

HPE PRESENTS

# THE GORILLA GUIDE TO...<sup>®</sup>



# The New Role of the IT Pro in a Hyperconverged World

Scott D. Lowe

---

## INSIDE THE GUIDE:

- Discover the basics of hyperconverged infrastructure as a technology
- Learn which resources get consolidated in a hyperconverged environment
- Understand the new ways that the IT Pro can thrive in a hyperconverged world

**HELPING YOU NAVIGATE  
THE TECHNOLOGY JUNGLE!**



ActualTech Media  
[www.actualtechmedia.com](http://www.actualtechmedia.com)

In Partnership With



**Hewlett Packard  
Enterprise**

# THE GORILLA GUIDE TO...

## The New Role of the IT Pro in a Hyperconverged World

### **AUTHOR**

Scott D. Lowe

### **EDITOR**

Keith Ward, ActualTech Media

### **LAYOUT AND DESIGN**

Olivia Thomson, ActualTech Media

Copyright © 2019 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.

### **ACTUALTECH MEDIA**

Okatie Village Ste 103-157

Bluffton, SC 29909

[www.actualtechmedia.com](http://www.actualtechmedia.com)

# ENTERING THE JUNGLE

<b>Chapter 1: Introduction to Hyperconverged Infrastructure</b> .....	6
Hyperconverged Infrastructure from 30,000 Feet.....	8
Resources to Consolidate.....	10
<b>Chapter 2: The New Role of the IT Pro in a Simplified Data Center</b> .....	14
The Rise of the Versatilist.....	20

# CALLOUTS USED IN THIS BOOK



The Gorilla is the professorial sort that enjoys helping people learn. In the School House callout, you'll gain insight into topics that may be outside the main subject but are still important.



This is a special place where you can learn a bit more about ancillary topics presented in the book.



When we have a great thought, we express them through a series of grunts in the Bright Idea section.



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



Discusses items of strategic interest to business leaders.

# ICONS USED IN THIS BOOK



## **DEFINITION**

Defines a word, phrase, or concept.



## **KNOWLEDGE CHECK**

Tests your knowledge of what you've read.



## **PAY ATTENTION**

We want to make sure you see this!



## **GPS**

We'll help you navigate your knowledge to the right place.



## **WATCH OUT!**

Make sure you read this so you don't make a critical error!

# CHAPTER 1

## Introduction to Hyperconverged Infrastructure

The world of technology is changing at both a broader and faster pace than ever before. In years past, IT practitioners and decision makers might have had to deal with a few new hardware options and some software upgrades. Today, changes come every day as companies break the data center walls in favor of hybrid cloud, as cloud providers continue to grow, and as software updates deploy silently behind the scenes.

As employees devour technologies such as smartphones, tablets, wearables, and other devices, and as they become more comfortable with solutions such as Slack, Dropbox, and Zoom, their demands on enterprise IT intensify. On top of this, management and other decision makers are also increasing their demands on enterprise IT to provide application environments that have higher levels of availability and easier growth capability, but with the kinds of pay-as-you-grow economics that they see from the cloud. Unfortunately, enterprise IT organizations often don't see much, if any, associated increases in funding to accommodate these demands, particularly as spending increases on other business priorities, such as improved security, analytics, and digital transformation efforts.

These demands have resulted in the need for IT organizations to attempt to mimic NASA's much-heralded "Faster, Better, Cheaper" operational campaign. As the name suggests, NASA made great attempts to build new missions far more quickly than was possible in the past, with greater levels of success, and with costs that were dramatically

lower than previous missions. NASA was largely successful in its efforts, but the new missions tended to look very different from the ones in the past. For example, the early missions were big and complicated with a ton of moving parts, while modern missions have been much smaller in scale with far more focused mission deliverables.

The same “Faster, Better, Cheaper” challenge is hitting enterprise IT, although even the hardest-working IT pros don’t usually have to make robots rove the surface of an inhospitable planet! Today’s IT departments must meet a quickly growing list of business needs while, at the same time, appeasing the decision makers who demand far more positive economic outcomes, either by cutting costs overall or doing more work within the existing budget.

Even as the public cloud continues to gain favor for certain workloads, the real center of workload action remains the on-premises data center. Unfortunately, traditional data center architectures actively work against modern goals, because with increasing complexity comes increased costs—and things have definitely become more complex. Virtualization was a fantastic opportunity for companies, but with virtualization came new challenges, including major issues with storage. With virtualization, enterprise IT moved from physical servers, where storage services could be configured on a per-server basis, to shared storage systems. These shared storage systems, while offering plenty of capacity, have often not been able to keep up in terms of performance, forcing IT departments to take corrective actions that don’t always align with good economic practices.

These are just some of the challenges that administrators of legacy data centers need to consider as well:

- **Hardware sprawl.** Data centers are littered with separate infrastructure silos that are all painstakingly cobbled together to form a complete solution. This hardware sprawl results in a data center that’s increasingly complex, decreasingly flexible, and expensive to maintain.

- **Policy sprawl.** The more variety of solutions in the data center, the more touch points that exist when it comes to applying consistent policies across all workloads.
- **Scaling Challenges.** Predictability is becoming extremely important. That is, being able to predict ongoing budgetary costs and how well a solution will perform after purchase are important. Legacy infrastructure and its lack of inherent feature-like scaling capability make both predictability metrics very difficult to achieve.
- **Desire for less technical overhead.** Businesses want analysts and employees that can help drive top-line revenue growth. Purely technical staff are often considered expenses that must be minimized. Businesses today are looking for ways to make the IT function easier to manage overall so that they can redeploy technical personnel to more business-facing needs. Legacy data centers are a major hurdle in this transition.
- **A focus on security.** Security has always been important, but never more than it is today. Increasingly a boardroom issue, security spending is accelerating, sometimes at the expense of other areas of IT, a situation that data center architects need to accommodate.

So, with all of this in mind, what are you to do?

## Hyperconverged Infrastructure from 30,000 Feet

A number of years ago, a new data center architectural option, dubbed hyperconverged infrastructure, came on the scene and exploded faster than anyone could have imagined. Hyperconvergence is a way to reduce your costs and better align enterprise IT with business needs. At its most basic, hyperconverged infrastructure is the conglomeration of the servers and storage devices that comprise the data center with modern iterations of the technology also swallowing up complex networking components. These systems are wrapped in comprehensive

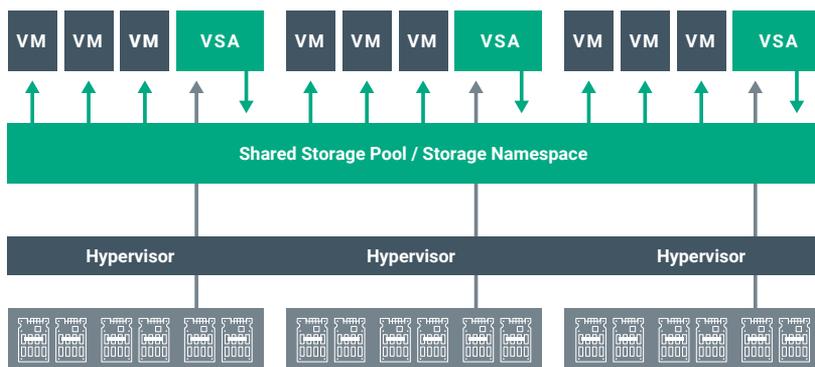
and easy-to-use management tools designed to help shield the administrator from much of the underlying architectural complexity while providing an infrastructure that's self-managing.

Why are storage and compute at the core of hyperconverged infrastructure? Simply put, storage has become an incredible challenge for many companies. It's one of—if not the—most expensive resources in the data center and often requires a highly skilled person or team to keep it running. Moreover, for many companies, it's a single point of failure. When storage fails, swaths of services are negatively impacted. Finally, with the increase in the growth of data volume, legacy storage architectures were beginning to crumble.

Combining storage with compute is in many ways a return to the past, but this time serious brains have been wrapped around it. Before virtualization and before SANs, many companies ran physical servers with direct-attached storage systems, and they tailored these storage systems to meet the unique needs for whatever applications might have been running on the physical servers. The problem with this approach was it created numerous “islands” of storage and compute resources without a centralized coordinating mechanism. Virtualization solved this resource-sharing problem, but introduced its own problems previously described.

Hyperconverged infrastructure distributes the storage resource among the various nodes that comprise a cluster. Built using standard server chassis and hardware, hyperconverged infrastructure nodes and appliances are bound together via Ethernet and a powerful software layer. The software layer often includes what is termed as a virtual storage appliance (VSA) that runs on each cluster node. Each VSA communicates with all of the other VSAs in the cluster over an Ethernet link, thus forming a distributed file system across which VMs (VMs) are run.

If you reread that last paragraph, you'll note that the network plays a critical supporting role in every hyperconverged solution on the market. However, few hyperconverged vendors treat the network as more



**Figure 1-1:** An overview of a Virtual Storage Appliance

than a support infrastructure. Today, as organizations scale beyond single racks, overlooking the criticality of the network introduces serious scaling limitations. Later in this book, you’ll discover how HPE has solved this challenge through the introduction of the HPE Composable Fabric, a software-defined networking solution that eliminates scaling challenges in hyperconverged infrastructure (and other) scenarios.

The fact that hyperconverged systems leverage standard off-the-shelf hardware is critical. The power behind hyperconverged infrastructure lies in its ability to corral resources—RAM, compute, data storage, and networking—from hardware that doesn’t all have to be custom-engineered. This is the basis for hyperconverged infrastructure’s ability to scale granularly and the beginnings of cost reduction processes.

## Resources to Consolidate

The basic combination of storage, servers, and networking is a good start, but once you look beyond the confines of this baseline definition, hyperconverged infrastructure begins to reveal its true power. The more hardware devices and software systems that can be collapsed into a hyperconverged solution, the easier it becomes to manage the solution and the less expensive it becomes to operate.

Here are some data center elements that can be integrated in a hyper-converged infrastructure.

## **Deduplication Appliances**

In order to achieve the most storage capacity, deduplication technologies are common in today's data center. Dedicated appliances are now available that handle complex and CPU-intensive deduplication tasks, ultimately reducing the amount of data that has to be housed on primary storage.

Deduplication services are also included with storage arrays in many cases. However, deduplication in both cases is not as comprehensive as it could be. As data moves around the organization, data is rehydrated into its original form and may or may not be reduced via deduplication as it moves between services.

## **SSD Caches/All-Flash Arrays**

To address storage performance issues, companies increasingly deploy either solid-state disk (SSD)-based caching systems or full SSD/flash-based storage arrays. However, both solutions have the potential to increase complexity as well as cost. When server-side PCI-e SSD cards are deployed, there also has to be a third-party software layer that allows them to act as a cache, if that's the desire. With all-flash arrays or flash-based stand-alone caching systems, administrators are asked to support new hardware in addition to everything else in the data center.

## **Backup Software**

Data protection in the form of backup and recovery remains a critical service provided by IT and is one that's often not meeting organizational needs. Recovery time objectives (RTO) and recovery point objectives (RPO)—both described in the deep-dive section, “The Ins and Outs of Backup and Recovery”—are both shrinking metrics upon which IT needs to improve.

The basics behind hyperconverged infrastructure should be well understood before proceeding with the remainder of this book. If you're new to hyperconverged infrastructure or are unfamiliar with the basics, please read "Hyperconverged Infrastructure for Dummies," available now for free from [www.hpe.com/HCI/fordummies](http://www.hpe.com/HCI/fordummies)



Using traditional hardware and software solutions to meet this need has been increasingly challenging. As RPO and RTO needs get shorter, costs get higher with traditional solutions.

With the right hyperconverged infrastructure solution, the picture changes a bit. In fact, some baseline solutions include comprehensive backup and recovery capability that can enable extremely short RTO windows, enabling very small RPO metrics—both *very* good characteristics!

## Data Replication

Data protection is about far more than just backup and recovery. What happens if the primary data center is lost? This is where replication comes into play. By making copies of data and replicating that data to remote sites, companies can rest assured that critical data won't be lost.

To enable these data replication services in traditional scenarios, companies implement a variety of other data center services. For example, to minimize replication impact on bandwidth, companies deploy WAN-acceleration devices intended to reduce the volume of data traversing the Internet to a secondary site. WAN accelerators are yet another device that needs to be managed, monitored, and maintained.

## The Ins & Outs of Backup & Recovery

There are critical recovery metrics—known as recovery time objective (RTO) and recovery point objective (RPO) that must be considered in your data protection plans. You can learn a lot more about these two metrics in *The Gorilla Guide to Hyperconverged Infrastructure Implementation Strategies*.



There are acquisition costs to procure these devices; there are operation costs in the form of staff time and training; and there are annual maintenance costs to make sure that these devices remain supported by the vendor.

### Up Next

With an understanding of hyperconverged infrastructure and knowledge about many of the resources that can be consolidated into such solutions, let's move on to discuss the ramifications that the technology may have on the career path of the IT Pro.

## CHAPTER 2

# The New Role of the IT Pro in a Simplified Data Center

Every technology paradigm shift is fraught with anxiety as each one introduces uncertainty into what people want to be a stable equation. Hyperconverged infrastructure brings to the market a solution that is truly revolutionary. With such change comes concern around what this means for careers.

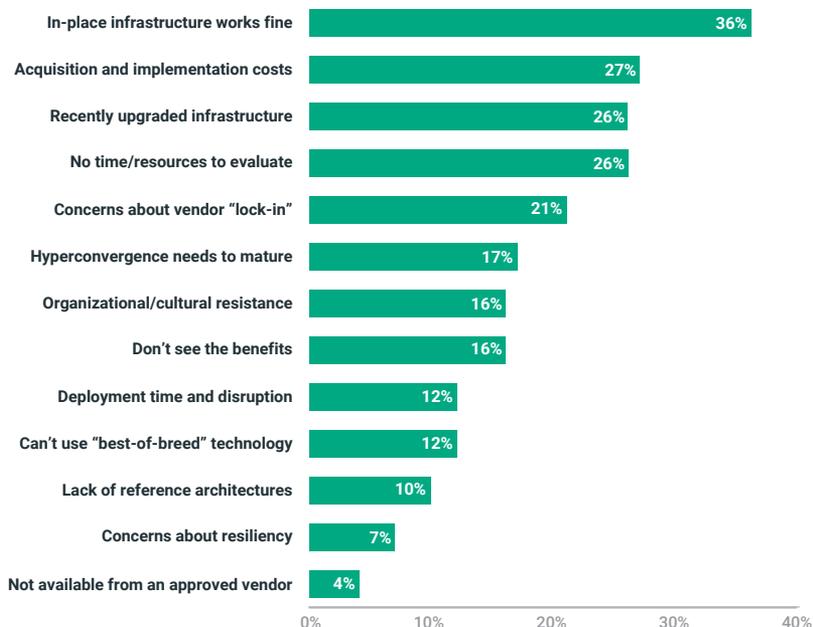
Simply put, hyperconvergence is a career opportunity for IT pros, even for those with specialized storage knowledge. While the change may give some the jitters, you should look at these feelings as growing pains. For admins looking to broaden their skills, there's big opportunity for job growth when you embrace new technologies.

But that doesn't mean modernization is *always* desirable. There can be perfectly valid reasons to delay adoption of new technologies.

In research conducted by ActualTech Media, IT pros identified a number of reasons why they're reluctant to move forward with hyperconverged infrastructure initiatives (**Figure 2-1**). In spite of the benefits it can bring, not everyone is enamored with the potential for hyperconverged infrastructure, for reasons that range from cost to resiliency concerns about this "new" architecture.

## Which are the primary reasons you have no interest in deploying hyperconverged infrastructure in the near term?

(N=290, Up to Three Responses Allowed)



**Figure 2-1: Primary reasons for not deploying hyperconverged infrastructure**

However, the top reasons why people aren't yet looking at the technology have nothing to do with the technology itself, but rather have to do with the business cycle:

- **Current solution works just fine.** The adage, "If it's not broke, why fix it?" holds true for many. However, that won't always be the case. Business priorities change on a dime, and understanding the significant benefits that may come from modernizing the data center with hyperconverged infrastructure solutions will be useful for the future.
- **Recently upgraded infrastructure.** This is a valid reason for businesses that have upgraded the entire infrastructure in the past 6 months. For everyone else, it should be a minor consideration.

Infrastructure components are constantly being refreshed – that’s a fact of business. Most hyperconverged infrastructure solutions do not require forklift upgrades. Different solutions offer varying degrees of integration opportunities but most can integrate into the existing environment at some level. In fact, the technology is designed for modular growth.

- **No current IT or business need.** Some people truly have no present infrastructure needs and are focusing their efforts in other areas. These folks can be considered as part of the “current solution works just fine” category. If your IT organization falls into this category, you might want to start looking for opportunities to introduce hyperconvergence at some point, to ease your infrastructure into the future. If new applications are being deployed or there’s a specific use case (such as VDI or edge), there may be an opportunity to introduce hyperconverged infrastructure into the environment. Many businesses start with a pilot program and expand it from there.

In this chapter, we’re going to tackle some of the more serious issues that people have brought up with regard to hyperconvergence.

## The IT Staffing Challenge

It should not come as a surprise that most companies want to limit the number of IT staff that they hire, at least in certain areas. Infrastructure is almost always one of those areas, because it’s directly associated with the expense side of the ledger.

As a result, many companies aren’t willing to scale their infrastructure staff at the same rate that they scale the infrastructure itself.

At that point, the IT staff is forced to take on more and more responsibility while working with dwindling resources. This is the very definition of that loathed phrase, “Doing more with less.”

This situation is one of the very reasons why people consider hyperconverged infrastructure. By massively reducing the variety of

hardware and software that needs to be managed and maintained, the same number of IT pros can manage an increasing number of appliances while providing all of the necessary services.

## **Hyperconverged Infrastructure and Your Job**

Job security is a huge concern for many IT pros, particularly storage experts. While server, storage and networking specialists all need to familiarize themselves with hyperconverged infrastructure when it's introduced, storage pros may feel more threatened, especially if their current job involves spending a whole lot of time on relatively mundane tasks.

New, simple, more efficient systems may appear threatening, but there are different ways to think about this. The IT organization's goal is to ensure that services are still provided, and while the environment is changing, it rarely changes overnight. Transforming from legacy infrastructure to hyperconvergence may at first mean a single hyperconverged cluster is deployed, instead of a new SAN.

The evolution of the infrastructure is an opportunity for the storage team to take the lead in choosing the right hyperconverged infrastructure solution—one that has the storage characteristics the organization needs. From there, if you're a storage pro, it will mean expanding beyond the storage realm into the world of servers and the hypervisor.

It's not just storage people that need to adapt, though. We are headed toward a world of IT versatilists, where you are the "director" of your world, and software and hardware handle the tedious, complex, repetitive tasks behind the scenes. Where an IT admin used to focus on servers or storage, now you'll be able to broaden your skillset and help with networking, too, and get a better grasp on software-defined services and systems. IT admins who are more efficient (thanks to new efficient systems) can work on less mindless tasks and participate in more innovative projects that can contribute to the bottom line.

## Solution Maturity

There's no question that we are headed into a more software-defined future. The overall packaging of hyperconverged infrastructure—that is, packing everything into appliances and stacking them together— isn't revolutionary, but the software that makes it all work truly is. Be an agent for change.

There's a simple way to ensure that your intended solution can meet your needs: Test the heck out of it.

Every hyperconverged infrastructure vendor on the planet will allow you to do proof-of-concept testing of their solution.

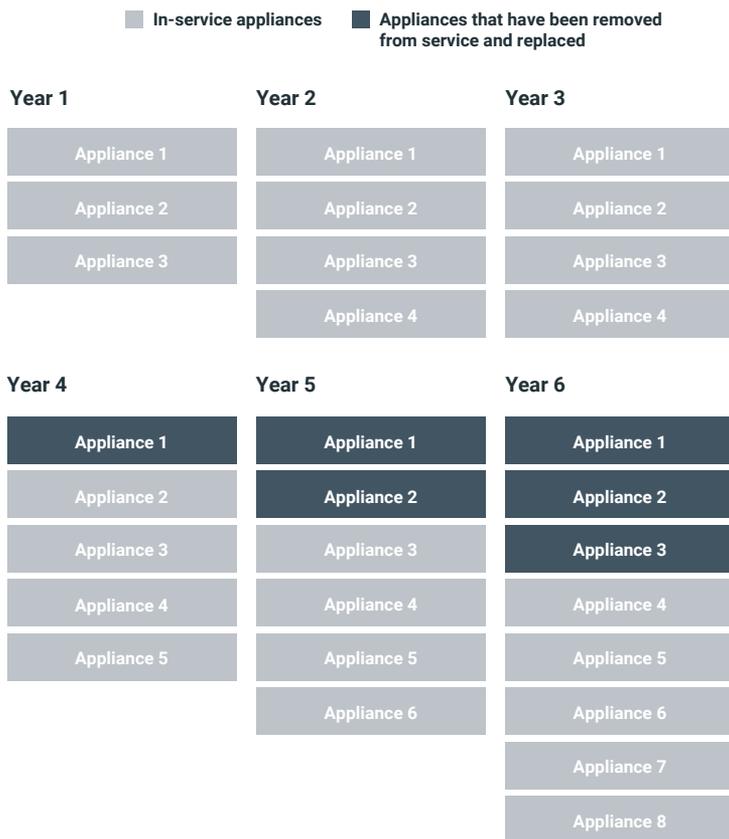
Your job is very clear. Test it. Put it through its paces and decide for yourself if the solution will meet your needs. Use real-world tests, not synthetic benchmarking tools to perform testing. After all, you won't be running synthetic tests in production; you'll be running real applications.

## Eliminating the Modern Refresh Cycle

The data center refresh cycle is out of control and is way too tied to capital budget dollars. Many organizations tackle individual resources separately. For example, they replace storage in Year 1 and then every four years thereafter. They replace a few servers each year. They replace supporting equipment, such as WAN accelerators and SSD caches, every three years. This mish-mash of replacement cycles creates confusion and doesn't always result in an equal distribution of budget dollars.

What if you could implement a real rolling refresh process and simply add new appliances as new business needs dictate? And, as a reasonable replacement cycle emerges, you simply remove an appliance a year and cycle in a new one. **Figure 2-2** gives you a look at once such scenario.

The figure also demonstrates how you scale this environment. As you need to add new appliances, you just add new appliances. As shown in



**Figure 2-2:** The hyperconvergence refresh process

**Figure 2-2**, you start with three appliances in Year 1 and scale to four appliances in Year 2 and Year 3. In Year 4, you maintain four appliances but begin to cycle in new appliances. The entire process consists of racking and stacking new appliances as well as removing old ones.

This is a pretty new way to handle every infrastructure element. You no longer have to juggle storage, servers, and other supporting systems, including WAN accelerators, backup storage, and the other myriad appliances that litter the data center.

Obviously, though, you can't just tear everything out of the data center today and replace it all with hyperconverged infrastructure. Like most things, you need to be able to phase in hyperconvergence as it is possible from a budgetary and logistical perspective.

There are a few ways you can do this:

- **Introductory project.** If you have a new need, such as VDI, test/dev, or a ROBO modernization initiative, you can introduce hyperconvergence by implementing your new project on hyperconverged infrastructure. From there, on your natural replacement cycles, begin to migrate your other workloads to the new environment. Eventually, you'll have fully phased out your traditional storage and server environment.
- **Pilot project.** If you have a workload that could use a boost in performance or efficiency, or a remote site that might benefit from modernization, use that as a model. A pilot is a great way to introduce the team to hyperconverged technology, and to learn about its capabilities.
- **Make coexistence a priority feature.** Some hyperconverged infrastructure solutions enable peaceful coexistence with legacy infrastructure. For example, you're able to use existing vSphere hosts with the storage in the hyperconverged cluster. Coexistence with legacy infrastructure isn't possible with all hyperconverged solutions, so choose carefully. Some solutions require you to implement hyperconvergence as a standalone silo of infrastructure.

## The Rise of the Versatilist

The IT "generalist" used to be considered a lesser role than those that required deeper subject matter expertise in a resource silo. The saying used to be "jack of all trades but master of none." Whether this was actually the case or not is open to debate, but the fact is that in today's world, a versatile generalist is exactly what the business needs in order to grow. It's also a great set of skills to have for IT pros.

As business requires increasing levels of nimbleness, silo-based solutions are eschewed in favor of ones that are more flexible and fit better into the overall equation. These kinds of solutions require knowledge of a broad range of technical concepts as well as of the business. The technical side of the equation isn't one that requires deep expertise, but rather understands how, at a high level, the pieces all fit together.

This high level thinking is in strong demand today and it's an incredible opportunity for IT pros to reshape their careers.

## Up Next

Understanding how to overcome inertia in order to make the right changes in your data center is really important. To learn more about how the technology impacts the IT pro and the organization and about key use cases for the technology, look for *The Gorilla Guide to Hyperconverged Infrastructure Implementation Strategies*.