier workload density. With containers, you can run more workloads on the same amount of hardware since there is far less workload duplication.

## **OpenStack**

OpenStack was created as an open private cloud platform geared toward developers and with hopes of achieving near Amazon-like simplicity to consume resources in the enterprise.

### **Eliminate Proprietary Components**

Touted as the path to escape vendor lock-in through open source innovation, OpenStack has emerged as a force in building public and private cloud infrastructure. This is achieved while running on nearly any hardware platform.

### **Improve Data Center Operations**

OpenStack is moving toward truly commoditizing the infrastructure and providing a single, widely supported set of APIs to consume and manage the components. The more that the infrastructure can be commoditized and consumed via APIs, the easier it becomes to automate IT operations.

Not only built as an IaaS (Infrastructure-as-a-Service) platform, OpenStack is also an enabling platform to house orchestration, containers, database (DBaaS), Big Data, and more. When leverage correctly, OpenStack can drive inefficiency out of the data center management equation.

# 5

## Simplifying and Securing the Heterogeneous Data Center

---

Today's data center has significantly different needs than the data centers of yesteryear. While struggling to maintain datacenter availability and performance, IT professionals must continue to adapt to the technology needs and ever-increasing expectations of the business. For IT Pros to have a chance at success, they require a data center infrastructure that can keep pace and that can quickly adapt to meet the ever-emerging needs of the business. These requirements are described in the following sections.

### **Hardware Heterogeneity and Independent Resource Scaling**

Scaling your data center systems as you grow organically as a business is the goal of web-scale architectures. Hyperconvergence and the heterogeneous data center means that forecasting is more aligned to the just in time models that the business is reaching toward.

A true modern data center solution also means being able to add storage independently of your compute, increase compute-to-target workloads that consume CPU, or create more resilient internal networks for application platforms.

Loosely coupled systems use software abstractions to ensure better manageability and upgrade protection. The decoupling of storage, compute, and network creates flexibility to deploy, grow, and manage your data center. Enabling RESTful API access into each component extends that manageability to nearly any orchestration or CMP (Cloud Management Platform).

Abstractions are created in order to bring the flexibility for the virtual application workload running inside the data center environment. The real goal of software-defined solutions is to entirely abstract the underlying infrastructure, regardless of brand or resource, and to present a single virtual hardware implementation to the virtual instances.

### **Any Vendor, Any Time**

Virtualization vendors leveraged the idea of creating independence from hardware. However, by virtualizing physical servers, the dependency just moved from the hardware up to the hypervisor, still leaving lock-in as a challenge. Virtualization was supposed to “cut the hardware cord,” but it has yet to completely do so.

Hardware that is purpose-built for specific workloads has become the norm. Using the “best of breed” hardware in order to meet application demand means that multi-vendor data centers are the new reality. You can no longer entirely replace one vendor with another very easily. The disruption of such moves is not necessary thanks to what hyperconvergence has brought to the ecosystem.

Long gone are the days where racks of hardware will need to be brought in for a total migration from one platform to another. It is just not a feasible to expect organizations to entirely buy-in to a single vendor and to continue to forklift the entire environment at every hardware refresh.

## **Hyperconverged Infrastructure Can Exacerbate the Problem**

As you look at appliance-based hyperconverged infrastructure solutions, bear in mind that you may adopt a solution that might not scale as granularly as you need. This may be an acceptable trade-off for the benefits provided by hyperconverged solution, but make sure you fully understand the options at your disposal.

Many vendors require specific hardware configurations. Different vendors and hardware offerings can have very distinct and individual requirements. This is highlighted by the introduction of hyperconverged infrastructure. Configuring and tuning hardware and the various hypervisors has led to frustration in trying to adhere to vendor-specific best practices across platforms.

The introduction of hyperconvergence brought the idea of true software-defined operating environments that could be run using alternate hardware platforms. Storage in particular is known to be challenging to tune and manage across platforms.

Despite the issues highlighted by disparate and heterogeneous environments, perhaps it has been viewed in the wrong light. The shift has begun away from providing a continuous vertical growth up to scale limits followed by building a disparate hardware pool. As physical hyperconvergence entered the market, a more scalable architecture was introduced.

The question is, why does this have to be done on one hardware platform only? Or doesn't it? This raises the importance of hyperconvergence and the capability to span disparate hardware solutions using a common abstraction. Hardware hyperconvergence did a lot to break down the traditional silos of separate compute and storage, tied together with physical and virtual networking. Another layer of abstraction is needed to truly create the pooled resources as network is integrated into the hyperconverged architecture.

### **Avoid the Forklift**

Despite some solutions providing more flexibility and attending to some of the common legacy challenges, there still exists a need to forklift into new hardware. New systems often promise the need to only have to do this one more time, but the reality is that many of these solutions present scale limitations requiring upgrades and replacements to grow along with the business needs.

### **"Set It and Forget It"**

Plug-and-play was among the features that made the Microsoft Windows desktop experience dominate the desktop market. Why can't our hardware be as simple as plugging in and having a zero-touch installation? This concept is what attracts many to hyperconverged solutions, which offer the low-touch experience that makes the adoption smoother and less invasive.

### **Granular Resource Control**

In the same way that an audiophile may opt for some more high-touch hardware to achieve the ultimate sound experience, it is often a similar case inside the data center. Some engineers enjoy being more involved in the control and tuning of resources, and there are genuine cases where

this is necessary to provide the granularity of control and performance for certain environments.

Just as there are hybrid clouds to deliver the right solutions for either side of the business case, the same holds true for the hyperconverged platforms. This can include granularity where needed, while simultaneously offering more wizard-driven approaches to deployment.

### **Heterogeneous Hardware as a Key Enabler**

When looking a hyperconverged infrastructure solution for your data center, consider the potential benefits that can come from a solution that can run on any hardware you choose. This heterogeneous-based approach carries with it a number of possible benefits. Hardware and the ever-growing requirements of vendor HCLs (Hardware Compatibility List) can become less of a limitation. Moving the abstraction to the right level above the hardware allows organizations to embrace heterogeneity in the data center by buying what they need when they need it.

### **Leverage Existing Hardware Resources**

Today's hardware solutions were architected for much longer physical longevity than leasing agreements might suggest. Storage hardware is particularly prone to this situation, with extremely long lifecycles due to the challenge of migrating workloads in and out. It's the hardest resource to migrate. As a result, it tends to stay in production for a long time.

Why would you continue to move from one forklift migration plan to another, when you can keep your existing hardware infrastructure and abstract it to preserve and extend its life? This approach can save you a lot of money over time.

## **Reduce Sunk Costs**

Imagine being able to extend the life of existing hardware, thus stretching the value of those sunk cost, physical resources. Removing the tight coupling to hardware provides the next level of value as hyperconvergence presents these resources as liquid pools of network, storage, and compute. Growth becomes less attached to the vertical scale limitations and challenges on existing platforms. Extending the resource pools with hyperconvergence can be done by adding any of a variety of vendor solutions, bolstering the elastic resource pools, and maximizing the value for existing hardware spending.

## **Adopt the “Just in time” Data Center**

Using hyperconvergence, you can now get the best of both worlds by fully leveraging existing infrastructure, while becoming agile with a “grow as you go” growth strategy.

These are the tenets of the public cloud that have drawn the eye of CIOs and CFOs alike. Being able to be more elastic with data center resources, as well as keeping the benefit of on-premises protection and controls is what draws organizations to private cloud deployments.

Hyperconvergence lets companies fully realize these capabilities by providing operational and growth flexibility. This is the next step in the evolution as budgets and hardware lifecycle become shorter and more adaptive to the business needs as they increase.

## **Extreme Abstraction**

Hardware abstraction has become something of a trend in its own right. VMware and others ushered in the modern abstraction trend by introducing to the world x86 virtualization technologies that completely commoditized the server. This

was accomplished by abstracting the hardware and allowing the hypervisor to become the application platform.

This trend is continuing as the market introduces abstraction options for storage and networking. It's all part of a larger need to simplify the data center environment and make it as invisible as possible to the business.

### **Data Center Hardware and Plumbing**

No one cares about the pipes behind the walls. The same thing applies when you walk into your data center. For the business, the data center holds the pipes that should operate behind the wall. Abstracting the underlying hardware and simplifying the delivery of services is the true path to bringing IT to the next generation. Attention on the plumbing that drives the data center is not where today's CIOs and operations teams want the focus to be.

In the same way that water comes out of the taps, there should be a way to simply have hyperconverged resources that make the plumbing invisible to the consumer. Many businesses have an office, but they don't employ plumbers, electricians, and full time building architects. They have those capabilities covered by building management teams that they engage as needed.

While it was necessary in the past to tune and customize every stage of configuration and operations, this is no longer the case. Modern systems administrators should be spending their engineering time on creating business value rather than architecting every cable and connection for resiliency and performance.

In the same way that a business doesn't spend time and effort managing the physical plumbing in their office, it is becoming obvious that the data center should follow suit.



## **A Focus on Outcomes**

Businesses want to create value for their customers, and spending unnecessary person hours building and managing technology solutions creates a cost and a challenge that should not be part of the business plan.

Infrastructure needs to be able to be available with immediacy, and operate with consistency, ultimately letting the business get back to what it is designed to do, which is to grow the business and satisfy customer demand.

## **Abstract It All**

Compute and storage were a good start. The advancements that came with virtualization laid the groundwork for the next generation of IT. This pathway was paved by abstracting compute and storage, with networking brought in next to create a more cohesive and consistent platform. Creating the next generation of data center platform with hyperconvergence, on industry tested products, is now the next step to true enablement for IT.

Where the public cloud had all the advantages in the beginning, private cloud capabilities are now entirely within reach thanks to hyperconvergence. This next leap in features with hyperconvergence is simplifying the consumption of IT resources, and enabling both development and operations teams to use a singular view and singular processes.

# 6

## Understanding the Network in the Hyperconverged World

---

Physical networking grew to find its limitations, followed by VLAN and logical networking reaching limitations for multi-tenant and overlapping network ranges. Now the hyperconverged platform needs to be the culmination of those advancements while also providing the ease-of-use for deployment and management.

As network and application complexity increases, it is apparent that many hyperconverged solutions have not architected with this in mind. Networking continues to be the third part of the equation often left behind by engineering within storage and compute.

While there are limitations and challenges in virtual networking deployments on current hypervisors, it is also a foundation that has driven thousands upon thousands of deployments in the industry today. In its current state, hypervisor networking has

brought a reasonable level of features and innovation in virtual deployments up to now which can be brought to hyperconverged solutions.

As more workloads are operated on hyperconverged solutions, more about the way to operate those workloads is being discovered. Boundaries will be pushed to create the true list of requirements that will drive the next stages of innovation in networking and hyperconvergence.

Compute and storage virtualization was the leading innovation that drove the need for hyperconvergence. Networking is now the final layer that will be driven to the next level, and with the advantages in a hyperconverged environment, a single toolset and unified management will be where it takes place.

### **Missed Networking Needs**

Embracing DevOps methodologies and more agile IT infrastructure management means that multi-tenant environments have become much more important and necessary. Private cloud platforms, often less popular due to complexity, have shown that there is an appetite to make multi-tenancy a common deployment technique in the data center. Many solutions are lacking in the ability to deliver true multi-tenancy to meet the needs of today's operations and development teams.

### **Security**

Beyond multi-tenancy as a complexity challenge, it is also a security challenge. Networking in hyperconverged platforms introduces potential for vulnerability in network security.

It is said that a majority of networking traffic in the coming years will be East-West traffic between virtual workloads.

As the shift toward multi-tenant, complex networks arrives, so does the importance of attending to security within these platforms.

### **Workload Isolation**

A key feature of modern platforms is the ability to enjoy lab deployments without negatively affecting the surrounding virtual resources. With multi-tenant platforms comes the ability to create workload isolation and maximize the safety within the platform.

This is quickly becoming a must-have item as modern application development practices embrace multiple environments, Infrastructure-as-Code, and DevOps methodologies. By providing a platform that lets teams work individually but use the same practices and tools on the same platform, means that organizations can get closer to the desire of rapid application deployment without compromising safety and security.

More opportunities are arising in network security and they are well up the stack from the typical Layer 2 and Layer 3 security (that has been where most attention has been put up to now in compute infrastructure). As the complexity within applications rises, so does the need to create true workload isolation for testing, security, and more.

### **Service Provider Needs**

Single organizations introduce the potential for challenges in multi-tenancy. Service providers and shared customer environments increase the criticality of robust, secure networking. It is not even something that could be entertained at the service provider if robust and flexible multi-tenancy was not available. This is core to the business of service

providers, and requiring custom built operations tools and staffing to support these environments, it would have a direct effect on the operational cost.

Imagine a telecom company that finds it has multiple customers on the same links without isolation, or a cloud provider who operates customers in comingled virtual data centers? Clearly this is not an option.

Security and separation of logical environments on hyperconvergence is done by design with these situations in mind. It is simply not feasible to have security and workload isolation as an afterthought.

Up to this point, service providers have typically relied on separate physical compute and storage to provide customer isolation. Multi-tenancy was mostly provided at the ingress and at the network layer, although much of that was also physical as well.

Hyperconvergence and managed private clouds have changed the needs for service providers and the requirements of customers have aligned with the capabilities that multi-tenant data center environments can provide.

## **Ensuring Security and Performance**

In order to absolutely ensure sufficient performance while maintaining appropriate levels of security, an infrastructure requires deep level of understanding of I/O patterns and pathways to avoid interference. This has not truly been possible when the infrastructure vendor doesn't own the hypervisor.

The primary goal of virtualization was to simply present virtual space in which to launch virtual machines. This is a far cry from a true intelligent, workload-aware platform which is needed for today's application environments.

By having hardware and software owned and operated separately in the past, this gap was always present and created limitations requiring either custom development of tools or purchase of additional software packages.

### **QoS and Dynamic Workload Rebalancing**

Where cloud-like deployment provided the flexibility of launching applications, it fell short on operational efficiencies. The simple round-robin schedulers and limited decision algorithms only operated at start-up.

Having a more enhanced view of workload operation and the ability to manage efficiency as usage patterns fluctuate is what is needed. The “set it and forget it” model does not satisfy enterprise needs and is just not appropriate on private cloud platforms.

### **Hyperconvergence and Software Defined Networking**

SDN has a distinct set of use-cases that will come to be important as hyperconvergence adoption increases. Cross-cluster and stretched cluster implementations are on more often on the wish list than the requirements list for today’s enterprise data center owners.

Within the cluster of course, there is a critical need to leverage abstractions at the network layer. In order to provide the flexibility and resiliency for the virtual workloads running inside the hyperconverged infrastructure, software-defined networking is key.

As the use-cases arise, so will the solutions for hyperconverged platform. When the adoption of SDN increases in the industry, the need to be more tightly tied to those platforms may arise.

At this point, networking in hyperconvergence is focused on providing the workload and the management components what they need to operate seamlessly across the environment. Despite the current hype of the industry flocking toward all-SDN implementation, the reality is that organizations have not yet found the use-cases to become heavy adopters of these products.

# 7

## Improving Operations: DevOps in the IT World

---

No longer the stuff of legend, DevOps methodologies are fast becoming more popular in adoption. Rapid time to market for applications is driving a new way of doing business in IT which opens the door for DevOps enablement.

This is one of the core reasons for the increase in cloud adoption in both the public and private cloud markets, because the API-driven cloud enables developers to get products and infrastructure on-demand and with less friction.

The primary challenge for DevOps adoption is that many data center environments are not designed for this paradigm. Doing DevOps right requires very fast auto-provisioning of resources to meet needs of rapid, iterative development cycles. Traditional virtualization solutions have not been able to meet the needs of



the rapid delivery of applications. Retrofitting existing virtualization architectures can provide some additional flexibility, but this is often met with resistance and challenge in enabling true cloud-like capabilities on legacy systems.

Today's development teams need rapid environment provisioning to do true test-driven development, and being able to meet the needs of multiple developers, and multiple development teams. Agility in operations needs to catch up to the agility that has been achieved in development practices.

In the same way that virtualization introduced the potential for virtual machine sprawl and the lack of control in monitoring and managing growth, hyperconvergence introduces simplicity in deployment that reduces the need for enhanced reactive monitoring.

Understanding what is happening within the environment needs to dive further into resource monitoring and give a much more granular view of application and virtualization layers. These tools help to provide the detailed information that operations teams need to manage their hyperconverged platform to the best possible efficiency.

# Building The Perfect World

---

There is a solution that can help you meet all of the challenges that have been described in this book. From ensuring that there is a consistent platform on which to run applications to enabling extreme resource expansion granularity to embracing the networking component, when done right, this solution can transform the data center and the IT department.

## **The Hyperconvergence Vendors Own It All**

Purpose-built to ensure interoperability, hyperconvergence has brought the heterogeneous data center concept to reality. Each layer in the data center stack is a part of the integrated solution, while maintaining the loosely coupled systems internally; truly the best of both worlds.

The attention to bringing all of the operational components together seamlessly was top of mind in designing a hyperconverged solution. It should not be necessary to have to go further than one interface and one platform in order to provide a versatile data center platform.

Every component within a hyperconverged solution is designed to work in conjunction with every other component. No innovation is happening in isolation because of the power of having every aspect of the solution handled in one place, by one vendor. Upgrades and updates of the operating environment are fully tested long before they make it to customer implementation. Also, by having control over compute, storage, and networking, the hyperconverged vendor is acutely aware of every moving part in the environment, and removes the reliance on other vendors outside the platform to further reduce risk.

### **Deep Integration**

The goal of providing a true OOB (Out of the Box) infrastructure means eliminating the poor integration of external services. With hyperconvergence comes the more seamless integration of the components to install, deploy, and manage the infrastructure.

It has been noted that difficult integration and finicky shims between systems creates tighter coupling, which will create rigidity, and more importantly, greater risk.

Modern solutions require fully and completely centralized management for all cluster resources, including internal networking. One management interface to see and manage the entire infrastructure stack has been the long sought after panacea of data center operations teams. While the term single pane of glass is often overused, the hyperconvergence solution delivers the closest to a true single interface for meeting all of the needs of systems administrators.

Managing storage, compute, networking, and day-to-day administration needs to be simple. With your operations teams

able to simplify the administration processes, they can get back to adding value in other areas of innovation within the data center.

Developers and operations teams alike will reap the benefits as hyperconvergence brings operational simplicity, clean and simple API interaction, monitoring and management, and most of all, consistency and performance.

# Discover how Stratoscale's complete hyperconverged infrastructure solution is engineered to simplify and transform IT

There are a number of worthy hyperconverged infrastructure solutions on the market, with just about all of them ably converging compute and storage resources to simplify management and introduce efficiency. Lost in the mix, however, is attention to the network and full ownership of the hypervisor layer, forcing manufacturers to bend to the hypervisor's will rather than embrace and extend it. Stratoscale brings to the market a highly efficiency hyperconverged infrastructure solution that, in addition to aggregating compute and storage resources, also harnesses the network and hypervisor. By doing so, Stratoscale provides a resource-complete data center stack and, via hypervisor ownership, far deeper insight into and management of virtual machines that run inside the hypervisor.

## Learn:

- Why it's important for the network to be considered a first-class resource in the world of hyperconverged infrastructure
- The critical reasons that full control of the hypervisor can revolutionize the nascent hyperconverged infrastructure market
- How to achieve your data center's "desired state"
- How OpenStack and containers fit into the hyperconvergence story
- Achieving the perfect world in your data center



ActualTech Media

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