

Making Everything Easier!™

HPE Special Edition

Hybrid IT

FOR
DUMMIES®
A Wiley Brand

Learn to:

- Discern the best location for individual workloads
- Identify how hybrid IT is the way of the future
- Build out your on-premises environment for maximum impact

Brought to you by



**Hewlett Packard
Enterprise**

Scott D. Lowe



About Hewlett Packard Enterprise

Hewlett Packard Enterprise is an industry-leading technology company that enables customers to go further, faster. With the industry's most comprehensive portfolio, spanning the cloud to the data center to workplace applications, our technology and services help customers around the world make IT more efficient, more productive, and more secure.

Hybrid IT
FOR
DUMMIES[®]
A Wiley Brand

HPE Special Edition

by Scott D. Lowe

FOR
DUMMIES[®]
A Wiley Brand

Hybrid IT For Dummies®, HPE Special Edition

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

Copyright © 2017 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. 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.

ISBN 978-1-119-44010-9 (pbk); ISBN 978-1-119-44011-6 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

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 BrandedRights&Licenses@Wiley.com.

Publisher's Acknowledgments

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

Development Editor: Elizabeth Kuball

Copy Editor: Elizabeth Kuball

Acquisitions Editor: Katie Mohr

Editorial Manager: Rev Mengle

Business Development Representative:

Karen Hattan

Production Editor: Vasanth Koilraj

Table of Contents

Introduction	1
About This Book.....	1
Foolish Assumptions	1
Icons Used in This Book.....	2
Beyond the Book.....	2
Chapter 1: Your Workloads, Your Options	3
Traditional IT Architectural Options.....	4
On-premises infrastructure.....	4
Managed service providers.....	5
Hosting and colocation.....	5
Private Cloud Features and Capabilities.....	6
The Public Cloud: Panacea or Not?	7
The real and perceived benefits of public cloud.....	8
Public cloud operating models explained.....	9
Bringing IT All Together: Hybrid IT	10
Chapter 2: Hybrid IT Critical Success Factors	13
Identifying the Problems with One-Size-Fits-All Infrastructure.....	13
Maintaining Performance Prerequisites.....	14
Containing Workload Runtime Environment Costs.....	15
On-premises	15
Public cloud.....	16
Corralling Infrastructure Control.....	16
Scaling to meet business demand	17
IT-as-a-Service	17
Safeguarding the Organization.....	17
Centralizing security control	18
Internet of Things/Device proliferation.....	18
Keeping IT legal: Ensuring regulatory compliance.....	19
Chapter 3: IT's All about the Workloads	21
Assessing Workload Infrastructure Needs	21
How application architecture drives cloud adoption decisions.....	21
Why the network is the public cloud's Achilles heel.....	22
Defining the right mix of infrastructure: Public, private, traditional	23
Understanding Workload Locality	23
Supporting legacy applications	24
Cloud-native application development.....	24
The role of bursting.....	25

Chapter 4: Workload Control	27
Setting the Scene: Why Workload Control Matters	27
Getting Clarity on Each Group's Role.....	28
Resolving Cloud Security Myths	28
Cloud is more secure than on-premises!.....	29
On-premises is more secure than cloud!.....	29
Understanding the Full Scope of On-Premises	
Physical Security.....	30
Denial-of-service attacks.....	30
Compromised firmware.....	31
Staff.....	31
Chapter 5: The Real Costs of Workloads	33
On-Premises Infrastructure Economic Models	33
Doing it yourself.....	34
Turning to managed service providers	35
Leasing	35
Cloud Services: Uncovering the Costs, Hidden and Not	36
Acquisition	36
Scaling	36
Data transfer and getting out of the cloud	37
Chapter 6: HPE Hybrid IT Infrastructure	39
Getting the Right Workload Environment Mix	39
An Array of Options	40
Infrastructure	40
Converged infrastructure	42
Hyperconverged infrastructure.....	42
Composable infrastructure	43
Bringing It all Together with HPE OneView	47
The Services Front: HPE Pointnext	48
Chapter 7: Enabling Hybrid IT	49
Solving the Staffing Situation	49
In-house IT	49
Consultants	50
Service providers.....	50
Understanding Hybrid IT Consumption Models	51
Getting the best of both worlds.....	51
Blending security.....	51
Gaining agility, flexibility, performance, scalability,	
and economic benefits	51
Maintaining Hybrid Control.....	52
Avoiding cloud lock-in	52
Enabling workload transparency.....	53
The perception of control	53
Chapter 8: Ten Key Hybrid IT Tips	55

Introduction

The world is changing! It seems like you can't load a news site without hearing that the cloud is going to overrun the on-premises data center. If you don't move your workloads *today*, you're a dinosaur and you're going to be left behind and replaced with someone more forward-thinking!

What?

This is where you start to think critically. You realize that your workloads are running just fine and, really, supporting workloads is one of the most critical IT functions. You may realize that, at times, the cloud can be a good option.

But, it's just that: an option. In this book, you see that there are two sides to the data center infrastructure equation and discover how each contributes to what will ultimately become the new normal in IT: hybrid IT.

About This Book

This book helps you understand both sides of the hybrid IT equation and how HPE can help your organization transform its IT operations and save time and money in the process. I delve into the worlds of security, economics, and operations to show you new ways to support your business workloads.

Foolish Assumptions

For this book, I assume that you have at least a basic understanding of data center computing, budgets, security, and operations. Although the general audience for this book is the IT decision maker, it's also useful for anyone in IT or the business who may want to learn more about hybrid IT.

Icons Used in This Book

Throughout this book, you find a number of icons intended to help you better understand and remember key concepts.



When you see the Remember icon, put that information in your back pocket to save for later.



When I share something that may save you time or money, I mark it with the Tip icon.



This book doesn't go super deep into technical stuff, but sometimes I wade into the weeds, and when I do, I use the Technical Stuff icon. You can safely skip anything marked with this icon without losing the main point.

Beyond the Book

I could only fit so much information in these 48 pages. You can learn much more about hybrid IT and HPE's support for it at www.hpe.com/us/en/solutions/transform-hybrid.html.

Chapter 1

Your Workloads, Your Options

In This Chapter

- ▶ Looking at the three most popular traditional IT deployment options
- ▶ Identifying the difference between virtualization and a private cloud
- ▶ Discovering the real versus perceived benefits of the cloud

Cloud. Does the very word make you cringe a little? Or does it pique your interest as you seek new and innovative ways to support the needs of your organization? Either reaction is understandable. For years, the term has been used, misused, and abused to a point where, for many people, it's tough to fully understand what it means and how it can really help propel an organization forward. As such, many organizations are still 100 percent on-premises, paralyzed by the mish-mash of guidance.

At the same time, though, all kinds of companies are screeching "To the cloud!" while advising other companies to simply dump their existing data centers and throw everything into some cloud provider's environment instead.



There is a happy, productive, cost-effective, and secure middle ground. This middle ground — dubbed *hybrid IT* — is quickly emerging as the new normal for enterprise IT departments. Hybrid IT combines the right mix of traditional IT, private cloud, and public cloud to meet your business and IT goals, so you can integrate new technologies where needed and maintain legacy systems where appropriate.

In this chapter, I offer a guided tour of the data center architecture options at your disposal.

But before I begin, I'll let you in on a little secret (okay, maybe it's not a secret!): You're probably already running workloads in a plethora of locations, including your on-premises data centers and various clouds. Most companies today have adopted a hybrid IT operating model, but they may still struggle with how best to define their environment.

Traditional IT Architectural Options

The data center has always been a somewhat complex place, with a plethora of options from which to choose how you want to operate workloads. The chosen path is generally selected after careful consideration of the costs and outcomes related to each approach.

On-premises infrastructure

You probably have an on-premises infrastructure of some kind in place today. These are your local data centers, server rooms, or server closets — whatever you happen to be operating. In a traditional environment, you and your staff decide what storage, servers, and networking gear you want to operate and then you go out and buy it, install it, and manage it on your own.



On-premises infrastructure is often unfairly maligned by cloud pushers. They say that traditional on-premises infrastructure is simply not flexible enough to meet the demands of modern businesses. They say that the need to buy everything upfront is simply not sustainable or desirable.

These people are correct . . . when an environment isn't well planned, well supported, well maintained, and well updated. A *good* on-premises environment can be just as cost-effective and reliable as any cloud provider, especially when supported by a highly trained staff. Plus, although many people discuss the rapid innovation in cloud, on-premises hasn't stood still and many new technologies are available to lower costs and

improve agility. Some examples include containers, hyperconverged, lowering cost of flash storage, and more. In addition, there are options available on the market that can assist with reducing upfront payments for hardware and software.

Managed service providers

People are the most expensive resource that you have in the budget. They can also be the most unreliable sometimes, and it seems like there are never enough of the right ones around when you need them. It can be exceedingly difficult to find people with the perfect combination of skills and experience to be able to meet all your business goals. That's where managed services come in.

Managed service providers are companies that can take some of the burden off you and your staff by augmenting your resources and, in some cases, take over complete management of certain services. As technology continues to become more complex and grows throughout an organization, a managed services contract can mean the difference between success and failure.

Hosting and colocation

Not every organization is gifted with plentiful physical space in which to build out a massive data center. Not every organization has the advantage of being located in a bandwidth-laden setting. For these and many other reasons, hosting and colocation have long been options for chief information officers (CIOs).

In a hosting or colocation scenario, you basically drive all or some of your servers, storage, and data center networking equipment to someone else's data center and you rack and stack everything there. The provider supplies you with all the physical space, power, cooling, and bandwidth you need to get your services up and running.



Colocation can provide you with a way to get services deployed without having to build out or expand an existing data center environment, which can be incredibly expensive. With such services, there are generally options available to

you for remote support and administration of your environment as well, making these kinds of services another potential augmentation for your existing staff.

Private Cloud Features and Capabilities

For options in which you decide to keep data center infrastructure privately owned, you can ultimately convert those services to *private cloud*. This doesn't mean that you wander into the data center with a magic wand, recite an enchantment, and magically create a new data center paradigm. In reality, it takes some planning and effort to get from a local data center to a private cloud.



But, what, exactly, is meant by the term *private cloud*? Let's start by telling you what it is *not*. Private cloud does not equate to highly virtualized. If you have a highly virtualized data center, that's all you have unless you've taken steps to add some additional capabilities, including the following:

- ✔ **Automation:** When you provision a service at a cloud provider, the process happens automatically. There's not some employee frantically banging away at a keyboard and mouse to fulfill your request along with the thousands of other requests that come in every hour. This automation enables cloud providers to focus far more on providing additional value to clients rather than just the basics.
- ✔ **Orchestration:** Closely related to automation, orchestration capabilities allow cloud providers to be able to link automated processes together so that even complex tasks can be handled without a human having to get involved.
- ✔ **Self-service capabilities:** Being able to automate and orchestrate processes means that users can be trusted to request their own resources. IT can place constraints around what is allowed to be provisioned and in what quantities. Thanks to the magic of automation and orchestration, when a user makes a request for resources, an IT staff person won't need to get involved, which enables

IT to focus on more important things while the basics are solved by users themselves.

- ✔ **Easy scalability:** One aspect of cloud is the ability to easily expand resource pools, but that requires underlying hardware that is easy to scale. At the same time, the individual workloads that users deploy should also be relatively easy to expand if the application calls for growth.
- ✔ **Chargeback/showback:** The public cloud makes it very clear how much you're spending, but many private environments don't yet have this capability. Even if you're not interested in charging individual departments for usage, it's still important for users to understand their actual usage and what it's costing the company. This also helps IT to better understand what's being used so they can make better decisions around where the infrastructure may need attention.

When you have these elements in place, you can officially consider yourself private cloudified!

The Public Cloud: Panacea or Not?

So many pixels and so much ink have been wasted on the idea that the public cloud is the end-all, be-all and only future of IT.

That's simply not the case. In fact, the future of IT is decidedly hybrid in nature. Some organizations will jump 100 percent into the public cloud; other organizations will stay 100 percent on-premises. Others will start in the cloud and go back to on-premises. But they'll be the outliers. The mainstream will be somewhere in the middle, with a mix of public cloud services and on-premises infrastructure.

Why aren't people just flocking to the cloud and never looking back? In many cases, it comes down to cold, hard cash. There is something of a cloud boomerang effect. At first, using cloud has some financial benefits. Over time, however, the operational costs associated with public cloud start to creep up and, eventually, you're staring a \$100,000 monthly invoice in the face.

Some people refer to this phenomenon as *cloud cliff*. The early party days of the public cloud resulted in some nasty hangovers, which have pushed people to rethink how they approach the service. They're not necessarily going all tee-totaler with regard to cloud, but they're warily adopting services in a far more targeted way, which is as it should be.



If someone tells you that a single solution is the only solution that is right 100 percent of the time, he's 100 percent misleading you.

The real and perceived benefits of public cloud

Now, don't put down the book thinking that I hate cloud. I don't! In fact, the public cloud carries with it a number of real and perceived benefits that are important to understand.

First, let's talk economics. For many cloud services, the initial capital expenditure that you need to fork out to get started is exactly \$0. You simply pay for what you use, a model that means that your operational budget is tapped rather than your capital budget. It also means that you're not paying for services that you don't directly use. For many chief financial officers (CFOs), this is a dream.

Instant scalability is another hallmark of the public cloud. If you need more of something — more virtual machines, more compute, more storage — you simply provision it and start paying for it. You don't need to wait for an IT staff person to order additional hardware and deploy it.

Sometimes, though, you don't need to scale *up*. Instead, you need to scale things *back* a bit as workload demands diminish. This is one area in which public clouds trounce on-premises infrastructure. Even in the most well-designed on-premises environment, when you've added hardware to the environment to handle new workload needs, you're probably not going to get rid of that hardware once that need has passed. The capability for the public cloud to support elastic resource consumption is a strength.

Public cloud operating models explained



There are a lot of different ways that organizations can consume public cloud resources. In fact, as you consider the totality of your personal and professional life, you're probably using at least one product or service from companies that operate in one of these spaces. There are a number of offerings from which you can choose:

- ✔ **Software as a Service (SaaS):** Under this simplest form of cloud computing, the provider controls everything and provides just an application layer interface for specific configuration items. As the customer, you don't need to worry about any underlying services except those that may extend the service.
- ✔ **Platform as a Service (PaaS):** Sometimes, you just need an infrastructure and an application development platform, and that's where PaaS comes in. Platforms typically include the capability to automate and deploy applications and include operating systems, databases, middleware, tools, and services. The customer is responsible for managing the application and data layers only.
- ✔ **Infrastructure as a Service (IaaS):** With an IaaS implementation, the vendor provides you with the underlying network, storage, compute resources, and virtualization technology. However, your IT staff will have the responsibility for configuring the resources. In addition, you have responsibility for managing security, databases, and applications.
- ✔ **Everything else (*aaS):** Everywhere you turn, you see all kinds of as-a-service products popping up. There are Virtual Desktops as a Service, Databases as a Service, and all kinds of services emerging. In most cases, these are just variations of one of the three items mentioned earlier (most often, repackaged SaaS offerings).



Let me give you a bit more of an overview of how these services are differentiated. As you can see in the On-Premises column in Table 1-1, when you have a traditional on-premises environment, you own and manage it all, from the network to the application.

Table 1-1 Cloud Offerings Responsibility Comparison

<i>Responsibility</i>	<i>On-Premises</i>	<i>IaaS</i>	<i>PaaS</i>	<i>SaaS</i>
Application	You	You	You	Provider
Data	You	You	You	Provider
Runtime	You	You	Provider	Provider
Middleware	You	You	Provider	Provider
OS	You	You	Provider	Provider
Hypervisor	You	Provider	Provider	Provider
Server	You	Provider	Provider	Provider
Storage	You	Provider	Provider	Provider
Networking	You	Provider	Provider	Provider

As you start to adopt as-a-service offerings, the provider you choose begins to manage certain aspects of the environment for you. For example, IaaS offerings may provide you with server, storage, networking, and a hypervisor upon which your staff handles just the operating system and up. At the extreme side of the equation are SaaS offerings, which are fully deployed and managed by the service provider. You simply consume the application. Table 1-1 gives you a rough overview of who holds primary responsibility for each aspect of an environment in various cloud scenarios.

Bringing IT All Together: Hybrid IT

In the previous sections, I tell you about both on-premises and cloud offerings. I explain that on-premises infrastructure, when designed correctly, can be very cost effective. I also explain that cloud carries with it a great number of operational benefits.

So, which should you choose?

The answer is that you should choose *both* and then leverage both sides of the equation when it makes sense. You might run particularly sensitive and latency-sensitive applications

on-premises and then choose to move certain workloads into the cloud. For example, Office 365 is increasingly selected to replace on-premises Exchange environments because it's simpler to deploy and far easier to administer than Exchange.



That's just one example. For every workload in your portfolio, you need to develop a routine process that analyzes the need for that application and that helps you to decide where to run that application. Throughout the rest of this book, I offer tips and insight into how to develop a hybrid IT model that blends public cloud and on-premises infrastructure to ensure that you're continuing to meet the needs of your organization in the best way possible.

Chapter 2

Hybrid IT Critical Success Factors

In This Chapter

- ▶ Working around the “one-size-fits-all” factor
 - ▶ Thinking outside the box
 - ▶ Keeping your organization safe
-

You understand on-premises in the cloud, but your CEO may still believe that you should put *everything* in the cloud, or that *nothing* should go into the dark, scary cloud. You need to be armed with responses to just about any question you’re asked about your strategy.

In this chapter, I provide you with the arguments you need so you can push back and create a coherent strategy that makes sense for your company rather than one that is based on the tweet of the week.

Identifying the Problems with One-Size-Fits-All Infrastructure

You probably know that “one size fits all” is often synonymous with “one size fits none.” Every organization on the planet is different. Likewise, the implementation of every workload is different, even for specific applications that are installed in a lot of different companies. The one thing that all workloads have in common is that they need an environment in which to operate.

The infrastructure on which these workloads operate needs to reflect this uniqueness. Let's take a quick walk through a couple of key points.

There are all kinds of generic, one-size-fits-all data center options out there intended to meet the needs of everyone. But in reality, they often have some major drawbacks, including being improperly sized to meet the needs of the company.

You may wonder how one size fits all can lead to improper sizing. Well, the answer emerges when you take off the lid and peer inside. Some data center resources are exhausted before others. For example, you may run out of storage capacity long before you run out of computing power. You need to tailor the resource allocation to meet the needs of your specific workload mix. Just as important, you need to make sure that this allocation supports those workloads with whatever uniqueness they may have.



In an ideal world, you'd be able to deploy and maintain resources in the way that makes the most sense for your workloads *and* the individual way that you're using those workloads.

Here's the short version: One-size-fits-all infrastructure may be good for one-size-fits-all applications, but it breaks down as soon as customizations and unique needs hit the scene.

Maintaining Performance Prerequisites

One of the challenges in maintaining a robust data center environment is maintaining necessary performance prerequisites while everything changes around you. You're constantly adding and removing applications, expanding storage, and responding to new business needs. Every time you add something new to the environment, there is potential that the whole thing just slows down.

Beyond just maintaining performance requirements necessary to operate workloads, users have become accustomed to

consistent performance. They'll notice and complain if applications suddenly start running slower than they used to, even if the application itself is still working perfectly well.

Containing Workload Runtime Environment Costs

No one wants to spend more than necessary on IT. That's one of the reasons that the cloud is so tempting to so many chief financial officers (CFOs), but there are a lot of other considerations to take into account. Both on-premises and cloud-based infrastructure have their benefits and drawbacks when it comes to cost.

On-premises

For on-premises infrastructure, there is generally a desire to maximize the impact of capital budget spend, which only comes around at replacement cycle time. To ensure that the infrastructure can meet the needs of the business for years to come, there is a tendency to overbuy. Simply put, we buy too much because we need to get through the next three to five years without more capital spend.

But that's nothing compared to out-of-cycle need to spend. The business doesn't plan the next quarter's needs for three to five years out, but IT has to plan infrastructure that way, and we're often forced to perform upgrades off-cycle, which can be expensive and difficult to manage.



For ongoing on-premises infrastructure deployments, organizations simply have to find ways to become more “cloudlike” in their approach. There are plenty of infrastructure offerings that can help with this, such as hyperconverged and composable solutions.

But don't think that the way you spend and maintain on-premises infrastructure is bad. Here's the great thing about it: You can absolutely abuse it and push it to its limits, and doing so doesn't cost you a penny more.

And, it gets better. You can have your infrastructure and cloud economics, too. There are pay-as-you-go financing opportunities from companies such as HPE that allow you to maintain an on-premises computing model while paying for that infrastructure just as you would the public cloud. Under HPE's Flexible Capacity program, HPE owns, and supports, your choice of on-premises equipment. You just pay for what you use based on metered usage and access self-service reporting to forecast future demand. An on-site buffer allows you to quickly scale up capacity when needed, and avoid spending on unused capacity.

Public cloud

To address capacity needs and to end the need to spend a ton of money upfront, the cloud is perfect. You simply buy what you need and no more. As you know, however, simply chucking workloads over the cloud wall isn't a viable strategy because, eventually, you'll find a workload that wrecks your budget.

Public cloud carries with it charges in all sorts of locations, from compute time to storage capacity used, but also in terms of network egress. The more data you pump *out* of the provider's environment, the bigger your bill. Workload uncertainty can wreak havoc on your cloud bill, as can unconstrained usage from business units that may have unfettered access.

Worse, as your costs go up, you may not have real visibility into why they're increasing. Once you've migrated a workload, it can become a full-time job trying to figure out how much you're paying and whether that payment is really accurate.

Corralling Infrastructure Control

In recent years, there has been a movement toward empowering users and business units with regard to technology. As people have become savvier with regard to technology, it's been an appropriate shift that can be helpful to the organization, but that also has its challenges.

Even back in the days when IT was the final arbiter in all things related to technology, although users hated it, it was a “safe” way to manage the environment. There were just a few key people who had control. Today there are new challenges to consider.

Scaling to meet business demand

As mentioned before, today’s IT department doesn’t just stand still while the business shifts around it. Even immediately after we’ve deployed the latest technology as a part of our most current refresh cycle, the business has new demands that we may not have known about at the time.



Every IT environment needs the ability to quickly and easily scale resources to meet new demand.

IT-as-a-Service

Savvy IT companies have not really decentralized by throwing the keys to the kingdom to users. Instead they have created new IT-as-a-Service models. IT actually manages the environments operationally across on-premises or off, based on the business needs and the business benefits. Most businesses don’t want to own the operational or IT part, but they want choice in IT as well as flexibility and agility.

Safeguarding the Organization

You may read what I’ve written about allowing users to manage the infrastructure and think something along the lines of “over my dead body.” And you’d be right to do so. When there is discussion around giving users the keys to the infrastructure, we’re talking about the *valet key*, not the driver’s key. As you may know, when you hand a valet key over to someone parking your car, he’s able to drive the vehicle, but he can’t access the trunk or the glove compartment.



Although IT departments are being pressed to provide more control to end users, that control only extends so far. There is no expectation that the finance department will suddenly start creating logical unit numbers (LUNs) or that sales will begin to create administrative user accounts. But these groups should be able to access resources in a sufficient enough way to be able to meet their goals.

Centralizing security control

Perhaps one of the biggest areas that should remain strictly within the confines of IT is security control. Given the prevalence of new and dangerous attacks on infrastructure, desktops, and everything else these days, security is too important to ignore or distribute to users.

By maintaining a centralized security posture, organizations can ensure that there is consistency in policy and application of policy. Plus, the services that are allowed to be performed by end users themselves can be restricted behind the scenes to those that are safe to perform and that won't impact other areas of the environment.

Internet of Things/ Device proliferation

Depending on who you talk to, by 2020, it is expected that the world will see some 20 to 30 billion (yes, billion with a *b*) devices connected to the Internet, *excluding* smartphones, tablets, and computers. The age of the Internet of Things (IoT) is truly upon us, and it's creating new challenges for IT departments as they struggle to figure out how to support the resulting storage and data analysis needs. As was the case when PCs began to supplant mainframes and proliferate throughout businesses, IT departments will have to make fundamental shifts in strategy and infrastructure to cope with such changes. This may include deploying more easily expanded storage systems or leveraging cloud resources to maintain these devices.

Keeping IT legal: Ensuring regulatory compliance

One of IT's most important jobs revolves around risk management, which involves the aforementioned security concerns, but also includes helping make sure the organization remains in compliance with regulatory requirements. Sometimes, regulatory requirements either preclude cloud or make the choice of a cloud provider more difficult.

In any case, though, the infrastructure environment has to be designed in such a way that, even with user self-service in place, people can't accidentally expose the organization.



Figuring all this out alone can be tough. Organizations need a partner that can help them find the right mix of on-premises and cloud resources and assist in securing those resources. You need a trusted advisor who can help you figure out hybrid IT.

Chapter 3

IT's All about the Workloads

.....

In This Chapter

- ▶ Paying attention to workloads
 - ▶ Keeping legacy workloads running
 - ▶ Saving time and money with cloud bursting
-

IT isn't about storage. It isn't about networking. It isn't about servers. IT's primary responsibility is ensuring that applications and the data are sufficiently supported to meet the burgeoning needs of the business. In other words, IT's all about the workloads.

Assessing Workload Infrastructure Needs

Every workload is different. Each one requires a carefully crafted environment in order to operate optimally. As you consider a hybrid IT operating mode, you need to understand some really important issues, and that's what I cover in this section.

How application architecture drives cloud adoption decisions

The ideal operating environment for your application is largely dictated by how that application is built and supported. Traditional client/server applications, in which there

is a centralized database server with the client side installed on client PCs, are not generally great options to move to the cloud, although there are exceptions.



Take Microsoft Exchange Server, for instance. In these environments, the server holds the database and Microsoft Outlook is the client, but these environments move really nicely into the cloud in the form of Office 365, hosted Exchange, and other services. However, in this case, latency (to a point) is okay. Users aren't sitting at their desktops manically mashing the Send/Receive button because an email is taking too long.

With other client/server applications, though, imposing latency on the application directly impacts the user's ability to get work done.

In addition, client/server applications often require full operating systems for deployment, whereas more modern applications often operate using a web browser that is available on just about any device.

Of course, browser access to an enterprise resource planning (ERP) application from anywhere may not be as important to you as being able to access email from any device, but that's just one of the determinations you'll need to make on an application-by-application basis.

Why the network is the public cloud's Achilles heel

Partially for the reasons discussed in the previous section, but also for other reasons, the network can be a major concern in the world of cloud. It's one reason that application workloads are often left in the local data center, close to the users — or at the edge, close to where it's created. After all, the closer applications and data are, the faster they can respond.



It's all about physics. Distance results in latency. The more distance between a user and an application server, the more latency there is. The more latency, the more the workload suffers and the circle is complete, and not in a positive way. Latency is *the* killer of workloads. And, in the cloud, this results in additional cost.

In addition to lamenting about latency, as you throw more networks between the user and an application, the likelihood of a network outage increases, which would bring down that service.

Performance and reliability are the lynchpins to successful IT. If you have a highly centralized workforce in a single building or a single campus, keeping workloads close by might make more sense. If, however, you have a highly distributed workforce and your users are mostly remote or you're a web application company, centralizing some key applications in the cloud could make sense, but you have to ensure that you don't damage performance and reliability by doing so.

Defining the right mix of infrastructure: Public, private, traditional

The right mix of infrastructure is different for every company, but it will include some combination of public cloud and on-premises infrastructure. In addition to application architecture and latency and network concerns, you also need to consider such factors as ease of support and cost to run a workload. Using Office 365 as an example, you may find it less expensive and less frustrating to move to Office 365 than to retain Exchange locally. If you've done that or you're considering such a move, you know that you went through a litany of checklist items before you made that fateful decision.

And that was for just a single application, albeit a big one. Now, as you consider your future hybrid IT strategy, you need to repeat that process for each and every application in your portfolio. It may sound daunting, but it's a worthwhile exercise.

Understanding Workload Locality

Workload and data locality are among the biggest decisions you have to make when you're trying to decide which applications should live locally and which ones should live in the cloud.

Supporting legacy applications

We chatted earlier about the latency reasons you may choose on-premises over cloud, but there are other considerations, too. If you have an application that is particularly sensitive, you may not want it to live beyond the confines of your data center.



Furthermore, there are geopolitical considerations and country-specific data governance laws to take into account. As you move beyond the borders of the United States, there is mistrust of U.S.-based clouds and U.S.-based services due to the potential for U.S. government spying, a scenario not likely to improve anytime soon.

Assuming that these aren't issues, there are ways to get around the client/server application latency issue, but it involves creating a series of virtual desktops at the cloud provider, installing the client on a virtual desktop, and then allowing a user to simply interact with that virtual desktop. As you might imagine, having to jump through such hoops can get expensive and complicated, so you have to take into account these processes during your application analysis.

Even in the realm of legacy apps, if you have those with large data requirements, there are cloud providers that have other cloud services and partners that can help you sift through it, sometimes making cloud a better fit for these needs.

Cloud-native application development

As we continue our unyielding trek into the future, we also change the way we do business. These are the kinds of applications that lend themselves very well to working in the cloud because the browser is the client. Plus, with new application development methodologies such as those based on microservices, cloud is a great target because these applications launch a lot of small short-lived services that are generally torn down very quickly. With cloud-native application development, you can allow your applications to scale to become very large and very distributed.

The role of bursting

Many industries are seasonal. There are seasons of peak computing need and, with traditional infrastructure, you need to build to that peak in order to satisfy application demands. For some industries, this may not be practical, so turning to the cloud to handle capacity needs beyond what's possible in an on-premises data center is often the answer.



This ability to burst to the cloud requires well-considered integration between your on-premises environment and whichever cloud service providers you select. You need to enable connectivity between these environments, as well as from any clients that will connect to these services.

Chapter 4

Workload Control

In This Chapter

- ▶ Discovering the workload control concerns that drive hybrid IT priorities
- ▶ Comparing the security of your local data center and the cloud
- ▶ Understanding why people can be the weakest link in your chain

In a world that is more interconnected each and every day, and as companies dip ever more deeply into the technology well, workload control is becoming a more important issue. Who is going to manage each individual aspect of the workload's life cycle? When it comes to working with on-premises and cloud-based environments, there are some things you need to remember.



As a part of workload control, remaining secure and compliant with regulatory agencies and other rules is generally considered a good outcome, especially if you want to avoid bad PR, huge fines, and jail sentences. Making sure you pay attention to these rules is always a good idea.

Setting the Scene: Why Workload Control Matters

CIOs and business decision makers *really* hate hearing, “I thought that was someone else’s job” when it comes to figuring out why something went wrong in a business-critical workload. As you move workloads around between on-premises and cloud environments, lines of support and authority can become blurry, leading to a lack of clear understanding of roles.

Plus, as workloads, which are prime candidates for the cloud, continue to rise, there may be a sense of loss of control from internal IT as a service provider assumes responsibility for certain functions. It can be a frustrating state, particularly when cloud-based workloads become unavailable.

In these scenarios, IT is largely on the sidelines as they await return to service from the provider. It's an uncomfortable and unfamiliar place for IT professionals, who, with on-premises infrastructure failure, can jump right into action themselves.

Getting Clarity on Each Group's Role

As you work through the appropriate location for each of your individual workloads, you'll also need to determine who is going to handle individual aspects of that workload. For the on-premises side of your hybrid IT environment, this is pretty much business as usual. Your internal IT team will likely handle most things, including ongoing management, software updates, data protection, and disaster recovery.

However, for your cloud-based workloads, you'll need to spend a bit more time figuring all this out. You may decide that your in-house IT staff are going to handle general application administration for a particular workload, but you're going to ask the provider to handle the data protection and disaster recovery responsibilities. With hybrid IT, you gain far more opportunity to assign roles in a granular fashion, enabling you to play to strengths instead of having to cover for weakness when you don't have sufficient internal skills.



Getting the role clarity part nailed down is really important. Any ambiguity here can result in extended downtime as internal IT and provider staff attempt to clarify who is supposed to be handling things.

Resolving Cloud Security Myths

We can't talk about workload control without discussing enduring cloud security myths. There are two diametrically opposed schools of thought when it comes to cloud security.

One exclaims that moving to cloud will solve all your security woes. The other side decries the state of cloud security, implying that you'd be crazy to move into the cloud. Let's debunk *both* sides of the security myth.

Cloud is more secure than on-premises!

For years, cloud providers and supporters have used security as an argument for moving to the cloud. After all, these folks say, cloud providers have a financial incentive and the resources to hire only the best and the brightest security minds in the world, and their services, as a result, are rock solid. It's the rare organization that will have more internal security expertise than cloud service providers.

And they're right. Cloud providers do tend to have a very strong security posture. They really have to. One of the reasons is that cloud providers are under constant attack. With all that juicy data, they're really attractive targets. And, as shared services, cloud providers have to protect more heavily from external threats, as well as those that originate inside their networks.

On-premises is more secure than cloud!

For those that are on-premises, part of the reason is often security. The feeling is that the company can do better on its own than trusting some cloud provider that doesn't have a vested interest in the company. The problem is that, although cloud companies may not necessarily have direct vested interest in your security, they do care about their reputation and what happens to their customers overall.

The other point to consider is that commodity public cloud does not necessarily disclose all security layers, and for some industries there are very specific security protocols that must be supported. Often, it's easier and cheaper for these specific workloads to reside in an environment in which the customer maintains an end-to-end security solution.



For those who operate in highly regulated arenas, such as the finance and healthcare industries, although there are hardened cloud services that you may be able to use, the fact is that your local data center will probably have more security and be more compliant than a commodity cloud provider. Commodity cloud providers look at the overall customer base and may not delve deeper to consider individual vertical needs.

Just like all other aspects of the hybrid cloud, the location you choose for individual workloads will depend on a number of factors, security included.

Understanding the Full Scope of On-Premises Physical Security

Let's talk about on-premises security for a minute. Like cloud providers, you may be under regular attack and just not know it. Or you may be painfully aware, as your Internet services slow to a crawl.

There are also a multitude of other risks that you need to keep in mind when it comes to on-premises physical security.

Denial-of-service attacks

Denial-of-service (DOS) attacks happen more frequently and are more serious than ever before. Some of this is due to hackers stepping up their game, but some is due to the fact that there are many, many more devices on the Internet that can be turned into so-called *zombies* and programmed to attack.

A DOS attack against your headquarters could be devastating, unless you have multiple redundant connections to the Internet — and even that may not help you. Cloud providers typically have multiple points of connectivity to the Internet. So, even though they're impacted by DOS attacks, the impact may be more muted. Plus, some terms of service with cloud providers may actually allow them to shut down your virtual machines if your site experiences a DOS attack. When a DOS attack does take place, although it may not impact every

customer on the service, the client under attack may actually face a steep increase in their monthly bandwidth bill.



There are all kinds of services available today, such as Cloudflare, that can help you protect yourself from DOS attacks, whether you're in the cloud or on-premises.

Compromised firmware

Firmware is critical to hardware functionality and gets updates every so often to correct bugs and to add new features. The unfortunate reality is this: It's not always that difficult to install compromised firmware updates to a server. Compromised firmware can intercept network communications, decode the contents of storage, and a whole lot more. There have been reports of compromised servers being shipped to unsuspecting customers. That's not good.

Firmware compromises generally require physical access to the server. This is just one reason that allowing only authorized people into a data center is so important. It's also important to log data center access so security events can be correlated against visitor logs.

Staff

The weakest link in any security system is the human element. People can make mistakes. People can be bribed. And without comment on whether this is right or wrong, people can unilaterally decide that their organizations are on the wrong side of history and release a treasure-trove of information that has wide-ranging impact.



When you're considering your information and security strategy, you need to keep *trust* in mind and ensure that there is always a second set of eyes reviewing decisions that are made, systems that are put in, and processes that are developed.

Chapter 5

The Real Costs of Workloads

In This Chapter

- ▶ Paying for your data center infrastructure
 - ▶ Considering the pitfalls of the cloud
-

Companies today are looking for balance in IT. They have dozens, hundreds, and even thousands of workloads already deployed, and they want to run these workloads in the most financially efficient way possible. Don't read that to mean that all companies want to do IT on the cheap. Instead, they want to ensure that workload operating environments are designed to support the current and ongoing needs of the business. At the same time, of course, they want to run those workloads in a way that makes financial sense.

On-Premises Infrastructure Economic Models

You probably already have a data center. It has servers, storage, networking, and workloads. You have people who manage it all. And you have a business that is constantly seeking more from you, but that isn't necessarily providing more financial or staffing resources to meet expanded goals.

Doing it yourself

Traditionally, companies have taken a do-it-yourself (DIY), pay-as-you-go approach to the data center. Through separate capital and operational budgets, hardware and software is procured, placed into service and depreciated over a period of, generally, three to five years (although this time frame is sometimes adjusted if needs are different).

The DIY method is sometimes considered the most resource-intensive option out there because, as the name suggests, you're doing everything yourself, and you have to hire people with appropriate skills across a variety of disciplines. Although many cloud advocates may deride DIY as a legacy approach, nothing could be further from the truth. Aspects of the approach can use some improvement, but other aspects work well. DIY success also depends on the tools at hand, and the tools for DIY'ers are getting much, much better. In fact, a lot of what enables cloud providers to do amazing things is open sourced, and that know-how is coming down to the on-premises DIY crowd.

On the plus side, because DIY is most often associated with on-premises or co-located data centers, there is also a recognition that it's easier to deal with many of the application performance issues that were discussed in previous chapters. After all, you're housing the application close to the user, so latency and the like shouldn't be a problem. Plus, when problems arise, time to respond is much faster in an on-premises environment.

On the downside, a DIY approach often implies that a massive capital purchase was made. However, as previously mentioned, this isn't always true anymore, either, as companies continue to expand their pay-as-you-go offerings.

On the staffing front, DIY also implies that you're hiring people to do all the work and that you're hiring a variety of skill sets. This is, perhaps, one of the biggest challenges in IT. And although enterprise IT vendors such as HPE are continually simplifying their products, you still need people to manage it all.

Turning to managed service providers

To help counter some of the people problems associated with deploying and managing data center technology, you may turn to managed services providers. These entities take care of specific technology areas for you, freeing up your most valuable resource — time.



The provider handles the installation and management of selected workloads and services. In this context, the firm is on the hook for making sure everything is done correctly. You may still be the person in charge, but the day-to-day hassle is someone else's problem.

Leasing

Remember that pesky business of having to pay for things you want in your data center? You may be able to get out of writing a huge check upfront by using leasing services. Leasing services allow you to rent your infrastructure rather than pay for it all upfront. In this way, it's a bit closer to how you pay for cloud services, although it's still not consumption based. It does, however, enable some financial flexibility, which may allow you to rethink how you implement data center services.

There are two primary ways to lease equipment:

- **Operating lease:** You truly *rent* equipment. You decide what to buy, and the leasing company buys it and rents it back to you for a monthly or annual payment. Of course, there is a *lease factor rate*, which is basically interest, included in the payment, but if you stick to a strict replacement schedule, you may find that the total of the payments is still less than what you would have paid had you purchased outright. With operating leases, payments are truly operating expenses, and the equipment is never placed on your company's books.
- **Capital lease:** This is closer to a loan, but you still make rental payments for the term of the lease. At the end of the lease period, you may have the option to buy the equipment at a then-current fair market value or for a \$1 buyout, at which point the ownership of the assets will transfer to you.



There is a lot more to leasing than I can fit in this chapter, but the short discussion point is this: Leasing can help you operationalize your data center expenses and get closer to pay as you go and may enable you to, from a financial perspective, more easily react to new business needs as they impact your on-premises data center infrastructure.

Cloud Services: Uncovering the Costs, Hidden and Not

Like on-premises infrastructure, cloud has economic models, but rather than a series of options, it's generally a series of charges that you incur at different points.

Acquisition

Cloud service acquisition costs are often pretty low, at least compared to buying a data center outright. You pay for what you need at the moment, and that's it. For some services, you may pay an activation charge of some kind, and you may choose to pay consultants or cloud provider staff to help you get started, but, again, on the hardware and software licensing front, you pay for just what you use.

Scaling

Eventually, you'll probably need to add storage or compute or some other service. It's easy in the cloud. You simply request more resources, and they magically appear! Of course, your payment goes up as well.



This is one of the big dangers in the cloud. There are stories of companies unknowingly racking up thousands and thousands of dollars in charges because someone made a mistake. Sometimes the provider will waive these charges, and sometimes they won't.

Data transfer and getting out of the cloud

Here's where things get dicey. Most cloud service providers don't charge you to bring data *in* to their environment from the Internet. But, when you start to move data *out*, things look a bit different.

Take a look at Table 5-1. Here, you see current outbound data transfer pricing from AWS. For this example, let's assume that, on average, you move 500TB per month of data out of the provider to a third-party data protection service. In this case, you'll incur almost \$29,000 per month in charges to perform this operation.

Table 5-1 A Sample Monthly Data Transfer Charge

<i>Outbound Data Transfer</i>	<i>U.S. Zone (\$/GB)</i>	<i>Move 500TB Out</i>
First 1 GB / month	\$0.000	\$0.00
Up to 10 TB / month	\$0.090	\$899.91
Next 40 TB / month	\$0.085	\$3,400.00
Next 100 TB / month	\$0.070	\$7,000.00
Next 350 TB / month	\$0.050	\$17,500.00
Total Per Month		\$28,799.91

Because of this and other hidden costs, more and more companies are pulling back from the cloud and returning some workloads to their on-premises data centers. Although the public cloud can provide a great deal of flexibility, that flexibility often comes with a very steep price tag that you may not fully understand until your eyes pop out of your head on receiving your monthly cloud bill.

Chapter 6

HPE Hybrid IT Infrastructure

In This Chapter

- ▶ Finding the right mix of on-premises and public cloud
- ▶ Ensuring on-premises success with HPE
- ▶ Finding the perfect hardware and software combination

As mentioned in previous chapters, when you have easily supported workloads that don't have or need much customization and that don't suffer from latency issues between the client and the server, you should consider the public cloud. For more guidance on public cloud services available to you, visit www.cloud28plus.com/na. For everything else, there's on-premises infrastructure. **Remember:** On-premises infrastructure encompasses everything in your local and remote data centers, as well as any co-located data centers you may lease.

Getting the Right Workload Environment Mix

Today's enterprises are balancing systems of record with the digitization of everything; they must operate in both the cloud and on-premises worlds. The traditional "systems of record" world co-exists with the "all things digital" world. Hybrid IT is the new reality; there is no one-size-fits-all IT strategy for today's workloads. Enterprises are operating IT environments based on a mix of traditional/on-premises environments, along with private and public clouds, and this will continue. It's complex. Now more than ever, IT must accelerate the business with the speed and flexibility to pivot with business demands across environments. The key to supporting these worlds requires the ability to compose resources for every

workload, as needed. To achieve this, enterprises must work to find the right mix of hybrid IT across traditional, private and public cloud environments. It's different for every enterprise, every industry, and it changes over time.

For enterprises to be successful, HPE believes speeding time to value across a hybrid world is a key priority. To do this, the new generation of workloads demands a technology platform strategy that is

✔ **Fast, flexible, scalable, and “composable,”** delivered in today's reality of a hybrid IT environment — whether your apps and data are in your data center, private and public clouds, multi-clouds and the edge of your network — and having the flexibility to operate seamlessly across these environments.

✔ **Able to capitalize on the convergence of IoT, apps, data closest to your customers, employees, users, and transactions** — the intelligent edge — where your enterprise touches customer, employees, and the outside world, and where data, new value, and experiences are being created. Beyond reliable and secure connections, it's critical the data collection and analytics be built in to provide context and insights that drive value.

An Array of Options

For on-premises infrastructure, HPE makes available a vast array of platforms and partners, which, when combined, can provide you with an on-premises workload environment that is affordable, easily maintained, and highly integrated, and that provides full support for even the most demanding applications.

Infrastructure

Although today's hyperconverged and composable stacks leverage best-in-class building blocks and are optimized for faster service delivery, they all start with the right core elements of compute, storage, and networking. HPE's software-defined infrastructure strategy delivers a common

management strategy that simplifies Hybrid IT, no matter the infrastructure you choose.

You can choose to modernize and scale data center elements independently or consume them as converged offerings, but either way, taking advantage of new compute innovations, storage data services, and low-latency networks is a critical part of a hybrid IT foundation.

HPE ProLiant

According to IDC, HPE ProLiant servers are the best-selling servers in the market. They remain the cornerstone of the data center market and are a rock-solid choice for those wanting to pursue a traditional infrastructure strategy.

HPE Storage

HPE Storage — including HPE 3PAR and HPE Nimble — is also a bedrock choice for the data center and forms the storage foundation for many organizations. Today's storage market is growing more complex all the time, with new media options and new data services opportunities. HPE Storage leads the way, with all-flash and hybrid flash arrays, containing capacity costs with comprehensive deduplication and compression features, and with HPE Recovery Manager Central and HPE StoreOnce — ensuring business continuity with reliable and cost-effective data protection.



These kinds of features are must-haves in the modern on-premises data center. Data growth is escalating, and these capabilities go a long way toward helping to keep spiraling capacity needs from massively impacting the IT budget.

HPE Networking

The network is the communications glue that holds the data center together. Binding together the servers, the storage, and the users and application with the Internet, the cloud, and other services, HPE's networking line provides organizations with a comprehensive set of capabilities that will help you meet your business goals.

Converged infrastructure

Traditional infrastructure can be the right choice, but for many organizations, it has some serious drawbacks. Such drawbacks generally revolve around the time it takes to get a complete infrastructure operational, particularly when you buy things from different vendors and cobble them all together. There is the potential for interoperability issues that can take a lot of time to resolve, which means that the time-to-value for the infrastructure purchase might be extended. *Time-to-value* is a measure of how long it takes from purchase order to operation. Most companies don't want a long figure here and want to be able to begin to gain benefit from their investments in an expedited manner.



This is where converged infrastructure, such as HPE ConvergedSystem, comes in. Converged infrastructure combines storage, servers, and networking into a single purchase from a single vendor so that you know everything will work exactly as it should. Converged infrastructure is often a simple bundling of existing products, but with heavy testing. Further, many converged systems arrive in your data center fully racked and cabled. You simply plug them in to your network and turn on the power.

Such solutions can massively reduce the time-to-value, and you can start running your business workloads almost immediately.

Hyperconverged infrastructure

But, maybe you don't want a converged solution or you want to start a bit smaller and be able to grow in more granular steps as your workload needs change. Or, perhaps you're looking for a solution that doesn't just transform the buying process, but that can also transform how you operate the infrastructure. Enter hyperconverged infrastructure.

Hyperconverged infrastructure is a data center model in which a simple, efficient, software-defined platform combines everything below the hypervisor — compute, storage, and data services — into a single appliance. The software on each appliance communicates with the software on all the other appliances in the hyperconverged cluster and creates

a software-defined single, shared pool of resources. Today's businesses need a hyperconverged solution that combines not only servers and storage, but all other IT devices below the hypervisor. That's exactly what the HPE SimpliVity 380 provides.

HPE SimpliVity 380 is an enterprise-grade solution that dramatically simplifies IT by combining all infrastructure and advanced data services for virtualized workloads — including guaranteed data efficiency, data protection, and VM-centric management and mobility — into one powerful hyperconverged platform.

At the core of HPE SimpliVity 380 is the HPE SimpliVity Data Virtualization Platform, which guarantees 90 percent capacity savings across VM storage and backup through de-duplication, compression, and optimization. The solution provides the highest levels of data integrity and availability ensured by built-in resilience, backup, and replication and includes simple, intuitive global VM-centric management and mobility that simplifies day-to-day operations and enables seamless data mobility.

The HPE SimpliVity 380 dramatically improve overall efficiency with TCO savings of 73 percent compared to traditional infrastructure and 49 percent compared to public cloud. HPE SimpliVity 380 is based on the highly configurable HPE ProLiant DL380 Gen9 Server. Combined with VMware vSphere, for increased storage, network, compute, and graphics options, this platform significantly improves flexibility in deployment and enables you to adjust storage, compute, and network capacity as your business needs change. As you grow, simply order identically configured nodes, and know your expansion will occur automatically at power up.

Composable infrastructure

At the top of the data center architecture pyramid of today lies composable infrastructure. Composable infrastructure carries with it a number of characteristics and has a number of points of contrast with the other options discussed in this chapter, which you can see in Table 6-1.

Table 6-1 Data Center Architecture Spectrum

	<i>Traditional</i>	<i>Converged</i>	<i>Hyperconverged</i>	<i>Composable</i>
Complexity	Low to high, depending on implementation	Moderate	Moderate	Low
Time-to-value	Weeks, months	Days	Hours, minutes	Minutes, seconds
Flexibility	Moderate	Moderate	Moderate	High
Scaling	Moderate to difficult	Moderate	Simple	Simple
Workloads supported	Physical, virtual, containers	Physical, virtual, containers	Virtual	Physical, virtual, containers

Here are four primary characteristics of composable infrastructure:

- ✔ **Hardware and software architected as one:** Unlike many other data center architectures, composable infrastructure systems are built right alongside the software that will manage the environment. This tight coupling ensures that the two sides of the house cooperate to make sure that business workloads are well supported.
- ✔ **Fluid IT:** With the right backing hardware, composable infrastructure makes it a breeze to manage resources, including compute, storage, and fabric.
- ✔ **Software-defined intelligence:** In general, hardware-bound architectures do not enjoy the kinds of flexibility and agility that are so critical to today's businesses. Composable infrastructure, however, puts all the enterprise brains into software, making it very easy for the system to manage the abstracted resources and ensure that they're used to the best possible effect.
- ✔ **Physical, virtual, and containerized workloads:** With composable infrastructure, you don't need to compromise on the kinds of workloads you need to run. Bring it on!

HPE Synergy

As a next-generation data center architectural option, composable infrastructure embraces and extends key concepts and traits from the architectures that have come before it, including converged and hyperconverged systems.

Simply trying to “bolt” composable onto one of these less-inclusive architectures wouldn’t have resulted in a complete, tightly integrated solution. With that in mind, HPE developed the first platform architected for composability focused solely on helping you achieve critical outcomes. It’s called HPE Synergy.

HPE Synergy is a custom-developed set of hardware that includes storage, servers, frames (chassis), and internetworking necessary to meet the key design principles associated with composable infrastructure.

HPE Synergy Image Streamer

Deploying a typical operating system or hypervisor can be time consuming because it requires customizing and/or copying each image for each compute module. HPE Synergy Image Streamer accelerates these processes through tight integration with HPE Synergy Composer templates.



The image payload consists of your compute module’s profile (from HPE Synergy Composer), your golden image (the operating environment with a bootable operating system and application and the I/O driver version), and the personality (the operating system and application configuration). Through the template, bootable images are deployed across stateless compute modules using a simple, consistent process. If you’re moving to a container-based architecture, these are exactly the kinds of tools you need to do it.

HPE Helion CloudSystem

HPE Synergy with HPE Helion CloudSystem 10 is the first solution that allows you to run both private cloud services and traditional applications on shared infrastructure, resulting in a more agile and cost-effective solution for hybrid IT.

Before Synergy, separate workloads would’ve required separate infrastructure. You would’ve needed one infrastructure for private cloud, another for your HR application, and another for an ERP application.

HPE Synergy's fluid pools of resources can be configured to support a variety of workloads simultaneously, including traditional applications and cloud services. As the needs of the business evolve, more resources can quickly be allocated to growing workloads, and obsolete workloads can return their resources to the available resource pool. Sharing infrastructure has the added benefit of cutting overall IT costs by minimizing over-provisioning.

HPE Helion OpenStack

HPE Helion OpenStack is a flexible cloud platform based on the world's largest open-source project, OpenStack, which adheres tightly to API standards and services. HPE Helion OpenStack offers core Infrastructure-as-a-Service (IaaS) functionality that is highly customizable to simplify the management experience and enable rapid application development and deployment. HPE Helion OpenStack also offers self-service containers orchestration and bare-metal multi-tenant networking. For the most sensitive workloads, your data is encrypted at rest and in transit.

HPE ProLiant for Microsoft Azure Stack

HPE ProLiant for Microsoft Azure Stack is a pretested, factory-integrated hybrid cloud solution that quickly transforms on-premises data-center resources into flexible hybrid cloud services, providing a faster, easier, and simpler development, management, and security experience that's consistent with Azure public cloud services. Co-engineered by HPE and Microsoft, the solution enables customers to rapidly deploy workloads across the Azure Stack private cloud and Azure public cloud, and then easily move those workloads as security, compliance, cost, and performance requirements change as your business grows.

HPE Composable Storage

HPE Synergy creates a pool of flexible storage capacity using multiple storage options that can be configured almost instantly to rapidly provision infrastructure for a broad range of applications and workloads. Multiple high-density storage options let you meet a wide range of application and workload requirements. These options include fully integrated internal storage modules, software-defined storage, and directly

connected, fully composable SAN storage, including HPE 3PAR StoreServ flash arrays. These flexible storage options allow you to achieve the right cost/performance mix based on your specific applications and workloads.

Bringing It all Together with HPE OneView

HPE OneView is the cross-architecture orchestration and management side of HPE's overall infrastructure vision, which include tightly coupled hardware and software components. The hardware side of the equation includes HPE's composable infrastructure portfolio, as well as the company's SimpliVity product line, along with more traditional blade systems and servers. All of those infrastructure components share a common element management vision and provisioning automation engine under HPE OneView.

OneView includes a number of features intended to help your organization stay ahead of the competition and fully embrace Hybrid IT and a cloudlike data center experience:

- ✔ **Faster infrastructure deployment with software-defined intelligence:** Through the use of templates, HPE OneView enables modeling of important settings in software such as RAID configuration, BIOS settings, Firmware baseline, network uplinks and downlinks, and SAN storage volumes and zoning.
- ✔ **Simplified life-cycle operations with proactive updates and frictionless insight:** An intuitive HPE OneView dashboard provides an at-a-glance view of the status of your servers, storage pools, and enclosures. Remote control allows rapid access to system power and remote console whether the server is running or offline.
- ✔ **Increased productivity with a unified API and a growing partner ecosystem:** Integrated with the unified API, HPE OneView supports both traditional environments and the IT environment needed for best practices and extends the power of the Composable Infrastructure to virtualization, facilities management, automation, cloud, and application development.

- ✓ **Accelerated time to value with HPE services and support:** HPE OneView 24/7 support provides the expertise required to address problems with product configuration or operations, as well as conduct defect analysis and resolution. Installation and startup as well as migration services simplify installation and configuration of OneView for a limited number of systems to deliver the best experience from day one.

The Services Front: HPE Pointnext

All the technology in the world won't help you if you don't have the staff or internal knowledge to manage it all. That's where HPE Pointnext comes in. HPE Pointnext is an organization aligned to your needs.

As part of the HPE Pointnext portfolio, there are three types of services across the different stages of your transformation journey: Advisory & Transformation, Professional, and Operational.

- ✓ **Advisory & Transformation Services:** HPE's Advisory & Transformation Services group is at the forefront, where we focus on your business outcomes and goals. HPE designs your transformation and build a road map tuned to your unique challenges to help you digitize the core, innovate offerings, and drive better experiences for your customers.
- ✓ **Professional Services:** The HPE Pointnext Professional Services team specializes in flawless and on-time implementation, on-budget execution, and creative configurations that get the most out of software and hardware alike.
- ✓ **Operational Services:** The Operational Services team understands that success means being accountable for the whole solution, accountable across your ecosystem, and accountable across your old and new infrastructure and apps.

Chapter 7

Enabling Hybrid IT

In This Chapter

- ▶ Making sure you get the right mix of people to get the job done
- ▶ Understanding the era of hybrid IT
- ▶ Leveraging the cloud and knowing how to exit

In this chapter, I cover the final items you need to keep in mind as you undertake your transformational hybrid IT journey: the people side of things, as well as how to avoid some of the risks associated with cloud and hybrid IT.

Solving the Staffing Situation

People make your IT department work. Although the business is often hesitant to add more critical people resources, you know how important people are to the equation, so you have a few options when it comes to supporting your hybrid journey.

In-house IT

You can push business decision makers to add appropriate staff to make sure that all aspects of IT can be supported. Make sure your staff is trained on both sides of the equation — the on-premises side, as well as the cloud side.



In-house IT staff can often be the most expensive long-term resources, but they also have the most incentive to make sure things get done right. More important, investing in your own team leads to long-term gains as they build up their base of knowledge. Treated right, your internal staff are loyal and will go to great lengths to make sure your business doesn't fail due to poor IT systems and support.

Consultants

Sometimes you need some expertise, training, or skills short term that would not make sense to hire internally. You may only need a specific set of skills for a short period of time. This is where consultants come in. Consultants are generally experts in their field and can often help you with both the strategic and tactical aspects of your hybrid IT undertaking. Good consultants will also share knowledge with your internal team through knowledge transfer.



If you're looking for consultants to help you in this, consider contacting the HPE Pointnext team. HPE Pointnext consultants can help you assess your infrastructure and all your applications, and provide you with recommendations on where workloads should operate. Plus, they'll be able to help you ensure that the on-premises portion of the environment is running the right infrastructure to meet application demands.

Service providers

You likely have a reseller or two that provides you with hardware and software for your operations. These value-added resellers (VARs) generally have on-staff experts that are intimately familiar with specific technology spaces and can easily augment your current staff. In addition, you may be able to get consulting assistance from your cloud provider.

Cloud provider specialization

There was a day when cloud service providers were jacks-of-all-trades, but masters of none. They supported very general workloads, but may not have always been experts on specific workloads and security concerns around those workloads.

Times have changed.

Today, cloud providers are becoming more specialized to

meet specific workload, industry security, and compliance needs. And HPE is helping to lead the way in working with these cloud providers. If you need help with cloud, visit www.cloud28plus.com. Cloud28+ is a resource intended to help you learn about the cloud and find providers that can help you with you specific needs.

Understanding Hybrid IT Consumption Models

I've talked about how to consume the individual sides of the hybrid IT model — on-premises versus cloud. There are a number of reasons that you should seriously consider a hybrid cloud undertaking of your own.

Getting the best of both worlds

When you have applications that need to be close to the users in your single campus buildings and these apps are latency sensitive or require intensive use of data, you should run them locally. If you need to ensure that all your customers around the world are able to access your e-commerce storefront, or have employees that work on a collaboration tool across your company, that application might best fit in the cloud.

Proper setup and management of a hybrid IT environment ensures that you always have the ability to operate workloads in the location that makes the most sense.

Blending security

You can decide on an environment that is suitable from a security perspective on an application-by-application basis. If you need super-tight local control on an application, keep it local. If you need to deploy an application into a cloud provider's hardened environment, go for it. Many providers now have specialized environments to meet the specific security needs of healthcare and government verticals.

Gaining agility, flexibility, performance, scalability, and economic benefits

Perhaps the most important thing to understand is that, with hybrid IT, you can get all kinds of core benefits. You can match

workloads with an infrastructure that makes financial sense *for that application*. You can blend on-premises security and performance with cloud agility and economics, so you can scale up to support changing business needs.

The point is that hybrid IT enables more flexibility than is possible with either on-premises IT or cloud alone.

Maintaining Hybrid Control

Perhaps the biggest challenge in considering cloud is the perception of loss of control. When you have infrastructure sitting next to your desk, you feel more in control. When you move to the cloud, there is a palpable loss of control because you can't just walk up to a server and do something to it anymore. You're working within the confines established by someone else.

Avoiding cloud lock-in

There's also a larger, overarching issue with cloud, and that's around lock-in. For decades, lock-in has been something to be avoided. IT doesn't want to feel trapped by a solution that could leave the department unable to react if a need changes or if a company goes out of business, changes its cost structure, or some other event transpires that's outside the customer's control. With on-premises infrastructure, at least, you still have your data tucked safely away on your local storage.



Cloud can take lock-in to whole new levels, particularly with the cost that can be incurred getting out if you decide to exit for some reason. You may decide that the provider isn't doing a good job, or you may decide that the workload is better suited elsewhere, or the cloud provider may simply shut down.

You shouldn't necessarily take this to an extreme, but you should consider a multi-cloud strategy that involves two or more providers. That way, if one provider fails, you can shift those workloads to another provider rather than back to your data center, where you may be delayed by the need to buy more hardware.

The major downside is one that you simply can't get around, and that's the cost of getting out. That outbound data transfer is a fee that you won't incur until you try to leave, so make sure you plan appropriately if you decide to exit.

Enabling workload transparency

When you're running services on-premises, it can be easy to think of them as "free" because you own the hardware. With the cloud, you see the specific charges for each and every workload and, in general, providers can tell you up front about what kind of pricing you should expect to see.

In this way, you may actually have more financial control over your workloads when they're running in the cloud rather than on-premises. Of course, if you have a huge sunk cost in infrastructure, the additive cost of a new workload is pretty close to free in a lot of cases, at least from an infrastructure perspective, but you don't have as much insight into what resources that workload may be consuming.

The perception of control

"The cloud provider is down and I feel helpless." This is sometimes cited as one of the challenges in cloud adoption. There is a feeling of helplessness when a provider goes down. If you have a local outage, your team can spring into action. There is a higher degree of control when workloads run on-premises.



In addition, cloud sprawl can become a security and compliance nightmare thanks to rogue business departments — sometimes referred to as *shadow IT*. For example, it's very easy to add credit cards to a cloud provider and standup workloads and have no thought to security, compliance, or data sovereignty. Such scenarios can also lead to cloud sprawl with no optimization at all, as one group does not know what the next is doing.

Chapter 8

Ten Key Hybrid IT Tips

In This Chapter

- ▶ Putting hybrid IT to work for you
- ▶ Avoiding common pitfalls

Here are ten tips to keep in mind as you embark upon your hybrid IT adventure:

- ✔ **Don't leave your people behind.** Any time you mention the word *cloud*, people may become concerned about their jobs and their roles in the organization. Make sure you prepare them and share your complete vision and strategy every step of the way.
- ✔ **Don't ignore training.** On the same note, make sure you support your people by ensuring that they receive training at every step as well. They're going to need to know how to manage the cloud side of the house and how all the integration pieces work.
- ✔ **Get the right on-premises infrastructure.** If you're planning a hybrid IT scenario, make sure that the infrastructure you install locally leverages the latest innovations and can work as closely as possible to what's operating in the cloud. Make sure you have cloudlike capability locally to ensure the smoothest integration between the two sides.
- ✔ **Don't go it alone.** Hybrid IT, at first, can be a complex undertaking, requiring a complete rethinking at both the strategic and tactical levels. Make sure you have a trusted partner, such as HPE Pointnext, standing by your side.

- ✔ **Embrace digital transformation.** Although laced with marketing, the term *digital transformation* is one that is resonating throughout the market. This transformation requires an infrastructure environment that is cloudlike, in the cloud, or a combination of both, the option that is expected to become the standard. The entire reason that organizations undertake hybrid IT implementations is often to support digital transformation efforts.
- ✔ **Turn IT into a services broker.** With hybrid IT, you have the opportunity to transform IT into an organizational service hub instead of simply reacting to requests and demands from the business. Help usher your organization into the future by becoming a hub for which on-premises is just one option among many.
- ✔ **Perform application audits.** When was the last time you reviewed your application portfolio? If you can't remember, it has been too long. Establish a routine review process so that you can make ongoing decisions regarding business applications in order to maintain the business for the long term. Ultimately, application audits will become a routine part of your operations as you constantly rebalance and optimize workloads.
- ✔ **Embrace security.** Security has become far too important to leave to chance. Make it a cornerstone of your hybrid IT strategy.
- ✔ **Embrace shadow IT.** There's a reason that business units go off on their own. As you transform IT, look for the structural weaknesses in IT or in governance that pushed business units to go beyond the borders. Take advantage of software-defined environments that keep your best practices and governances in place while allowing the business user to be agile.
- ✔ **Start today.** Don't wait. Your journey is a long one, so don't wait until tomorrow to get started!

HPE Synergy

Future first. Ready now.

The world's first composable infrastructure, HPE Synergy, is a new hybrid IT engine fully adaptable and ready for everything.

HPE Synergy is designed to bridge traditional and cloud-native applications with a unified API enabling infrastructure to be programmed like code. Transform traditional apps, modernize today's workloads and run tomorrow's workloads—all on one future-proofed engine for hybrid IT.

Sharpen your competitive edge.
Learn more at hpe.com/synergy



Hewlett Packard
Enterprise



Open the book and find:

- Comprehensive overviews of your cloud options
- Concrete definitions of public cloud, private cloud, and hybrid IT
- The criticality of your workloads and why they must be the focus of your efforts
- How you can gain public cloud pricing with private cloud outcomes
- The HPE product portfolio and how it drives your hybrid IT success

The journey to hybrid IT begins here!

In a few years, you won't be able to imagine an environment that forces workloads to run where they don't make sense. Hybrid IT has become the new standard in enterprise IT. Discover how HPE can help guide your journey to this new normal.

- **Hybrid IT** — *It's real, it's important, it's time to get a jump on it*
- **The HPE advantage** — *Your journey to hybrid IT requires leadership and expertise; HPE can help guide you along the correct path*
- **Critical success factors** — *Ensuring your success in hybrid IT*
- **Leveling up** — *How hybrid IT helps you take your company to new heights*

Go to [Dummies.com](https://www.dummies.com)
for more!

Scott D. Lowe is an IT veteran, with ten+ years spent as a CIO and more than ten in the trenches. Today, Scott is the CEO of ActualTech Media, a recognized speaker and author, and a consultant.

WILEY END USER LICENSE AGREEMENT

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