

Making Everything Easier!™

SimpliVity Special Edition

Hyperconverged Infrastructure

FOR
DUMMIES®
A Wiley Brand

Learn to:

- Use hyperconverged infrastructure to simplify IT and reduce TCO
- Achieve economics and agility of large-scale web companies in your data center
- Solve capacity, performance, data protection, mobility, and management challenges

Brought to you by

simpli**ity** 

Scott D. Lowe

Consultant and industry veteran



About SimpliVity

To truly simplify IT in post-virtualization data centers, the “data problem” must be addressed. That means performance, capacity, mobility, manageability, and protection of data are ensured without driving up CAPEX and OPEX. This requires a new data architecture — one that takes a VM-centric approach, ensures adequate IOPS to fuel application requirements, and drives dramatic efficiency. SimpliVity’s Data Virtualization Platform delivered on hyperconverged infrastructure simplifies IT and operations, and uniquely enables data mobility and efficiency.

SimpliVity reduces IOPS, capacity, and bandwidth requirements, and provides native data protection, eliminating additional hardware and tools — resulting in a 3x TCO savings. SimpliVity’s hyperconverged infrastructure allows you to realize the economics, scale and agility of large-scale web companies without compromising the performance, reliability, and enterprise functionality you need in your data center.

Launched in 2013, SimpliVity is used by hundreds of enterprises and service providers worldwide.

Hyperconverged Infrastructure

FOR
DUMMIES®
A Wiley Brand

SimpliVity Special Edition

by Scott D. Lowe

FOR
DUMMIES®
A Wiley Brand

Hyperconverged Infrastructure For Dummies®, SimpliVity Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2014 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. SimpliVity and the SimpliVity logo are trademarks of SimpliVity Corporation. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact [Branded Rights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN 978-1-119-01394-5 (pbk); ISBN 978-1-119-01395-2 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Development Editor: Kathy Simpson

Senior Project Editor: Zoë Wykes

Acquisitions Editor: Katie Mohr

Editorial Manager: Rev Mengle

Business Development Representative: Sue Blessing

Project Coordinator: Melissa Cossell

SimpliVity contributors: Jesse St. Laurent, Lauren Whitehouse

Table of Contents

Introduction	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book.....	2
Beyond the Book.....	2
Where to Go from Here	2
Chapter 1: Hyperconverged Infrastructure Basics	3
Defining Hyperconvergence	3
Hyperconvergence Constructs	4
Chapter 2: Virtualization Challenges	5
Infrastructure Innovation.....	5
Underused Resources	6
Multiple Management Interfaces.....	7
Deployment Difficulty and Delays.....	8
Storage.....	10
IO blender	10
VDI boot storm.....	12
Mitigation	12
Multiple Touchpoints	13
Policy Misalignment	16
Chapter 3: Welcome to the SDDC	17
Virtualization	17
Automation	18
IT as a Service.....	19
Chapter 4: What Businesses Want from IT	21
Increased Efficiency.....	22
Using time better	22
Matching skills to tasks.....	22
Managing resources wisely	23
Meeting budgets	24
Reduced Risk	24
Improved Agility.....	25

Chapter 5: How the Cloud Is Changing IT27

Scaling and Sharing Resources	27
Software-centric design	28
Economies of scale.....	28
Resource flexibility.....	29
Enabling Automation	29
Abstracting policy from infrastructure.....	30
Taking a VM-centric approach.....	30
Understanding Cloud Economics	31

Chapter 6: Understanding Converged Infrastructure . . .33

The Evolution of Convergence	34
Integrated systems	34
Converged infrastructure	34
Hyperconverged infrastructure	35
Convergence Characteristics	36

**Chapter 7: Ten Things Hyperconvergence
Can Do for You.37**

Software Focus	38
Use of Commodity x86 Hardware	38
Centralized Systems and Management	38
Enhanced Agility	39
Scalability and Efficiency	39
Low Cost.....	40
Easy Automation	40
Focus on VMs	40
Shared Resources	41
Data Protection	42

**Chapter 8: Seven Ways to Apply
Hyperconvergence43**

Introduction



Over the past few years, a lot has happened to IT and to the technology industry. First and foremost, business leaders continue to demand more from their IT functions, including better and faster service, with seemingly endless expectations for increased efficiency.

In addition, several industry trends have had significant effects on enterprise IT:

- ✓ The proliferation of purpose-built devices
- ✓ The rise of the software-defined data center (SDDC)
- ✓ The emergence of cloud computing — or at least cloud principles and economics — as a key driver of enterprise IT
- ✓ The emergence of converged infrastructure
- ✓ The use of flash storage in certain enterprise applications

At the center of all these trends lies *hyperconvergence*, which is the third generation of a series of converged opportunities that have hit the market. Hyperconverged infrastructure (also known as hyperconvergence) is a data center architecture that embraces cloud principles and economics. Based on software, hyperconverged infrastructure consolidates server compute, storage, network switch, hypervisor, data protection, data efficiency, global management, and other enterprise functionality on commodity x86 building blocks to simplify IT, increase efficiency, enable seamless scalability, improve agility, and reduce costs.

About This Book

The hyperconvergence story has many chapters. In this small book, I discuss the trends that are leading modern IT to hyperconverged infrastructure. I also discuss both the technical and business benefits that come from implementing a data center based on hyperconverged infrastructure.

Foolish Assumptions

In this book, I assume that you know a little something about virtualization and cloud computing trends and models. As such, this book is written primarily for IT executives and managers such as Chief Information Officers (CIOs), Chief Technology Officers (CTOs), IT directors, and technical managers.

Icons Used in This Book

Throughout the book, you'll find several icons in the margins. Here's a rundown of what these icons mean.



Anything that has a Remember icon is well worth committing to memory.



The Technical Stuff icon indicates extra-credit reading. You can skip it if you like (but I hope you won't).



The Tip icon points out helpful information.



The Warning icon alerts you to risks of various kinds.

Beyond the Book

There's only so much I can cover in such little space. If you're eager to learn more about hyperconverged infrastructure after reading this book, visit www.hyperconverged.org.

Where to Go from Here

Like all *For Dummies* books, this book is designed to be read in any order you choose. Start with the chapter that interests you most, or read straight through. Hey, it's your book, so it's totally up to you.

Chapter 1

Hyperconverged Infrastructure Basics

In This Chapter

- ▶ Understanding what hyperconvergence is
- ▶ Discussing the many forms of hyperconvergence

Corporate technology undergoes a massive shift every so often as new models emerge to meet changing business needs. This chapter is about hyperconverged infrastructure, which is the culmination and conglomeration of a number of trends, all of which provide specific value to the modern enterprise.

Defining Hyperconvergence

So, what is hyperconvergence? At the highest level, hyperconvergence is a way to enable cloudlike economics and scale without compromising the performance, reliability, and availability you expect in your own data center. Hyperconverged infrastructure provides significant benefits:

- ✓ **Data efficiency:** Hyperconverged infrastructure reduces storage, bandwidth, and IOPS requirements.
- ✓ **Elasticity:** Hyperconvergence makes it easy to scale out/in resources as required by business demands.
- ✓ **VM-centricity:** A focus on the virtual machine (VM) or workload as the cornerstone of enterprise IT, with all supporting constructs revolving around individual VMs.
- ✓ **Data protection:** Ensuring that data can be restored in the event of loss or corruption is a key IT requirement, made far easier by hyperconverged infrastructure.

- ✓ **VM mobility:** Hyperconvergence enables greater application/workload mobility.
- ✓ **High availability:** Hyperconvergence enables higher levels of availability than possible in legacy systems.
- ✓ **Cost efficiency:** Hyperconverged infrastructure brings to IT a sustainable step-based economic model that eliminates waste.

Hyperconvergence is the ultimate in an overall trend of convergence that has hit the market in recent years. Convergence is intended to bring simplicity to increasingly complex data centers.

Hyperconvergence Constructs

Convergence comes in many forms. At its most basic, convergence simply brings together existing individual storage, compute, and network switching products into pre-tested, pre-validated solutions sold as a single solution. However, this level of convergence only simplifies the purchase and upgrade cycle. It fails to address ongoing operational challenges that have been introduced with the advent of virtualization. There are still LUNs to create, WAN optimizers to acquire and configure, and third-party backup and replication products to purchase and maintain.

Hyperconvergence is a ground-up rethinking of all the services that comprise the data center. With a focus on the virtual machine or workload, all the elements of the hyperconverged infrastructure support the virtual machine as the basic construct of the data center.

The results are significant and include lower CAPEX as a result of lower upfront prices for infrastructure, lower OPEX through reductions in operational expenses and personnel, and faster time-to-value for new business needs. On the technical side, newly emerging infrastructure generalists — people with broad knowledge of infrastructure and business needs — can easily support hyperconverged infrastructure. No longer do organizations need to maintain separate islands of resource engineers to manage each aspect of the data center.

To fully understand hyperconvergence, it's important to understand the trends that have led the industry to this point. These include post-virtualization headaches, the rise of the software-defined data center, and cloud computing.

Chapter 2

Virtualization Challenges

In This Chapter

- ▶ Keeping up with innovation
- ▶ Using resources efficiently
- ▶ Looking at management interfaces
- ▶ Discussing touchpoints
- ▶ Setting effective policies

Fact: Virtualization fundamentally and permanently changed IT and the data center. Today, most services are running inside virtual environments, and IT often takes a “virtualized first” approach to new application and service deployment. That is, administrators consider the virtual environment for running new applications rather than just building a new physical environment.

Although virtualization offers significant benefits, it also introduces challenges that IT must overcome to help propel the business forward. This chapter describes those challenges.

Infrastructure Innovation

Every time a startup company releases a great new product, enterprises scramble to implement that solution. The proliferation of purpose-built devices has created unnecessary complexity — and the result has been data center chaos.

Innovation is great, and we all want it to continue, but eventually, data centers have so much stuff that they become unmanageable. It’s time to clean out the closet, so to speak.

Over the past decade, IT departments have focused on solving the storage capacity problem, deploying all kinds of technologies to tame the capacity beast, such as WAN optimization and backup deduplication appliances. As a result, data efficiency technologies have become standard features of many different products.

But what happens when you put these products together in the data center? You end up constantly deduplicating and hydrating data as it moves between various devices. Storage deduplicates data; then you read the data to back it up, where it requires hydration (to return it to a state that the backup application understands) and often re-deduplicating it somewhere in the backup data path. The CPU cost to reprocess the same data is enormous, not to mention the bandwidth cost of all that hydrated data.



Deduplication is the process in which data is examined for common blocks. When identified, these common blocks are replaced with a tiny little pointer to the unique copy of the data already stored on disk — which takes up significantly less capacity when written to storage. Deduplication has a tremendous savings on storage capacity and, importantly, on input/output operations per second (IOPS) since fewer writes occur to disk. *Hydration* is reversing the deduplication process, such as when moving the data to a new system that doesn't support deduplicated data.

I discuss this challenge in more depth later in this chapter.

Underused Resources

Virtualization helped organizations consolidate many of their servers to run on a common platform: the hypervisor software layer. This move has helped IT departments make much better use of their server resources. Before virtualization, it was common for server utilization to average just 15 percent. Virtualization has pushed that average much higher. As a result, organizations now enjoy a much better return on their server investments. Moreover, they usually don't need to buy as many physical servers as they did in the past.

Virtualization has changed the game when it comes to server resources. Unfortunately, IT departments often need to maintain separate groups of people to manage separate hardware resources. One group manages storage, for example; another group manages the server side; a third group handles networking. When an issue arises, it's not uncommon to see a lot of finger-pointing.

Further, emerging workloads are creating resource challenges that push IT departments to develop infrastructure environments on a per-service basis. Virtual desktop infrastructure (VDI) environments, for example, have vastly different resource usage patterns from server virtualization projects. To meet user expectations with VDI, IT professionals often implement completely separate environments, from servers on down to storage.

Aren't resource islands the very problems that virtualization is meant to solve? These islands are among the biggest culprits of underutilization. Virtualization is supposed to result in a single resource pool from which resources are carved out to meet application needs, thereby maximizing the use of those resources.

Multiple Management Interfaces

Storage devices. Optimizers. Hypervisors. Load balancers. What do they have in common? Each of these disparate components features its own management interface. If you use multiple components, each with separate management consoles (and policy engines) rather than a single, centralized, easy-to-use administrative system, you may experience the following challenges:

- ✓ Vendors blaming each other when something goes wrong.
- ✓ The inability to scale your data center environment easily and linearly.
- ✓ Greater complexity due to policies and management being tied to IT components versus workloads.

Deployment Difficulty and Delays

Resource challenges represent the number-one reason why organizations continue to have problems deploying new applications and services. A close second is administrative overhead. Allow me to explain.



Converting to flat IT

The legacy data center is very delicate in many ways. Any change at any level has the potential to disrupt the overall structure. With lessons and tactics learned from the big cloud vendors, hyperconvergence vendors are replacing tiered and resource-siloed data centers with a much flatter IT structure. As practically all of the formerly separated data center hardware gets folded into the hyperconverged environment, the IT department needs to shift its focus, changing its resourcing structures and skill sets. Rather than having staff with deep subject matter knowledge in each resource area, hyperconvergence can give rise to infrastructure generalists.

Emerging infrastructure generalists don't have deep knowledge of individual resources, but they have broad knowledge of all resource areas. They don't need to have deep knowledge. In a hyperconverged world, the most complex stuff is handled under the hood. Infrastructure generalists need to

have enough broad knowledge to meet business needs and to manage the entire environment through a single administrative interface. In many ways, these people are far more application-focused than their island-based predecessors were. They just need to know how to apply infrastructure resources to meet individual application needs.

This development offers several bits of really good news for IT departments that have struggled to align IT operations with business needs:

- ✓ This new structure paves the way to eliminating the inefficient resource management islands that have emerged in IT.
- ✓ A flat data center managed by an infrastructure engineering group provides improved economies of scale compared with old resource islands.
- ✓ Infrastructure generalists are far closer to applications than specialists of old were.



Flash arrays as quick fixes

As flash-based (and really fast) storage has emerged at relatively affordable prices, new arms of the storage market have sprung up. One such arm provides storage arrays based solely on flash storage.

Although vendors in this all-flash space offer compelling products, many of these products are quick fixes designed to solve single-application problems — think VDI and Big Data analytics. For the vast majority of enterprise workloads, though, all-flash arrays are the very definition of throwing hardware at a performance problem. Storage solutions based on a combination of flash storage and spinning disks provide

a far more balanced and reasonable approach to meeting workload needs. In addition, the cost per gigabyte of flash storage is pretty expensive compared with other storage options. That said, for applications that need to achieve hundreds of thousands or even millions of IOPS in a tiny amount of rack space, all-flash arrays can't be beat. For everything else, consider more balanced storage options. Remember, a flash array really is a higher-performing storage array. It doesn't address the resource islands, infrastructure management, interoperability challenges, or scalability issues of the modern data center.

Multiple challenges exist on the resource front, including the following:

- **IO blender:** The consolidation of virtual machines (VMs) contributes to a random IO workload — each with its own pattern for reading/writing data to storage. I discuss the *IO blender* in detail later in this chapter.
- **Capacity:** Another challenge is ensuring adequate capacity as the organization grows. With resources divvied up and islands of resources strewn about the data center, managing ongoing capacity so that there's enough to go around becomes increasingly difficult.
- **Overhead:** Even if you have enough resources to deploy a new application (see the preceding bullet), the administrative overhead involved in the process presents its own challenges:
 - A new logical unit number (LUN) must be provisioned to support the new application. If tiers of

storage are involved, this process could require multiple steps.

- One or more new VMs must be provisioned.
- Networking for those new VMs has to be configured.
- Load balancers and wide-area network (WAN) optimization devices need to be managed to support the new VMs.
- Data protection mechanisms must be implemented for the new services.

Whew! That's a lot to do. All of it is time-consuming, and all of it involves different teams of people in IT. Good luck!

Storage

Virtualization is heavily dependent on storage, but this resource has wreaked havoc in companies that are working hard to achieve 100 percent virtualized status. Here's why.

Consider your old-school, physical server-based workloads. When you built those application environments, you carefully tailored each server to meet the unique requirements for each individual application. Database servers were awarded two sets of disks — one for database files and one for log files — with different redundant array of independent disks (RAID) structures. File servers got RAID 5 to maximize capacity, while still providing data protection.

Now consider your virtualized environment. You've taken all these carefully constructed application environments and chucked them into a single shared-storage environment. Each application still has specific storage needs, but you've basically asked the storage to sort out everything for you, and it hasn't always done a good job.

10 blender

In the old days, storage systems were optimized around LUN management. LUNs were replicated from a controller in one storage array to a LUN attached to a controller in a second array. The storage systems took snapshots of LUNs, and LUNs could be moved from one host to another host.

Today, servers have been replaced by VMs, and many VMs are running on a single host and many hosts are using a single LUN to store VMs. This means that the storage system has dozens (or hundreds) of logical servers (VMs) all stored in the same LUN. A single application or host or VM can no longer be managed from the storage system perspective.

A VM-centric platform cuts through this IO blender — a term that's been coined to describe environments in which mixed IO patterns are vying for limited storage resources — and allows you to optimize individual VMs. Policies can be applied to individual VMs. Performance can be optimized for individual VMs. Backups can be managed per VM, and replication is configured per VM.

Do you see a pattern emerging here?

When all your applications attempt to work together on the same LUN, the IO blender is created. Here are some ways that common services contribute to the IO blender:

- ✓ **Databases:** Databases feature random IO patterns. The system has to jump all over the disk to find what you're looking for.
- ✓ **Database log files:** Log files are sequential in nature. Usually, you just write to log files — again, sequentially.
- ✓ **Random file storage:** File servers are very random when it comes to IO. You never know when a user will be saving a new file or opening an old one.
- ✓ **Enterprise-level applications:** Applications such as Microsoft Exchange and SharePoint are sensitive in terms of storage configuration, and each application often includes a mix of random and sequential IO.
- ✓ **VDI:** VDI is one of the primary disruptors in the storage market. VDI storage needs are varied. Sometimes, you need only 10 to 20 IOPS per user. At other times, such as when you're handling boot storms and login storms, IOPS needs can skyrocket.

What the industry has done over the years is combine all these varied workloads. In other words, their very efforts to consolidate environments have created a storage monster.

Many storage-area network (SAN)-based storage environments suffer big-time problems due to this IO blender:

- ✓ Continued consolidation of VMs contributes to random IO workloads, each with its own pattern for reading and writing data to underlying storage.
- ✓ Highly random IO streams adversely affect overall performance as VMs contend for disk resources.

VDI boot storm

One situation that perfectly demonstrates a relatively new phenomenon in storage is the *VDI boot storm*, which occurs when many users attempt to boot their virtual desktops at the same time. The result: Storage devices can't keep up with the sheer number of requests.

It's the beginning of the day that really kills storage, though. As computers boot, the operating system has to read a ton of data and move it to memory so that the system can be used. Now imagine what happens when hundreds or thousands of users boot up their virtual desktops at the same time. Legacy storage systems crumble under the IO weight, and it can end up taking a long time for users to fully boot their systems.



The situation is largely mitigated by the use of solid-state storage as a caching layer. Adding this kind of service without considering the administrative inefficiencies that it introduces has been standard operating procedure for quite a while and is one of the main reasons why people implement resource islands when they want to do VDI.

Mitigation

Administrators can work inside their legacy environments in various ways to try to solve the serious IO issues that arise with shared workloads. Following are a few of those ways:

- ✓ Buy a separate environment to support each application.
- ✓ Buy complex stand-alone storage devices that include automatic tiering features.
- ✓ Buy multiple tiers of storage and manage them separately.



Read/Write IO

In traditional data center environments with shared storage, the difference in performance between reading and writing data is incredible. Reading generally is quick and can be accomplished with just a single IO operation. Writing data is a different story; it can take up to six IO operations to accomplish.

As administrators have moved from RAID 5 to RAID 6 for better data protection, they have introduced

additional overhead to the storage equation. A RAID 6 write operation requires six IOs to complete. The reason: It's not just the data that has to be written, but also the parity information that has to be written multiple times in RAID 6. RAID calculations also tend to require a lot of CPU overhead to perform the actual parity calculations needed for data protection.

What do these mitigation techniques have in common? They require administrators to overprovision storage, which requires more investment in storage hardware. They also require additional time by the administrator to configure and manage. Eventually, these models become unsustainable.

Multiple Touchpoints

Touching data many times in a virtualized environment isn't so great. Consider the following scenario: A legacy, but heavily virtualized, data center has many VMware vSphere servers connected to a SAN. The SAN has data deduplication mechanisms. The company backs up its data by using a local disk-to-disk-to-tape method; it also copies certain VMs to a remote data center each day. This way, the company maintains local backups and gains disaster recovery (DR) capabilities.

Quite a bit of redundancy is built into this scenario. Figure 2-1 examines the path the data travels as it wends its way through the various processes associated with the scenario.




Figure 2-1: Hyperconverged infrastructure requires far less CPU power and network bandwidth than nonconverged systems.

Every time the data has to be hydrated and then re-deduplicated as it makes its way to different components, the CPU must be engaged. Deduplication can be an expensive operation and constantly treating data in different locations has several drawbacks:

- ✓ Constant CPU use to treat data multiple times limits the number of VMs that can be run in the environment.
- ✓ Data takes more time to traverse the network because it isn't in reduced form as it moves between services.
- ✓ WAN bandwidth costs are significant as data travels across wide-area connections in unreduced form.

It gets worse. Many storage systems — including those related to data protection — use a *post-process* deduplication method, as opposed to what is known as an *inline* deduplication process. Post-process deduplication means that data is not deduplicated until after it's actually been written to disk. Here are the steps:

1. **Write data to disk undeduplicated.** Requires availability capacity and uses IOPS.
2. **Read data back from disk later.** The data then needs to be touched again by the post-process deduplication engine, consuming yet more IOPS and CPU resources.
3. **Invest CPU to deduplicate or compress.** Once read, the data then needs to be processed again using more CPU.

This means that the data will be replicated before deduplication and then all the dedupe work to save capacity must happen at both the primary and disaster recovery sites. This consumes additional resources, including CPU, IOPS, capacity, and network bandwidth. Post-process deduplication invests all these resources to get a reduction in storage capacity. The tradeoff isn't a positive one.

The results: greater costs and lower efficiency. The best outcome in any environment is to eliminate writes to disk before they even happen. In a hyperconverged environment, because of caching in RAM, many operations don't have to touch storage.

In the modern data center, data efficiency is about IOPS and WAN bandwidth, not storage capacity. Capacity has become plentiful as vendors release bigger drives (6TB and more!). Originally, data efficiency technologies were focused on the backup market. The objective was to provide an alternative for tape-based technology. In order to make the economics work, the primary objective was to fit more data on the disk while delivering the throughput needed for backups. In short, pack 20 pounds of data into a 5-pound bag. This was the right solution at the time.

While disks have gotten bigger, performance drives have barely improved. Organizations don't have a capacity problem; they have an IOPS problem, which manifests as poor performance. With the addition of DR in most customer environments, the demand for WAN bandwidth has increased and we also have a bandwidth challenge. Data reduction technologies, such as deduplication, are intended to address emerging resource challenges, including WAN bandwidth needs.

Given this reality, in a primary storage environment, infrastructure needs to optimize for performance and latency, not capacity and throughput. This requires new technology and an approach to data efficiency that is systemic — which is one of the hallmarks of hyperconverged infrastructure.

Inline deduplication provides the level of efficiency needed and consists of only two steps: process data and write data. Inline deduplication invests CPU just one time and gets a reduction in IOPS, WAN bandwidth, and storage capacity.

These are critical resources, but these resources are only conserved when the data efficiency is delivered inline.

In the modern data center, data efficiency is also about mobility and data protection, particularly when talking about online backup and restore. Data efficiency saves all the IO traditionally required to do backup and restore operations.

Policy Misalignment

In addition to facing the performance challenges of the post-virtualization world, virtualized organizations face policy challenges in both the physical and virtual worlds.

- ✓ **Physical:** Physical servers have a direct mapping from application to server to storage array to LUN to storage policy. This results in an environment where an application policy is directly linked to an internal construct of the storage array. There is no abstraction. This approach is what makes storage upgrades so complex. For example, a replication policy is applied to a LUN in storage array X at IP address Y and tells that LUN to replicate to storage array A at IP address B. Imagine the complexity of an array replacement when there are a couple of arrays in a couple of locations, and the replication policies are all tangled together. No wonder there are so many storage administrators in IT.
- ✓ **Virtual:** In the virtualized world, there are many applications on a host, and many hosts on a single LUN. It isn't efficient to apply a policy to a single LUN if that LUN represents the data for many applications (and hosts). In a hyperconverged environment, backup and replication policies are applied directly to individual applications (or VMs). There are no LUNs or RAID sets to manage. Replication policies specify a destination — in this case, a data center — that is abstracted away from the infrastructure. This allows an administrator to perform a platform upgrade without any policy reconfiguration or data migration, which increases efficiency and decreases risk.

Chapter 3

Welcome to the SDDC

In This Chapter

- ▶ Virtualizing everything
 - ▶ Introducing automation
 - ▶ Using IT as a Service
-

W

hat a concept. Software for the modern data center.

Consider the situation even five years ago. Legacy data centers were hardware-centric. Storage companies created their own chips and boxes to ship to customers. Networking vendors took a similar approach, creating individual circuits and arrays for their products. Although this approach wasn't necessarily bad, the resulting hardware products were relatively inflexible, and the flexible software layer played a supporting role.

In this chapter, I introduce the new data center standard: the software-defined data center (SDDC), in which software becomes the focus over hardware. Because SDDCs have several defining characteristics, including virtualization, automation, and the use of IT as a Service (ITaaS), I take a look at these characteristics in detail.

Virtualization

Every SDDC employs a high degree of virtualization, which goes beyond virtualizing servers. Everything is sucked up into the virtualization vacuum cleaner: storage, servers, and even

supporting services such as load balancers, wide-area network (WAN) optimization devices, and deduplication engines. Nothing is spared. This eliminates the islands of CPU, memory, storage, and networking resources that are traditionally locked within a single-purpose device, such as a backup to disk device, and creates a single shared resource pool for both business and infrastructure applications.

Virtualization abstracts the hardware components of the data center and overlays them with a common software layer: the virtualization layer, which manages the underlying hardware. The hardware can be a mix-and-match mess, but it doesn't matter anymore, thanks to the virtualization layer. All the data center administrator has to worry about is making sure that applications are running as expected. The virtualization layer handles the heavy lifting.

Automation

Many boardrooms today are asking companies to do more with less. One of the fastest ways to improve efficiency (and reduce costs) is to automate routine operations as much as possible.

Until now, many legacy IT architectures have been so varied and complex that automation has been a pipe dream only. SDDC brings the dream one step closer to reality.

Software-imposed normalization of the data center architecture permits higher degrees of automation. Moreover, the software layer itself is often chock-full of automation helpers, such as application programming interfaces (APIs). With this kind of assistance, automation becomes far easier to achieve.

IT as a Service

When resources are abstracted away from hardware and plenty of automation techniques are in place, companies often discover that they can treat many IT services as exactly that — services.

As they do with all other services, companies that use ITaaS have certain expectations:

- ✓ **Predictability:** The service should operate in a predictable way at a predictable cost. The SDDC can provide this conformity.
- ✓ **Scalability:** Business needs today may be very different from tomorrow, and the data center can't be a limiting factor when expansion becomes necessary. In fact, a data center should be an *enabler* of business expansion.
- ✓ **Improved utilization:** Companies expect to get maximum benefit from the services they buy. Because a hyperconvergence-powered SDDC is built on common components that eliminate the islands of resources traditionally trapped within infrastructure appliances, high utilization rates are exceedingly easy to achieve.
- ✓ **Fewer personnel:** With SDDC, a company can operate a data center with fewer people. The reason is simple: SDDC banishes traditional resource islands in favor of the new software-powered matrix.

Having fewer personnel translates directly to lower costs. In fact, research by Avaya suggests that an efficient SDDC can lower personnel costs from 40 percent of total cost of ownership to a mere 20 percent.
- ✓ **Reduced provisioning time:** A company that invests in SDDC expects to receive business benefits. SDDC offers agility and flexibility, which reduce provisioning times for the new services that business units require.



Hardware in a software world

When people hear the phrase *software-defined data center*, their first question usually concerns where the software for the SDDC is supposed to run. The answer is simple: The software layer runs on hardware.

But if SDDC is software-centric, why is hardware still required? Again, the answer is simple: You can't run a SDDC without hardware.

Most hardware in SDDCs looks quite different from hardware in traditional environments. Whereas legacy data centers have lots of proprietary hardware to manage myriad devices,

a SDDC uses mostly commodity hardware.

If a SDDC contains any proprietary hardware, the software leverages it to carry out important functions. In the world of hyperconvergence, this kind of hardware essentially becomes part of the data center's standard operations. Because it's identical hardware (and not unique to each device), it scales well as new appliances are added to the data center. Software is still top dog in such an environment, but without the hardware, nothing would happen.

Chapter 4

What Businesses Want from IT

In This Chapter

- ▶ Boosting efficiency
 - ▶ Keeping down risk
 - ▶ Getting — and staying — agile
-

Did you know that the IT department doesn't exist just to play with technology? Who knew? Apparently, it's far more important for this increasingly critical department to take its eye off the gear and turn a bit more toward the business.

This attention shift isn't just a good idea; it's a trend being driven hard by CEOs and business-unit leaders who have important needs to meet. Technology pros that yearn to stay ahead of the curve need to hone their business chops.

Expectations of high returns on large data center investments are climbing ever higher, and companies are much less willing to assume risk. They want a data center that has these three characteristics:

- ✓ Improves operational efficiency
- ✓ Reduces business risk
- ✓ Is flexible and agile enough to support changing business needs

This chapter examines these characteristics.

Increased Efficiency

Does your boss ever walk into your office and say something like this?

“Bob, we need to have a talk about your performance. You’re just too darn efficient, and we need you to dial that back a few notches. If you could do that on Saturday, that would be greeeeeeeat.”

I didn’t think so. If anything, IT departments are under growing pressure to increase efficiency. Improving efficiency generally means changing the way that IT operates — changes that involve anything from small-course corrections to major initiatives.



One of the greatest benefits of hyperconverged architecture is that it generates efficiency benefits without significantly disrupting operations.

Using time better

As poet Delmore Schwartz put it, “Time is the fire in which we burn.” For those who slog through repetitive, mundane tasks every day, truer words were never written. When it comes to business, any time wasted on mundane work really is burned time — time that could have been spent advancing business objectives.

Management wants IT to spend its time wisely. Traditional IT operations simply won’t cut it anymore. Neither will prolonged product evaluation and integration processes or extended return-on-investment metrics. IT needs to be faster and leaner than ever before.

Matching skills to tasks

Step back for a second to think about what the IT staff really has to deal with on a day-to-day basis: servers, hypervisors, storage devices, network accelerators, backup software, backup appliances, replication technology, and a whole lot more. Forget for a moment about the physical effects of this plethora of equipment on the data center. Instead, consider the human toll.

Every one of these devices has a separate administrative console that operators have to learn. Also — let's face reality — not every device plays nicely with every other device.

When each device requires vastly different sets of skills to operate, each skill requires ongoing training. Even when you can get a few people in IT trained on everything in the data center, at some point those people may move on, and you may have trouble finding new employees who have the same set of skills.

In addition, every time you bring a unique resource into the environment, you need staff to manage it. As that resource grows, you may need even more staff to keep up with the workload. In essence, you're creating resource islands as you forge ahead.



Resource islands are inherently inefficient. The broader you can make the IT environment, the easier it is to achieve operational economies of scale.

The bottom line: IT staffs are being crushed under the weight of legacy infrastructure. Each unique resource requires unique skills, and companies aren't adding IT staff at a pace that keeps up with technical needs.

Managing resources wisely

The laws of physics are immutable forces in the land of IT, and these natural laws become apparent when you walk around a data center:

- ✓ You find that two objects can't occupy the same space at the same time, which is why equipment occupies separate racks.
- ✓ You discover the laws of thermodynamics as you walk behind a rack and feel the blast of hot air, and then walk under the cooling equipment and feel a chill.
- ✓ Finally, you encounter electromagnetism as you watch electricity-sucking hardware whirl away. (On the plus side, all the blinking LEDs put on a pretty impressive light show.)

All these physical resources — space, power, and cooling — require cash to keep running. Every piece of equipment added to a data center requires the use of all these resources.



Continuing to add equipment without considering your use of resources will do nothing to increase your overall efficiency.

Meeting budgets

Think back to your first virtualization project. Why did you undertake it? I'm willing to bet that you wanted to make better use of your resources and improve on what used to be a 15 percent average utilization rate for servers.

If you're not using all your equipment at a reasonable level, you're leaving money on the table. You may not be enjoying the maximum return on an incredibly expensive investment.



You can make your IT budget more efficient and serve the business better by rethinking the way you provide data center services. Don't think about each individual resource as its own island. Instead, think at a higher level. Rather than focusing on individual resources, focus on overall scale of all resources as your standard operating procedure.

Reduced Risk

Risk can worm its way into an otherwise flawless IT system in several places:

✓ **Procurement:** With so much hardware to maintain in the data center, staying on top of all the little details can be difficult. Before you purchase any new equipment, ask yourself questions such as these:

- Are you sure that the storage-area network (SAN) you've selected has enough capacity in terms of terabytes *and* performance in terms of IO operations per second (IOPS)?
- If you're expanding an existing system, will the expansion create a risk of downtime or data loss?

- If you're upgrading an existing storage array, are all the component firmware revisions the same or supported by the new hardware?

✓ **Operations:** Generally, you need two of everything in a data center to maintain the levels of availability that the business expects. If you don't, you run the risk of prolonged outages, which management tends to dislike.

Redundancy is the norm in IT, but it's an expensive standard. Moreover, it requires personnel who have specialized skills to maintain the different high-availability features that ship with each product.

✓ **Data protection:** Too many companies don't plan their data protection mechanisms carefully or rely on many services provided by many companies. As a result, the blame game typically occurs when something unexpected happens. When recovery is Job One, no one wants vendors to fight about who's at fault. Keep in mind: Data protection isn't about backup; it's about *recovery*.



Chief Information Officers (CIOs) and IT staff want — and need — to lower risk in their organizations. Systems and applications must remain highly available, and data must be safe. Unfortunately, as more diverse hardware is installed, achieving these critical goals becomes more difficult.



Companies can reduce risk by adopting a hyperconverged architecture. Hyperconverged systems include all the components required to make a data center operate without the complexity of legacy solutions.

Improved Agility

Getting a new product or service to market quickly often results in long-term business advantages. Today's commerce markets are cutthroat places, and IT issues can't be allowed to get in the way of critical business operations. It's one of the reasons that the public cloud is such a draw for business users.

Any lag in delivering what the business needs could result in the business taking IT matters into its own hands — the so-called “Shadow IT” problem, which happens more often than you think. If a business constituent requests a function

or service that IT can't be as responsive to, then the business constituent, with the swipe of a credit card, can procure a cloud-based service without IT's involvement — or IT's knowledge that it exists. Shadow IT practices have the potential to introduce risk.

To meet the demands of business units for more speed and agility, many IT departments simply build larger and larger infrastructures, which are inflexible and difficult to manage. Eventually, such a system turns into a house of cards, and the slightest problem can topple it. This scenario is hardly the hallmark of an agile IT infrastructure.



Agility and *speed* are two mantras that every IT person should adopt. It's critical to ensure that the infrastructure is agile so that IT can quickly and easily deploy new applications and services in response to business demands.

Chapter 5

How the Cloud Is Changing IT

In This Chapter

- ▶ Employing shared, scalable resources
- ▶ Emphasizing automation
- ▶ Managing purchases wisely

Virtualization is just one major trend impacting IT. It's tough to deny that the cloud has also affected IT. Consider this:

- ✔ Business units and IT departments can acquire services simply by using a credit card.
- ✔ Major cloud service providers like Google and Facebook are changing expectations of how a data center should operate because their massive environments are nothing like legacy data centers. Even though most organizations don't need anything of that scale, the best architectural design elements from these clouds have been brought to the hyperconverged world and packaged for affordability.

For hyperconverged infrastructure, only the second trend is really pertinent and is what this chapter is about.

Scaling and Sharing Resources

The hallmarks of Google's and Facebook's environments are, among other things, sheer scalability and reasonable economics. Many of these cloud principles have been adapted for use

in smaller environments and packaged in hyperconverged products that any company can buy.

Software-centric design

As I mention in Chapter 3, software overtaking hardware in the data center has the potential to lead to very good things. Companies such as Google discovered this potential years ago and tamed their hardware beast by wrapping it inside software layers. A data file inside Google is managed by the company's massively distributed, software-based global file system. This file system doesn't care about the underlying hardware. It simply abides by the rules built into the software layer that ensures that the file is saved and with the right data protection levels. Even with expansion of Google's infrastructure, the administrator isn't concerned about where that file resides.

Economies of scale

In a legacy data center environment, growing the environment can be expensive due to the proprietary nature of each individual piece of hardware. The more diverse the environment, the more difficult it is to maintain.

Commodity hardware

Companies such as Google and Facebook scale their environments without relying on expensive proprietary components. Instead, they leverage commodity hardware. To some people, the word *commodity*, when associated with the data center environment, is a synonym for *cheap* or *unreliable*. You know what? To a point, they're right.



When you consider the role of commodity hardware in a hyperconverged environment, however, keep in mind that the hardware takes a back seat to the software. In this environment, the software layer is built with the understanding that hardware can — and ultimately *will* — fail. The software-based architecture is designed to anticipate and handle any hardware failure that takes place.

Commodity hardware isn't cheap, but it's much less expensive than proprietary hardware. Also, it's interchangeable with other components. A hyperconvergence vendor can switch

its hardware platform without recoding the entire system. By leveraging commodity hardware, hyperconvergence vendors ensure that their customers get affordable hardware without disruption because change is so quick and easy.

Bite-sized scalability

Think about how you procure your data center technology now, particularly when it comes to storage and other non-server equipment. For the expected life cycle for that equipment, you probably buy as much horsepower and capacity that you'll need, with a little extra "just in case" capacity.

How long will it take you to use all that pre-purchased capacity? You may never use it. What a waste! But on the other hand, you may find it necessary to expand your environment sooner than anticipated. Cloud companies don't create complex infrastructure update plans each time they expand. They simply add more standardized units of infrastructure to the environment. This is their scale model; it's all about being able to step to the next level of infrastructure in small increments, as needed.

Resource flexibility

Hyperconverged infrastructure takes a bite-sized approach to data center scalability. Customers no longer need to expand just one component or hardware rack at a time; they simply add another appliance-based node to a homogenous environment. The entire environment is one huge virtualized resource pool. As needs dictate, customers can expand this pool quickly and easily, in a way that makes economic sense.

Enabling Automation

Do you think that the typical cloud provider spends an inordinate amount of time rebuilding individual policies and processes each time it makes changes or adds equipment to its data center? Of course not. In the cloud, change is constant, so ensuring that changes are made without disruption is critical. In the world of enterprise IT, things should work the same way. A change in data center hardware shouldn't necessitate reconfiguration of all your virtual machines and policies.

Abstracting policy from infrastructure

Since hardware isn't the focus in the software-defined data center (SDDC; see Chapter 3), why would you write policies that target specific hardware devices? Further, because enterprise workloads leverage virtual machines (VMs) as their basic constructs, why should a VM be beholden to policies tied to underlying infrastructure components?

Consider a scenario in which you define policies that move workloads between specific logical unit numbers (LUNs) for replication purposes. Now multiply that policy by 1,000. When it comes time to swap out an affected LUN, you end up with LUN-to-LUN policy spaghetti. You need to find each individual policy and reconfigure it to point to new hardware.

Policies should be far more general in nature, allowing the infrastructure to make the granular decisions. Instead of getting specific in a LUN policy, for example, policies should be as simple as, "Replicate VM to San Francisco."

Why is this good? With such a generalized policy, you can perform a complete technology refresh in San Francisco without migrating any data or reconfiguring any policies. Since policy is defined at the data center level, all inbound and outbound policies stay in place. Beautiful.

Taking a VM-centric approach

The workload takes center stage in the cloud. In the case of enterprise IT, these workloads are individual VMs. When it comes to policies in cloud-based environments, the VM is the center of the world. It's all about applying policies to VMs — not to LUNs, shares, datastores, or any other constructs. Bear in mind the plight of the VM administrator, who is VM-centric. Why wouldn't the administrator assign backup, quality-of-service, and replication policies to a VM?

The need to apply policies across individual resource domains creates fundamental operational issues in IT. In the cloud and in hyperconvergence, policies are simply called *policies*. There are no LUN policies, caching policies, replication policies, and so on — just policies.

Understanding Cloud Economics

Cloud providers and enterprise IT organizations operate their environments using very different economic models. Chief Information Officers (CIOs) expect enterprise IT infrastructure to last many years, so they buy enough capacity and performance to last that long. In many cases, however, the full potential of the infrastructure buy is never realized. CIOs often overbuy to ensure that capacity lasts the full life cycle.

Conversely, by design or by mistake, CIOs often underbuy infrastructure. The organization then needs to buy individual resources when they begin to run low. This leads to watching individual resources constantly, reacting when necessary, and hoping that your existing product doesn't have end-of-life status.

Now consider cloud vendors, who don't make one huge buy every five years. Doing so would be insane in a few ways:

- ✓ A *lot* of hardware would have to be purchased up front.
- ✓ Accurately planning three to five years' worth of resource needs in these kinds of environments may be impossible.

Instead, cloud organizations pay as they grow. Operational scale and homogeneity are core parts of cloud providers' DNA, so they simply add more standard resources as needed.

The public cloud is highly appealing to the enterprise. The instant-on service is elastic and only costs a few cents an hour. What could be better? It's not for all applications; it presents major challenges for many, particularly when it comes to cost predictability. The true cost of public cloud increases dramatically when compared against the cost of predictable storage performance, high availability, backup, disaster recovery, private networking, and more. IT ends up paying for a server that is running at 15 percent utilization, and the cloud provider is benefiting from packing those VMs onto a single host.

IT applications are designed expecting high-availability infrastructure, disaster recovery, backup and recovery, and other necessary services (this is why internal IT isn't like Facebook

and Google). IT applications place a different set of demands on the infrastructure. Therefore, any hyperconverged infrastructure *must* deliver on these requirements.

Business units may not understand these nuances and may be compelled to buy cloud services without having a grasp on the full picture. This rise of Shadow IT — where non-IT units create their own systems (see Chapter 4) — is real, and the cloud is complicit in enabling this trend. Shadow IT exists, however, either because IT isn't able to provide the kinds of services business units demand or because they aren't responsive enough. So, these units turn to the cloud to meet individual needs, giving rise to fragmented services, potential security breaches, and overall data-quality issues.

Hyperconvergence brings cloud-type consumption-based infrastructure economics and flexibility to enterprise IT without compromising on performance, reliability, and availability. Rather than making huge buys every few years, IT simply adds building blocks of infrastructure to the data center as needed. This approach gives the business much faster time-to-value for the expanded environment.

Figure 5-1 shows integrated systems (which scale via large steps; see Chapter 6) and hyperconvergence (which scale via bite-sized steps; see Chapters 6 and 7), with wasted resources above the wavy line.




Figure 5-1: The unit of scale has a significant impact on economics.

Chapter 6

Understanding Converged Infrastructure

In This Chapter

- ▶ Seeing how convergence evolved
- ▶ Introducing hyperconverged infrastructure
- ▶ Making the most of differences

The IT infrastructure market is undergoing unprecedented transformation. The most significant transformation is reflected by two major trends: convergence and software-defined data centers (SDDCs). Both trends are responses to the IT realities of infrastructure clutter, complexity, and high cost; they represent attempts to simplify IT and reduce the overall cost of infrastructure ownership. Today's infrastructure environments are typically comprised of 8 to 12 hardware and software products from as many vendors, with each product offering a different management interface and requiring different training.

Each product in this type of legacy stack, however, is grossly overprovisioned, using its own resources (CPU, DRAM, storage, and so on) to address the resident applications' intermittent peak workloads. The value of a single shared resource pool, offered by server virtualization, is still limited to the server layer. All other products are islands of overprovisioned resources that aren't shared. Therefore, low utilization of the overall stack results in the ripple effects of high acquisition, space, and power costs. Simply put, too many resources are wasted in today's legacy environments.

This chapter explores a leading solution: convergence, which ultimately leads to hyperconvergence.

The Evolution of Convergence

The following sections describe the evolution of convergence over the past few years.

Integrated systems

The earliest infrastructure convergence solutions have complete network, compute, storage, and virtualization capabilities, but they're simply conglomerations of existing hardware and software, with little to no actual innovation in product features to be leveraged. Many people look at the solutions offered by major original equipment manufacturers as being ways to lock customers into their technology stacks.

Integrated systems do offer a few benefits. Most notably, customers gain a single point of contact for their infrastructure, from purchase to end of life. These systems are always tested and almost always arrive at the customer site fully racked and cabled, so they're ready to go.

On the downside, these systems often have a big step size. When you need more power, you have to buy a big chunk of infrastructure. Also, these products usually don't solve the serious challenges that so many organizations face (see Chapter 2).

Converged infrastructure

The converged infrastructure products combine the server and storage components in a single appliance, effectively eliminating the need for dedicated storage area network (SAN)-based storage.

These systems provide a localized single resource pool solution, offering simplified management and faster time to deployment. They have effectively virtualized the storage layer and allowed it to run in the virtualization platform. Overall acquisition cost is lower, and management (at least,

for the server and storage resources) is simplified. With these systems, overall resource utilization is higher than with a legacy island-based infrastructure.

Converged infrastructure has some limitations, however:

- ✓ The systems include just the server and storage resource components.
- ✓ The fundamental data management challenges have not been solved. It is the functionality of a traditional storage array migrated into the virtualization platform.
- ✓ Resource ratios (such as CPU:storage:network) are fixed, making them less flexible than some organizations require.
- ✓ The products can't always be used by existing infrastructure. In other words, you can't use a converged infrastructure appliance's storage from existing legacy systems. In essence, you're forced to create a separate resource island.

For these reasons, converged infrastructure systems don't efficiently address performance problems in the legacy infrastructure.

Likewise, on the data front, the systems don't address all data problems, because not all data efficiency appliances are converged. Management may be improved, but it's not unified global management.

Hyperconverged infrastructure

Enter hyperconvergence, also known as *hyperconverged infrastructure*, which represents the logical next step in the evolution of infrastructure convergence. Hyperconvergence delivers simplification and savings by consolidating all required functionality into a single infrastructure stack running on an efficient, elastic pool of x86 resources. The underlying data architecture has been completely reinvented, allowing data management to be radically simplified. As a result, hyperconverged infrastructure delivers on the promise of the SDDC at the technological level. It also carries forward the benefits of convergence, including a single shared resource pool.

Hyperconvergence goes far beyond servers and storage, bringing into the convergence fold many services that make some legacy services obsolete, including the following:

- ✓ Data protection products (backup, replication)
- ✓ Deduplication appliances
- ✓ Wide-area network (WAN) optimization appliances
- ✓ Solid-state drive (SSD) arrays
- ✓ SSD cache arrays
- ✓ Public cloud gateways
- ✓ Replication appliances or software

In Chapter 7, I delve into how hyperconvergence takes convergence to the next level in the data center and provides a plethora of benefits to both IT and to the business.

Convergence Characteristics

The preceding convergence options build on one another, with each having key differences. Figure 6-1 describes the high-level characteristics that define each convergence type.

	Technical Attributes			Organizational Benefits	
	Data Efficiency	Single Shared Resource Pool	Global Management	TCO Improvements	Simplification
Integrated systems	☐	☐ Resource pooling limited to server layer	☐	☐	☐ Some time to deployment and administrative gains
Convergence	☐	☒ Limited to primary server and storage resources; other resources not included	☐	☒ TCO gains primarily due to reduction of legacy gear; does not address backup, replication, and DR	☒ Fewer products to manage
Hyperconvergence	☒ Data architecture begins with one-time deduplication, compression, and optimization of data	☒ All data center resources are brought into the resource stack	☒ Complete management of all infrastructure resources and virtual machines; single point of administration	☒ Major TCO gains through reduction of hardware resources, streamlined operations, and automation	☒ Reduces hardware littered across data centers, eases management, VM-centricity
☐ Not Supported ☒ Partially Supported ☒ Fully Supported					

Figure 6-1: Compares improvements as convergence has evolved.

Each characteristic is critical to realizing all the traits that business demands of IT in the modern era: lean, mean, and green.

Chapter 7

Ten Things Hyperconvergence Can Do for You

In This Chapter

- ▶ Focusing on software
- ▶ Centralizing management
- ▶ Enhancing data protection

How does hyperconvergence bring together all the important trends that enterprise IT struggles to handle? Here's a look:

- ✓ Hyperconvergence is the embodiment of the software-defined data center (SDDC; see Chapter 3). Based in software, it provides the flexibility and agility that business demands from IT.
- ✓ Cloud operators have their economic model figured out. Hyperconvergence brings to enterprise IT a cloud-like economic model that provides faster time to value for data center expenditures and lower total cost of ownership for the entire solution. Hyperconvergence offers the economic benefits of the cloud while delivering the performance, high availability, and reliability the enterprise demands.
- ✓ Flash solves performance issues but has to be used carefully. Hyperconverged options combine flash and spinning disk to provide the ideal blend of capacity and performance, allowing IT to eliminate resource islands (see Chapter 2).

- ✓ The converged infrastructure market provides a single-vendor approach to procurement, implementation, and operation. There's no more vendor blame game, and there's just one number to call when a data center problem arises.

In this chapter, I dive a bit deeper into hyperconvergence, showing you ten ways that it solves the challenges inherent in virtualized data centers (see Chapter 2).

Software Focus

Hyperconvergence is the epitome of the software-defined data center (SDDC), discussed in Chapter 3. The software-based nature of hyperconvergence provides the flexibility required to meet current and future business needs without having to rip and replace infrastructure components. Better yet, as vendors add new features in updated software releases, customers gain the benefits of those features immediately, without having to replace hardware.

Use of Commodity x86 Hardware

Commodity hardware equals lower cost. The software layer is designed to accommodate the fact that hardware will eventually fail. Customers get the benefit of these failure avoidance/availability options without having to break the bank to get the hardware.

Financial officers really like the opportunity to save a few bucks. The business gets to enjoy better outcomes than legacy data center systems offer, often at a much lower cost.

Centralized Systems and Management

In hyperconvergence, all components — compute, storage, backup to disk, cloud gateway functionality, and so on — are combined in a single shared resource pool with hypervisor

technology. This simple, efficient design enables IT to manage aggregated resources across individual nodes as a single federated system.

Mass centralization and integration also happen at the management level. Regardless of how widespread physical resources happen to be, hyperconverged systems handle them as though they were all sitting next to one another. Resources spread across multiple physical data centers are managed from a single, centralized interface. All system and data management functions are handled within this interface, too.

Enhanced Agility

Agility is a big deal in modern IT. Business expects IT to respond quickly as new needs arise, yet legacy environments force IT to employ myriad resources to meet those needs. Hyperconverged infrastructure enables IT to achieve positive outcomes much faster.

Part of being agile is being able to move workloads as necessary. In a hyperconverged world, all resources in all physical data centers reside under a single administrative umbrella (see the preceding section). Workload migration in such environments is a breeze, particularly in a solution that enables consistent deduplication as a core part of its offering. Reduced data is far easier to work with than fully expanded data and helps IT get things done faster.

Scalability and Efficiency

Hyperconvergence is a scalable building-block approach that allows IT to expand by adding units, just like in a LEGO set. Granular scalability is one of the hallmarks of this infrastructure. Unlike integrated systems products, which often require large investments, hyperconverged solutions have a much smaller step size. *Step size* is the amount of infrastructure that a company needs to buy to get to the next level of infrastructure. The bigger the step size, the bigger the up-front cost.



The bigger the step size, the longer it takes to fully utilize new resources added through the expansion. A smaller step size results in a far more efficient use of resources. As new resources are required, it's easy to add nodes to a hyperconverged infrastructure.

Very Borg-like, eh (in a good way)?

Low Cost

Hyperconverged systems have a low cost of entry compared with their integrated system counterparts and legacy infrastructure. (For more on integrated systems, see Chapter 6.)

Easy Automation

Automation is a key component of the SDDC (see Chapter 3) and goes hand in hand with hyperconvergence. When all resources are truly combined and when centralized management tools are in place, administrative functionality includes scheduling opportunities as well as scripting options.

Also, IT doesn't need to worry about trying to create automated structures with hardware from different manufacturers or product lines. Everything is encapsulated in one nice, neat environment.

Focus on VMs

Virtualization is the foundation of the SDDC (see Chapter 3). Hyperconverged infrastructure options use virtual machines (VMs) as the most basic constructs of the environment. All other resources — storage, backup, replication, load balancing, and so on — support individual VMs.

As a result, policy in the hyperconverged environment also revolves around VMs, as do all the management options available in the system, such as data protection policies, which are

often defined in third-party tools in legacy environments. With hyperconvergence, integrated data protection policies and control happen right at VM level. (I discuss data protection later in this chapter.)

VM-centricity is also apparent as workloads need to be moved around to different data centers and between services, such as backup and replication. The administrator always works with the virtual machine as the focus, not the data center and not underlying services, such as storage.

Shared Resources

Hyperconvergence enables organizations to deploy many kinds of applications in a single shared resource pool without worrying about the dreaded IO blender effect (see Chapter 2), which wrecks VM performance.

How does hyperconvergence make this type of deployment possible? Hyperconverged systems include multiple kinds of storage — both solid-state storage and spinning-disk — in each appliance. A single appliance can have multiple terabytes of each kind of storage installed. Because multiple appliances are necessary to achieve full environment redundancy and data protection, there's plenty of both kinds of storage to go around. The focus on the VM in hyperconverged systems also allows the system to see through the IO blender and to optimize based on the IO profile of the individual VM.

Hyperconverged infrastructure's mix of storage enables systems to handle both random and sequential workloads deftly. Even better, with so many solid-state storage devices in a hyperconverged cluster, there are more than enough IO operations per second (IOPS) to support even the most intensive workloads — including virtual desktop infrastructure (VDI) boot and login storms (see Chapter 2).

The shared resource pool also enables efficient use of resources for improved performance and capacity, just like those very first server consolidation initiatives that you

undertook on your initial journey into virtualization. Along the way, though, you may have created new islands thanks to the post-virtualization challenges discussed earlier. Resource islands carry with them the same utilization challenges that your old physical environments featured. With hyperconvergence, you get to move away from the need to create resource islands just to meet IO needs of particular applications. The environment itself handles all the CPU, RAM, capacity, and IOPS assignments so that administrators can focus on the application and not individual resource needs.

The business benefits as IT spends less while providing improved overall service. On the performance front, the environment handles far more varied workloads than legacy infrastructure can. IT itself performs better, with more focus on the business and less on the technology.

Data Protection

Although it's not always the most enjoyable task in the world, protecting data is critical. The sad fact is that many organizations do only the bare minimum to protect their critical data. There are two main reasons why: Comprehensive data protection can be really expensive and really complex.

To provide data protection in a legacy system, you have to make many decisions and purchase a wide selection of products. In a hyperconverged environment, however, backup, recovery, and disaster recovery are built in. They're part of the infrastructure, not third-party afterthoughts to be integrated.

The benefits of hyperconvergence are clear:

- ✓ Comprehensive backup and recovery, and affordable disaster recovery.
- ✓ Efficient protection without data rehydration and re-deduplication — and the inefficient use of resources that result.
- ✓ A single centralized console that enables IT to respond quickly.

Chapter 8

Seven Ways to Apply Hyperconvergence

In This Chapter

- ▶ Consolidating and modernizing
- ▶ Preparing for the worst

The great thing about hyperconvergence is that it doesn't necessarily require you to replace existing infrastructure in order to be of immediate value. Here are seven ways that you can gain hyperconvergence benefits starting now:

- ✓ **Consolidating servers and data center.** Are you tackling a new consolidation project or building a new data center? Leading hyperconvergence vendors provide products that integrate seamlessly with your existing environment. The right hyperconverged solution can solve your immediate challenges and produce major benefits.
- ✓ **Modernizing technology smoothly.** The beauty of hyperconvergence is its nondisruptive implementation. The hyperconverged environment is part of your overall environment, so you can phase in new architecture while you phase out the old, implementing and expanding as funds allow. If applications in the legacy environment need the storage performance provided by the hyperconverged environment, they can leverage those resources.
- ✓ **Deploying new tier-1 applications.** Is your existing environment suitable for new tier-1 workloads? Rather than simply throwing more resources at an outdated environment, deploy the new workload in a hyperconverged environment to gain the inherent operational benefits. As time goes on, you can start bringing the rest of your

infrastructure into the same architecture with easy-to-add, LEGO-like efficiency.

- ✔ **Deploying VDI.** Resource islands are established in large part due to virtual desktop infrastructure (VDI) needs. However, the way that IT implements these resource islands means that they're forever separate. By deploying your VDI project in a hyperconverged infrastructure, you don't face resource challenges that require you to create these islands. When the VDI project is out of the way and the rest of your environment becomes eligible for renewal, you can slide everything to the hyperconverged environment with ease. For more on VDI, see Chapter 2.
- ✔ **Managing sites remotely.** In a hyperconverged environment, the entire infrastructure is controlled by a single management system. Remote resources can be managed as though they were local resources. There's no need to have remote personnel perform manual operations such as running backup jobs, or creating logical unit numbers (LUNs) or quality-of-service policies. The data efficiency technology allows backups to be simplified in the remote office and offsite copies to be automatically sent to headquarters, another remote office, or even the cloud. This enables centralization of administrative resources, providing staffing economies of scale.
- ✔ **Performing testing and development.** Many organizations operate test and development (test/dev) environments so that bad code isn't released into production. Hyperconvergence supports test/dev and production needs, with management tools that can help you create logical separations between these functions.

Unfortunately, many organizations give short shrift to test/dev and run it on lower-class hardware, which doesn't really make sense. This greater IT agility will keep developers in house instead of building their own Shadow IT in the public cloud.
- ✔ **Modernizing backup and implementing disaster recovery.** If you don't do a good job with either backup or disaster recovery, run — don't walk — toward hyperconvergence as your infrastructure architecture. Hyperconverged infrastructure eliminates the complexity that can be inherent in these operations. In IT, simplicity is the new mantra, and hyperconvergence is one of the simplest ways to achieve your backup and disaster recovery goals.

Cloud Economics. Without The Cloud.

SimpliVity's hyperconverged infrastructure allows you to realize the economics and agility of large-scale web companies—right in your own data center. With SimpliVity, you address your most challenging performance, capacity, and scale issues with a pay-as-you-grow building block approach to IT infrastructure.

SimpliVity takes a VM-centric approach to managing your environment and delivers a data architecture that deduplicates, compresses and optimizes data inline to reduce your capacity footprint and IOPS, addressing your most annoying—and costly—capacity and performance issues.

SimpliVity's hyperconvergence simplifies IT, reduces TCO by 3x, and provides your IT organization with the agility to keep pace with business demands. No more overprovisioning as an insurance policy to meet IOPS demands. No more management at the hardware component level. No more sacrificing the resiliency, data protection and other enterprise IT functionality you know and love.

For more information, visit [**www.simplivity.com/dummies**](http://www.simplivity.com/dummies)



Transform your data center with hyperconvergence to simplify IT and lower TCO

You've converged your workloads, so why not your infrastructure too? Collapsing the infrastructure stack into a single consumable hyperconverged IT system streamlines operations. Hyperconvergence resolves interoperability issues, improves management, reduces costs, and brings cloudlike efficiency to your enterprise. Learn the ins and outs.

- **Hyperconverge everything** — understand why stopping at just servers and storage leaves opportunities on the table
- **Take a VM-centric approach** — leave siloed infrastructure and hardware-level management behind
- **Introduce a new economic model** — gain cloudlike consumption-based economics in your data center
- **Create new levels of efficiency** — increase performance, lower capacity requirements, improve agility, and dramatically lower costs

Scott D. Lowe is a former CIO, the founder and managing consultant of The 1610 Group, and cofounder of ActualTech Media. Scott has written thousands of articles for various publications and regularly contributes to VirtualizationSoftware.com and EnterpriseStorageGuide.com.



Open the book and find:

- Why virtualization on siloed infrastructure created unforeseen challenges
- How hyperconverged infrastructure goes above and beyond converged infrastructure
- How SDDC, ITaaS, and cloud are impacting your data center
- Ways that x86 commodity building blocks enable scale, agility, efficiency, and cost savings

Go to Dummies.com
for videos, step-by-step examples,
how-to articles, or to shop!



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.