

LEARNING MADE EASY



Hewlett Packard Enterprise
and Intel Special Edition

SMB Hybrid IT

for
dummies[®]
A Wiley Brand



See how hybrid IT
is the way of the future

Identify the best location to
run individual workloads

Build your on-premises
environment

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.

About Intel

Intel (NASDAQ: INTC) expands the boundaries of technology to make the most amazing experiences possible.



SMB Hybrid IT

Hewlett Packard Enterprise
and Intel Special Edition

by Scott D. Lowe

for
dummies[®]
A Wiley Brand

SMB Hybrid IT For Dummies[®], Hewlett Packard Enterprise and Intel Special Edition

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

Copyright © 2018 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 978-1-119-52104-4 (pbk); ISBN 978-1-119-52101-3 (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

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 server. 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! Small and medium businesses (SMBs) have it worse than their larger brethren, too. They have smaller budgets, smaller staffs, and fewer IT resources, but they're expected to get the same — or more — work done in the same amount of time.

The answer for many has been to turn to the cloud. Or is it?

To hear some tell it, the cloud is the only way forward for IT. I disagree. Even for smaller organizations, cloud is just an option. Many companies will — and should — choose to retain an on-premises computing environment in addition to the choices that they make for the cloud.

In this book, you see that there are two sides to the server 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 or server room computing, budgets, security, and operations, as well as the cloud. 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.



REMEMBER

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



TIP

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



TECHNICAL
STUFF

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 64 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.

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

Chapter 1

Your Workloads, Your Options

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 mishmash 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 servers and throw everything into some cloud provider's environment instead. It's madness, I tell you!



REMEMBER

Fortunately, there is a happy, productive, cost-effective, and secure middle ground that straddles the cloud and your traditional approach to IT. 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 and the 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 server room 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 server rooms 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 server room 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. Increasingly, organizations are turning to various types of cloud service and application providers to augment their local capabilities, but these are just a few options among many.

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. For smaller businesses, it's not as desirable to operate these types of environments as it once was, though.



REMEMBER

For this and other reasons, 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. They say that the need to have staffing resources that can handle various aspects of on-premises infrastructure is simply too expensive, particularly for smaller organizations.

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 the right staff and when that organization makes smart hardware and software decisions. Plus, although many people discuss the rapid innovation in cloud as a key driver for cloud adoption, on-premises hasn't exactly stood still and many new technologies are available to lower costs and improve agility. Some examples include containers, hyperconverged infrastructure, increasingly inexpensive flash storage, and more. In addition, there are options available on the market that can assist with reducing upfront payments for hardware and software, thus allowing customers to enjoy a cloud-like financial experience with their on-premises server environments.

Managed service providers

People are the most expensive resource that you have in the budget. And for small and medium-size organizations for which technology is not the primary business, hiring a plethora of these expensive people isn't always desirable. In addition to being expensive, they can also be unreliable at times, and it seems like there are never enough of the right ones around when you need them. Finding people with the perfect combination of skills and experience to be able to meet all your business goals can be exceedingly difficult. 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.

Managed service providers may directly manage a specific service for you, but they're sometimes used to bring in skills that would be otherwise difficult to obtain. In such cases, organizations should be prepared to pay more, per capita, for managed service provider resources, at least in general. After all, there is typically a profit motive behind managed service providers, and they often have high expenses in maintaining the skills of the people they employ. As a customer, you should expect to pay a bit of a premium for these services. However, you get the benefit of someone else paying to maintain skills, and you don't need to expend effort in identifying and hiring talent. Plus, if you don't need a full-time person in a particular role, use a managed service provider to provide fractional staffing.

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 and stay there; other organizations will stay 100 percent on-premises. Others will start as very small businesses in the cloud and will then implement on-premises services as they grow. The mainstream will be somewhere in the middle, with a mix of public cloud services and on-premises infrastructure. It's really akin to the historical adoption of virtualization. Where it has made sense, companies have adopted technologies that help them improve the bottom line and drive the business.

You'd be forgiven if you thought that everyone was shutting down on-premises servers and shifting everything to the cloud. There has been a lot of talk about how much of the market public cloud providers will own in the future.

But why *aren't* people just flocking to the cloud and never looking back? In many cases, it comes down to cold, hard cash. In fact, as you peel back the covers a bit, you discover that 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. That has a sobering impact.

Some people refer to this phenomenon as the *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 teetotaler with regard to cloud, but they're warily adopting services in a far more targeted way, which is as it should be.



REMEMBER

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

WORKLOAD MIGRATION: THE REAL DEAL

So often, when talking about the benefits of cloud, people espouse the fact that you can just expand and contract workloads whenever you like. Even better, you can just “lift and shift” your on-premises applications to the cloud — and back — any time you want.

Technically, these are correct statements. But, as is so often the case, there are some important caveats.

Let’s start with expanding and contracting workloads at will. In theory, this sounds, well, awesome! In practice, though, some challenges emerge. First off, not all applications can flex. Consider a legacy database application that you run on-premises. Can you *really* expect it to seamlessly support part of the database operating locally and another part operating in the cloud? Of course, for modern applications that are designed with sharing and cloud in mind, this may be no problem, but the world still relies on a whole lot of applications that don’t have this kind of intelligence.

Second, let’s talk about the idea that you can just shift workloads back and forth between on-premises and cloud. Sure, you *can*, but only if your budget has done something to offend you and you’re looking for a way to punish it, because doing this will wreck it. Much like another industry, albeit an illegal one, there are certain free hooks that make it really easy to buy into the market. But, as soon as you try getting out, bad things happen. In the world of cloud, you can place as much as you want in the cloud and just pay for what you use. In many cases, if you try to go the other way — get your stuff *out* — you quickly discover the concept of *egress* charges. Egress charges are fees paid based on how much data you move out of a provider.

Maybe cloud is more like a Roach Motel, but for enterprise IT applications? Workloads go in, but they don’t come out!

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, and for many companies, the public cloud is the perfect solution.

Let's start with the 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



TECHNICAL
STUFF

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.
- » **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 two items mentioned earlier (most often, repackaged SaaS offerings).



REMEMBER

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 Comparing the Responsibilities of Cloud Offerings

Responsibility	On-Premises	IaaS	SaaS
Application	You	You	Provider
Data	You	You	Provider
Runtime	You	You	Provider
Middleware	You	You	Provider
OS	You	You	Provider
Hypervisor	You	Provider	Provider
Server	You	Provider	Provider
Storage	You	Provider	Provider
Networking	You	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, but also some risks.

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 often selected to replace on-premises Exchange environments because it's simpler to deploy and far easier to administer than Exchange.



TIP

Exchange is just one example. For every workload in your portfolio, you need to develop a routine process that analyzes the need for that application and 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.

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

Chapter 2

Hybrid IT Critical Success Factors

Have you ever heard someone say something like, “No more hardware! We should put *everything* in the cloud!” Or, alternatively, maybe you’ve heard that “*nothing* should go into the dark, scary cloud!”

You need to be armed with responses to these kinds of blanket statements and to just about any question you’re asked about your strategy, because, in 99 percent of cases, neither statement is particularly helpful.

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. That's where infrastructure comes in.

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 on-premises servers options out there intended to meet the needs of everyone. But in reality, they often have some drawbacks, including being improperly sized to meet the needs of every 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.



REMEMBER

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.

Or does it?

I want to make it clear that I'm a fan of hyperconverged infrastructure solutions such as HPE's SimpliVity product. These kinds of solutions are often described as "one size fits all" and "data center in a box" products. Based on what I just said, you may wonder why I'd believe that these products are good for anyone.

There are two reasons:

» I believe that for small and medium-size companies, there are other driving factors beyond just how efficient you can be with every single discrete resource in the environment. There comes a point at which trying to

perfectly finagle every resource into complete submission takes too much time and effort — and, as a result, cost — and it just makes more sense to take a cookie-cutter approach and add more appliances as you need more resources.

» **Although hyperconverged solutions are often discussed as “one size fits all,” they actually do have some level of resource configurability, so you’re not just buying a solution that has a single SKU and then having to live with it.** With that in mind, it becomes easier to understand how these kinds of solutions can accommodate the kinds of flexibility that are needed in today’s SMB environment.

Maintaining Performance Prerequisites

One of the challenges in maintaining a robust on-premises data center environment is maintaining necessary performance prerequisites even 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 for no apparent reason.

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.

WORKLOAD PERFORMANCE PERCEPTION AND RESPONSIBILITY

By the way, these users won’t care if it’s actually you providing the environment or it’s someone else. They’ll also raise concerns to *you* about workloads you’ve outsourced to a managed services provider or for which you’ve subscribed to a SaaS solution such as Office 365. The point is that ensuring consistent performance isn’t just an on-premises need. You need to carefully choose providers and ensure that contract terms, where possible and appropriate, include performance guarantees.

Containing Workload Runtime Environment Costs

You've probably heard that many companies — frankly, often in the SMB space — view IT as nothing more than a cost center to be managed. As a result, very few want to spend more than absolutely necessary on IT. Unfortunately, this commodity-centric view of IT ends up having real-world impacts on the business, but, in the end, price often wins.

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.

The fact is that 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. Of course, that assumes that there is even a formal replacement cycle. There are SMB environments out there that push the limits of hardware replacement in pretty extreme ways.

Let's talk about how equipment is purchased in an SMB that uses traditional funding methods and a simple capital expenditures (CapEx)/operating expenditures (OpEx) methodology. At some point, the chief information officer (CIO), IT director, CFO, or managed services provider partner indicates that an asset is due for renewal. Knowing that the funding cycle for that asset only comes around every three to five years, that company buys what it thinks it'll need until it's time to replace that asset again.

To ensure that the infrastructure can meet the needs of the business for years to come, there is a tendency to overbuy. Simply put, the company sometimes buys too much because it doesn't want to have to spend again. This isn't always the case, of course. Sometimes, SMB budgets are simply so constrained that companies can only focus on the here and now, and so they buy what they need right now with the idea that they can revisit later if needs change.

DEPRECIATION SCHEDULE VERSUS REPLACEMENT SCHEDULE

Some companies tend to use their depreciation schedule as a marker for their replacement schedule. If that new storage array is fully depreciated after a three-year period, that's considered its useful life. Larger companies are often better about sticking to formalized replacement schedules than smaller ones, but, as is the case in all things, it really does come down to money.

Just because something is fully depreciated doesn't mean it's lost all of its operational value. For companies that are working hard to minimize IT costs, it can be desirable to extend the replacement schedule for equipment that remains viable. Doing this means that these companies get some "free" years of service, after which they can make the decision to eventually replace the equipment once it's no longer meeting business needs.

But that's nothing compared to out-of-cycle need to spend. SMB organizations don't always plan the next quarter's needs for three to five years out, but IT has to plan infrastructure that way, and if they miss on their forecasts or if business needs unexpectedly change, they're often forced to perform upgrades off-cycle, which can be expensive and difficult to manage.



REMEMBER

For ongoing on-premises infrastructure deployments, organizations simply have to find ways to become more "cloud-like" in their approach, at least financially. There are plenty of infrastructure offerings that can help with this, such as hyperconverged and composable solutions, which help make it possible to get closer to a "pay as you go" methodology.

With all the focus on getting to this much-ballyhooed pay-as-you-go state, you may think that your quaint "buy everything I need and pay for it all upfront" approach is antiquated. In reality, you have a bit of a secret weapon: You can absolutely abuse it and push it to its limits, and doing so doesn't cost you a penny more. You don't pay *any* additional usage fees for stuff you already own.

If, however, you gaze longingly at ads for pay-as-you-go cloud services, did you know that you can have your infrastructure and cloud economics, too? There are numerous 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. For example, under HPE's Greenlake program, HPE retains ownership of and supports your choice of on-premises equipment. You just deploy it and 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. When 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. Most SMBs don't have the resources necessary to fully monitor and manage everything necessary to ensure complete cost compliance.

Corralling Infrastructure Control

In recent years, there has been a movement toward empowering users 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.



TIP

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. An MSP or an internal IT generalist actually manages the environments operationally across on-premises or off, based on the business needs and the business benefits. Most users and departments don't want to own the operational or IT part, but they want choice in IT as well as flexibility and agility. The move to shift IT to the role of a services broker while transferring operational administration to users and departments is a good first step toward organizations trying out the hybrid cloud waters.

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. The valet key allows that driver to accomplish his job while having some reasonable limits.



TECHNICAL
STUFF

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) on storage arrays or that sales will begin to create administrative user accounts in Active Directory. But these groups should be able to access resources

in a sufficient enough way to be able to meet their goals. They need an IT valet key that can help them get their jobs done more efficiently.

Centralizing security control

If you read nothing else, read this section!

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.

It's generally understood that many, many SMBs don't have chief information security officers (CISOs). They do often have staff with security as a part of their responsibilities, though. Plus, some SMBs have contracts with outside managed services providers that help the company reduce security risk.

For departments that may have some of their own resources, one area that can't be negotiable is generally IT security. The business faces too much risk otherwise, and there does need to be some "command and control" to prevent incidents. That said, for specialized security or compliance needs, the entity responsible for overall security should be ensuring that there is continued collaboration to make sure that everyone remains aware of requirements and potential risks so that appropriate decisions can be made.

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. Even small and medium-size businesses face challenges related to IoT, if not to the extent of their larger brethren. IoT and device proliferation, in addition to posing some new resource challenges associated with storage and the like, also exposes the potential for new security challenges that may not be otherwise present.

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.



REMEMBER

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.

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

Chapter 3

IT's All about the Workloads

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 and enabling the business as a whole. In an SMB environment, this can be incredibly difficult for reasons that I allude to elsewhere in this book, but skill set remains a key inhibitor. Depending on location, it can be hard for smaller organizations to find people with the mix of skills they need. A larger company may be able to afford to hire someone with 25 years of storage knowledge, but a smaller company may need a single person that can handle storage plus virtualization plus cloud.

In this chapter, you see how a focus on workloads remains a hallmark for IT departments, regardless of company size.

Assessing Workload Infrastructure Needs

Every workload is different. Even the *same workload* in a different company 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. For example, traditional client/server applications, in which there is a centralized database server with the client side installed on local client PCs, are not generally great options to move to the cloud.



TECHNICAL
STUFF

Of course, 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. Plus, the Outlook client is purpose-built to support potentially high-latency environments.

With other client/server applications, though, imposing latency on the application directly impacts the user's ability to get work done. It slows them down because the client portion of the application can't elegantly handle latency.

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. With the need to run a full client, there is a limit for how far away the server side may exist.

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 preceding section, but also for other reasons, the network can be a major concern in the world of cloud. In fact, concerns about the network are one of the main reasons that application workloads are often left in the on-premises servers, close to the users, or at the edge, close to where data is being generated or created. After all, the closer applications are to their data, the faster those applications can respond.



REMEMBER

It's all about physics. Distance results in latency, and there's absolutely nothing you can do to change that. 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 and dissatisfied users and possibly customers.

But there's more: Imagine that you have a satellite view of both your on-premises servers and the cloud provider you intend to use. You may be able to draw a straight line between the two points, but in reality, there are myriad other networks in between those two locations. In addition to increasing latency, as you throw more networks between the user and an application, you're increasing the likelihood of a network outage, which would bring down access to that service.

Performance and reliability are the linchpins to successful IT. If you have a highly centralized workforce in a single building or a single campus, keeping workloads close by may 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.

You may think that these distributed enterprises are the purview of only very large companies, but modern businesses are very different from companies from the past. In fact, I work with two very distributed small/medium organizations that don't even have physical headquarters; everything happens virtually. So, the need to support a distributed workforce is a problem faced by companies large, medium, and small.

Defining the right mix of infrastructure: Public and traditional

For many SMB organizations, it's probably safe to say that there is a significant desire to adopt cloud in some form, but it may not be adopting the likes of Amazon Web Services (AWS). Instead, Software-as-a-Service (SaaS) tools would be most applicable for these types of businesses.

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.

THE SaaSy WORLD OF THE SMB

For a lot of SMBs, SaaS is where it's at. Some SMBs may need the likes of AWS or Microsoft Azure services, but many just want to adopt a series of SaaS-based applications that they can then easily integrate with their on-premises environments to provide users with a consistent experience. More and more, companies are working hard toward integrating their Active Directory environments with SaaS services to, for example, enable single sign-on (SSO) capabilities. In some cases, these customers may run an instance of their Active Directory environment in Azure, to make this integration a bit easier.

In these SaaS-centric hybrid IT environments, IT is responsible for ensuring that the various integrations that need to take place are actually working. Such integrations include centralized authorization and authentication (often accomplished via services such as Active Directory Federated Services [ADFS]), identity management/service provisioning (to bring SaaS apps to first-class citizen status), and data protection (not all SaaS providers do backups in a way that meets everyone's requirements).

So, even though SMBs may not be running the full suite of Amazon services, there's still plenty of hybrid IT to rein into submission.

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.

Of course, as I mention earlier, there may be a series of SaaS services that work better for your business than some of your existing traditional applications. In other cases, you may actually want to move certain workloads to AWS or Microsoft Azure. And in others, you may want to keep them right where they are.

Understanding Workload Locality

What is locality, anyway? Simply put, *locality* defines where various parts of your applications will run. 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 should live in the cloud. There are serious operational, performance, and even regulatory details that you need to consider before you finalize your locality decisions.

Supporting legacy applications

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



REMEMBER

Plus, 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 your data security needs don't conflict with regulations and you want to move your legacy client/service applications to a cloud provider, you can get around the resulting application latency issues, 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 may 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.



TIP

I've deployed exactly these kinds of systems for clients, and I'm generally somewhat hesitant to do so unless there is a compelling cost or operational reason for moving forward in this way. The resulting complexity is often much greater than what would result by just retaining the workload in-house. When I have deployed such systems, I've generally recommended that the client — mostly SMB clients — retain the services of a managed service provider to help them maintain the solution because tracking down problems can be somewhat difficult.

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.

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.



TIP

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.

How you burst really depends on the underlying application architecture. If you're operating a very traditional client/server-centric database application, "bursting" may mean lifting and shifting a workload to the cloud for busy season and moving it back when busy season is over. By doing so, you gain the ability to more quickly scale the resources assigned to the application.

If you have a traditional client/server environment that is architected in a standard three-tier architecture — front-end/web server, middleware, and back-end/database servers — you can sometimes choose to spin up more resources in the cloud as you need to add more resources to any of the tiers. However, you need some pretty comprehensive networking rules established between the cloud provider and your on-premises environment.

For bursting to be most effective, you need to operate a more modern application type that has some cloud-native intelligence. These may be virtual machine- or container-based workloads that can seamlessly integrate cloud and on-premises components to form a cohesive application for the user. Unfortunately, these kinds of applications often require specialized developers, something that many SMBs lack because they prefer to buy rather than build their workload operating environments.

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

Chapter 4

Workload Control

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 will 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.



TIP

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

Chief information officers (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 that are prime candidates for the cloud continue to rise, there may be a sense of loss of control from internal IT as an external service provider assumes responsibility for certain functions. IT has always been about control, and that loss can be difficult to overcome, particularly when cloud-based workloads become unavailable, and IT is left awaiting a response from a provider rather than jumping into action to take corrective measures. In these scenarios, IT is largely on the sidelines as they wait for return to service from the provider. It's an uncomfortable and unfamiliar place for IT professionals, who, with on-premises infrastructure failure, are used to doing things on their own.

Getting Clarity on Each Group's Role

As you work through determining the appropriate location for each of your individual workloads, you also need to decide who will 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 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.



REMEMBER

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. In the biz, we refer to this as “finger pointing,” and it's not a desirable activity.

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 says that moving to the 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 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 companies that choose on-premises infrastructure, 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.



REMEMBER

For those who operate in highly regulated arenas, such as the finance and healthcare industries and the federal government, 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. Security can be hard work, and even the most severe attacks can go unnoticed for months. In fact, many organizations are under regular attack.

Like cloud providers, you may be under regular attack and not even know it. Or you may be painfully aware, as your Internet services slow to a crawl or other bad things happen, like a ransomware scammer locking up all your files.

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 much larger and 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.



TECHNICAL
STUFF

A *zombie* is a computer or other smart device, such as a router, thermostat, HVAC controller, or anything else with a processor and operating system, that has been taken over by bad guys with the intent to do harm. These devices may be programmed to lay in wait until they're activated by the bad guys, at which point they stage coordinate assaults against their targets.

THE SOBERING STATISTICS BEHIND SMB SECURITY COMPROMISES

No company wants to suffer a security compromise. For SMBs, the risk is even greater than it is for larger organizations. SMBs often don't have the resources they need to combat the bad guys, so they're disproportionately impacted when something goes awry. The U.S National Cyber Security Alliance has a statistic that may keep you awake at night. Did you know that 60 percent of small businesses close their doors within six months of a cyber attack? For those that are able to weather the storm, the average recovery cost is a staggering \$690,000. And, in further research, IBM discovered that small businesses are the ones hit by 62 percent of all cyber attacks.

When it comes to certain kinds of attacks, the news is even worse. SMBs are the unlucky recipients of 75 percent of all spear-phishing attacks. SMBs are disproportionately impacted by ransomware attacks as well.

It's clear that the bad guys see SMBs as soft targets and place their efforts there. Don't be a statistic!

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, depending on the nature of the attack.

Just about any cloud provider of any reasonable size typically has multiple points of connectivity to the Internet. So, even though they're impacted by DOS attacks, the impact may be more muted.

You do have to be careful, though, because your cloud provider's terms of service can work against you in certain cases. Some service terms with cloud providers may actually allow them to shut down your virtual machines and other workloads if your site experiences a DOS attack. You won't have any control over this. Bear in mind that the cloud provider has to service *all* its customers. If your workload is having a negative impact thanks to an attack, you may pay the price.

Worse, when a DOS attack does take place, you may be in for a nasty surprise at the end of the month. Although an attack may not impact every customer on the service, the client that did get attacked may actually face a steep increase in its monthly bandwidth bill because the attack used resources that were allocated to that client. Pretty nasty.



TIP

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. If you're planning to have resources that live in the cloud, look for ways to protect yourself.

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. Even worse, you don't even need to do anything to be affected. There have been reports of compromised servers being shipped to unsuspecting customers. That's not good.

In some cases, firmware compromises require physical access to the server. This is just one reason that allowing only authorized people into a data center is so important, even in small organizations. It's also important to log data center access so security events can be correlated against visitor logs. You may not want to *blame* someone for a firmware compromise, but you may want to be able to correlate so that you can determine how a compromise was introduced so you can prevent it from happening again.

But there's more to the security story, and it begins before you even accept delivery of a server into your data center. You need to make sure you've chosen a vendor that has a mechanism in place to ensure that no one and nothing has tampered with your server prior to arrival in the data center. Not all server vendors can make that promise, so choose wisely!

Staff

The weakest link in any security system is the human element. People can make mistakes. People *will* 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.



REMEMBER

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

In a hybrid IT environment in an SMB, more people will have more interactions with more workloads. The time is now to begin creating and adopting stricter data and application access policies and procedures that you can put into effect to protect your organization.

IN THIS CHAPTER

- » Paying for your data center infrastructure
- » Managing the people side of the infrastructure equation
- » Considering the pitfalls of the cloud

Chapter 5

The Real Costs of Workloads

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. You probably have an owner who is consistently pushing back on requests for new

expenditures and is always looking for ways for you to squeeze more life out of the systems you've already deployed.

Some owners are more concerned with how money is spent than with what it's spent on. If this sounds familiar, you may need to develop financing models that help your owner achieve company goals. In other cases, you may need to turn to completely different models to make it all work. Plus, as you develop these models, you need to remain mindful of the impact that your decisions could have on IT staffing requirements.

There's a lot to think about. The following sections outline what you need to keep in mind.

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 are 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 implies, 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 DIYers 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.



TIP

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.

Of course, there is a cost — a direct one — in going this route, and there is always a trade-off. The great part of using a managed service provider is that you get access to the skills you need. Plus, if you really just needed a fractional resource — perhaps a part-time storage administrator — you no longer need to hire full-time people. But if you compare the costs, you're probably paying a bit more for the service on that basis. After all, the provider has to turn a profit as well.

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.

Regardless of which path you choose, remember this: There are a minimal number of payments you must make before you can terminate your lease. With a cloud-based operating paradigm, you can generally just end the service whenever you like. With a lease, however, you'll need to buy it out if you want to be able to terminate it early.



REMEMBER

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. Each of these points is described in the following sections. Before getting started, though, remember that many cloud providers are very public about their pricing. They often have calculators or spreadsheets available for download that you can use to estimate your spend. Sometimes,

these are called “simple” calculators, but you may find them anything but. Even though estimating on-premises utilization can be a hard task, doing the same for cloud can be even harder. For example, locally, you probably don’t care about network ingress and egress. Those costs are “free” because you own the pipe. With cloud, though, network egress is anything but. More on that in the “Data transfer and getting out of the cloud” section.

Acquisition

Cloud service acquisition costs are often pretty low, at least compared to buying servers 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.



TIP

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 (more often) they won’t. It’s up to you to figure out how much you need and when you need it.

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 Amazon Web Services (AWS). For this example, let’s assume that, on average, you move 500TB of data per month 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

Outbound Data Transfer	U.S. Zone (\$/GB)	Move 500TB Out
First 1GB per month	\$0.000	\$0.00
Up to 10TB per month	\$0.090	\$899.91
Next 40TB per month	\$0.085	\$3,400.00
Next 100TB per month	\$0.070	\$7,000.00
Next 350TB per 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 servers. 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.

Remember that calculator that I mentioned? In order to even begin to estimate your cloud spend, you need to know a whole lot about your workloads. You need to understand exactly the resources that your workloads will use, including how much data they'll incur in egress bandwidth. As you move toward a hybrid IT model, you're going to need to understand this kind of networking capacity so that you don't end up on the floor after getting a bill.

For all the simplicity that cloud can bring from an operations standpoint, it doubles complexity on the financial projection side of the house. It's not uncommon for organizations going to the cloud to need to hire people that are dedicated to managing cloud spend so that nasty surprises are kept to a minimum.

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

Chapter 6

HPE Hybrid IT Infrastructure

The cloud isn't a perfect place for all workloads, but then again, neither is your on-premises server. This book shows you the importance of analyzing your workloads across a variety of characteristics to decide where it's best to place each one. 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 everything else, there's on-premises infrastructure.



REMEMBER

On-premises infrastructure encompasses everything on your on-premises server.

Getting the Right Workload Environment Mix

Modern enterprises are incredibly dynamic places with often conflicting needs. On the one hand, today's enterprises are supporting systems of record that contain transactional data critical to the business's health with the digitization of everything else,

regardless of source. Most enterprises need to figure out how to operate in both the cloud and on-premises worlds and bring the two together in a successful way.

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 public clouds, and this will continue well into the future. It's not an easy scenario to support. It's complex. But IT really has no choice. Now more than ever, IT is required to rapidly accelerate the business and ensure the flexibility necessary to allow the business to pivot as demand changes.

The key to supporting these dynamic worlds really lies in the ability to compose resources that can support every workload, as needed. To achieve this, enterprises must work to find the right mix of hybrid IT across traditional and public cloud environments. It's different for every enterprise, every industry, and it changes over time. Some may need an environment that emphasizes the public cloud, while other may place more value in on-premises solutions.

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

- » **Simple, secure, and affordable**, delivered in today's reality of a hybrid IT environment — whether your apps and data are on your on-premises server, the public cloud, or multiple clouds — and with the flexibility to operate seamlessly across these environments.
- » **Able to capitalize on the convergence of Internet of Things (IoT); apps; and data closest to your customers, employees, users, and transactions** — the intelligent edge — where your enterprise touches customers, 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.

WHAT IS THE “EDGE”?

More and more, vendors are using the term *edge computing* to describe environments in which data maintains residency closer to the source that generated that data. Edge computing is rising in popularity thanks to such trends as the rise in IoT devices and even autonomous vehicles. Edge computing, like other environments, can stand alone, be connected to a larger on-premises environment, or even connect to the public cloud. The point here is that the discussion around on-premises versus in the cloud isn't complete until and unless you develop a plan for how you're going to handle edge devices that may require a microcosm of your server room environment to be distributed across a potentially vast landscape.

An Array of Options

For on-premises infrastructure, HPE makes available a wide array of platforms supported by a variety of 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. You're not stuck with a one-size-fits-all application environment, nor are you constrained by one-size-fits-all partners. You get to choose the platform and partner combination that makes the most sense for the workloads you need to support balanced with the needs and capabilities you have around IT staffing.

Although today's hyperconverged infrastructure 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.

HYPERCONVERGED AND COMPOSABLE INFRASTRUCTURE: GETTING EDUCATED

You've probably heard of hyperconverged infrastructure, but *composable infrastructure* may be a new term for you. Regardless of your existing level of knowledge around these two terms, I recommend that you bone up on these technologies to learn what they can do for you. The following links point you to two resources on these very subjects:

- www.hpe.com/us/en/resources/integrated-systems/hyperconverged-infrastructure-dummies.html
- www.hpe.com/us/en/resources/composable-infrastructure-for-dummies.html

You can choose to modernize and scale server room elements independently in a traditional way or consume them as a part of more modern solutions, such as converged, hyperconverged, and composable offerings, but regardless of the direction you go, taking advantage of new compute innovations, storage data services, and low-latency networks is a critical part of a hybrid IT foundation. The following sections describe the product portfolio elements that HPE brings to bear to help you.

HPE ProLiant

HPE ProLiant servers are a cornerstone in many server rooms and are a rock-solid choice for those wanting to pursue a traditional infrastructure strategy. They can help lower costs, boost productivity, and keep small or midsize businesses running smoothly.

The HPE ProLiant ML family of servers delivers simple, efficient business value and is the ideal choice for small businesses and for larger organizations that need to support remote or branch office environments.

HPE Gen 10 servers also feature cutting-edge security capabilities intended to help protect SMBs from the activities of nefarious netizens. There are a number of features of note:

- » **Silicon root-of-trust:** Ensure your Gen10 ProLiant server can never boot with compromised firmware through an unalterable digital fingerprint in the silicon that verifies the firmware code is valid and uncompromised.
- » **Server system restore:** Recover the server's essential firmware, firmware configuration settings, operating system, and host environments back to an operational state quickly.
- » **Firmware threat detection:** Leverage runtime firmware verification to regularly (as often as daily) check firmware validity, and automatically recover essential firmware to a known good state after detection of compromised code.

For the SMB, HPE has also developed a new operating model called HPE SMB Offers, which focuses on putting the customer first, driving the best opportunities, and simplifying choice through targeted offers for the volume market. SMB Offers delivers simple, pre-configured hardware and complete solutions that include the software applications today for these use cases:

- » **Unified Threat Management Solution:** Provides multiple gateway applications that can replace physical network gateway devices to keep your network secure.
- » **Storage and Backup Solution:** Stores, backs up, and shares media and other large files and applications with a secure central server, not a cloud service. This can help protect you from the potential for loss from individual devices, which are susceptible to viruses, malware, and theft.
- » **Virtualization Solution:** Enables your business with multiple line of business apps to increase efficiency and enjoy the benefits of virtualization.
- » **Multi-function Solution:** Enables your workforce and increases employee productivity with multiple apps for communication, collaboration, and email.

The complete server package is, of course, the critical element, but server packages are empowered by the choice of the components that comprise them. In this case, Intel Xeon Processor-based servers enable small and medium business IT needs. These processors, when paired with HPE server hardware, provide the performance, security features, and cloud-readiness small businesses need to thrive and grow. They offer the value and flexibility

that provide a strong foundation for success for a growing SMB and carry with them a number of features:

- » Reliability, security, and performance needs for a really small business server
- » Affordable, on-premises server based on the Intel Xeon Processor E3 family
- » On-premises server plus cloud, for the best small business solution



TIP

Learn more about Intel's SMB-centric solutions at <http://intel.ly/smbserver>.

Intel is more than just processors. The company's solid state drives — Intel 3D NAND SSDs, powered by Intel technology — help growing businesses transform the economics of storage with trusted, breakthrough 3D NAND technology. 3D NAND drives are architected for capacity while also maintaining the performance characteristics of flash storage. Better yet, such devices help you to reduce your costs. By cramming more data into less space, 3D NAND drives are space and power efficient at capacity, a combination that reduces overall total cost of ownership (TCO). The partnership between HPE and Intel is a strong one. HPE provides a comprehensive server package that includes components from Intel's rapidly expanding portfolio, and you get to enjoy the results of this combination.

HPE storage

HPE storage — including HPE 3PAR and HPE Nimble — is a bed-rock choice for the server room 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, reining in capacity costs with comprehensive deduplication and compression features. With HPE Recovery Manager Central and HPE StoreOnce, HPE Storage solutions ensure business continuity with reliable and cost-effective data protection.



REMEMBER

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

HPE data center networking

All this talk of servers, storage, and cloud makes it easy to forget that all of it has to connect together somehow. That's the job of the network. More powerful compute and faster storage require a new breed of data center networking solutions. Whether you're refreshing traditional applications or rolling out new cloud applications, HPE FlexFabric switches can help you accelerate performance, simplify operations, and contain costs.

Specifically designed to meet the increased capacity and resiliency requirements of today's on-demand data centers, HPE FlexFabric 10/25GbE switches are ideal for your most bandwidth-intensive and latency-sensitive applications. HPE FlexFabric 10/25GbE solutions can help eliminate sprawl and simplify your data center environment by converging compute and storage networks, and ensure fast, reliable connectivity for IP-based (iSCSI or FC) server or storage networking requirements. HPE FlexFabric switches offer simplified licensing and flexible consumption models and don't require additional costly add-on feature licenses because all advanced features are included by default.

HPE's server and networking solution just work better together. The combination helps you avoid potential interoperability issues and vendor finger-pointing by leveraging a single point of contact for compute, hypervisor, storage, and networking solutions.

HPE Aruba Networking

HPE's data center networking solutions may be overkill for some SMBs. That's where HPE Aruba's SMB networking solutions come into play. Aruba provides an integrated wired and wireless access portfolio, simplified network management, and built-in security to support a growing number of mobile devices and cloud-based applications while minimizing business disruption.

HPE Aruba's solutions feature enterprise-grade Wi-Fi, switching with HPE Smart Rate ports — a multi-gigabit (1Gbps, 2.5Gbps, 5Gbps, 10Gbps) twisted-pair network interface — zero-touch provisioning, and flexible management options for wired and wireless connections. Built-in security, fine-grained access control, and smart application handling allows organizations to adopt a bring your own device (BYOD) environment and prioritize business-critical applications to improve workplace productivity.

Whether wireless or wired, it's networking made simple, reliable, and affordable, so organizations can focus on their business, not managing their network. Regardless of company size, modern networks need to feature some critical characteristics, including the following:

- » **Cost effectiveness:** Controllerless wireless access points (APs) eliminate hardware and maintenance expenses, but integrate all critical components for ease of management.
- » **Reliability:** Accommodates a range of device and traffic types, including data, voice, and video, and prioritizes business-critical applications to keep your business online and productive.
- » **Efficiency:** Flexible management options including local, operating expenditure (OpEx)-friendly cloud subscriptions with always up-to-date software or on-premises multi-vendor wired/wireless management.
- » **Security:** Incorporates advanced secure technology, such as built-in firewall, built-in web security, and wireless intrusion detection and prevention.
- » **Futureproofing:** A flexible architecture provides investment protection as your business grows.

HPE iLO: Bringing it all together

HPE Integrated Lights-Out (iLO) allows you to configure, monitor, and update your HPE servers seamlessly from anywhere in the world. Providing you with consistent insight into the health and operation of your servers, HPE iLO arms you with the tools to resolve issues and keep your business running. Featuring the latest innovations in simplified operations, performance, and security, HPE iLO allows you to manage your entire server environment with ease. Upgrade your licenses for additional functionality, such as graphical remote console, multi-user collaboration, video record/playback, and much more. Learn more by visiting www.hpe.com/info/iLO and reviewing the Licensing Guide.

iLO ADVANCED PREMIUM SECURITY EDITION

Exclusively available on HPE ProLiant Gen10 servers, the iLO Advanced Premium Security Edition license delivers all the management capabilities of the iLO Advanced license with premium

security features like Server System Restore, Commercial National Security Algorithm (CNSA) mode, Runtime Firmware Verification, Automatic Secure Recovery, and Secure Erase of NAND/User data. Learn more at www.hpe.com/servers/ilopremium and try a free 60-day trial.

iLO ADVANCED

For larger businesses, this license provides advanced remote functionality and all the HPE iLO features to improve speed, scale, and simplicity. Key features include Integrated Remote Console, Virtual Media, and iLO Federation (rapid discovery, inventory and management at scale). Learn more at www.hpe.com/servers/ioloadvanced and try a free 60-day trial.

iLO ESSENTIALS

This license offers remote server management features that are uniquely designed for small business at an affordable price. Key features include Integrated Remote Console, Virtual Media, and email alerting. Learn more at www.hpe.com/servers/ilo/essentials and try a free 60-day trial.

HPE Pointnext: The services front

All the technology in the world won't help you if you don't have the staff or internal knowledge to manage it. That's where HPE Pointnext comes in. HPE Pointnext has the expertise and experience to help you simplify operations and easily extend your IT staff, helping them do more with less. HPE Pointnext services help you innovate to keep ahead of risks like competitive trends and security threats alike. Providing a variety of offerings — from advising on best practices and project-based engagements, to operational support services — HPE Pointnext helps you keep your environment running efficiently and reliably.

As part of the HPE Pointnext portfolio, there are three types of services across the different stages of your transformation journey:

» **Advisory & Transformation Services:** The Advisory & Transformation Services group focuses on your business outcomes and goals. HPE designs your transformation and builds 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 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.

HPE Financial Services: The financial services front

Earlier, I mention the importance of having the option to leverage a financial structure that makes sense for your business. HPE Financial Services helps to make this possible and provides you the flexibility to procure infrastructure that allows you to run applications where it best serves your business — on-premises or in the cloud.

To make Hybrid IT work seamlessly, timely access to new, advanced technology is essential. But IT budgets and staffs are stretched across many priorities and demands. HPE Financial Services can help you gain affordable access to new, advanced technology, but in a way that preserves cash and helps you be more efficient and handle change. Such investments can accelerate your digital transformation plans and align your IT and financial needs.

HPE Financial Services can help in lots of ways:

- » **Trade in the old.** Make room for new equipment and services by shipping existing equipment back and get money for the technology.
- » **Make a big transition easier.** Securely retire and recover value from your old equipment.
- » **Simple to consume.** Subscribe to a complete technology solution from predefined options for a predictable monthly payment.

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

Chapter 7

Enabling Hybrid IT

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 businesses are 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 that your staff is trained on both sides of the equation — the on-premises side, as well as the cloud side.



REMEMBER

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 that 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.



TIP

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.

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.cloud28p1us.com. Cloud28+ is a resource intended to help you learn about the cloud and find providers that can help you with your specific needs.

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.

Understanding Hybrid IT Consumption Models

I've explained 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 you have employees who work on a collaboration tool across your company, that application may best fit in the cloud.



REMEMBER

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 of 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, if a company goes out of business or changes its cost structure, or if 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.



REMEMBER

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 server room, 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 upfront 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.



TIP

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 stand up workloads and give no thought to security, compliance, or data sovereignty. Such scenarios can also lead to cloud sprawl with no optimization at all, because one group does not know what the next is doing.

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

Chapter 8

Ten (Or So) Key Hybrid IT Tips

Here are ten (or so) tips to keep in mind as you embark upon your SMB 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.** Make sure you support your people by ensuring that they receive training at every step. 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 hardware.** With access to HPE's entire portfolio — from ML series servers to hyperconverged solutions to composable infrastructure — there's a solution that fits every workload need.
- » **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 cloud-like capability locally to ensure the smoothest integration between the two sides.

- » **Don't go it alone.** At first, hybrid IT 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 cloud-like, 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. Even an SMB IT department can deploy services like the big guys.
- » **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.

- » **Don't forget the network.** HPE's data center networking and Aruba wired and wireless solutions provide the glue that binds your on-premises and cloud-based environments. Choose a product that meets your needs and that can grow with you.
- » **Deploy the right financial model.** With HPE Financial Services, you can buy or lease. Get the hardware and services you need to get started right away with a model that fits the needs and financial capabilities of your business.
- » **Start today.** Don't wait! Your journey is a long one, so don't wait until tomorrow to get started!

WORLD'S MOST RESPONSIVE DATA CENTER SSD¹



Intel® Optane™ SSD DC P4800X Series



5-8x faster at low
Queue Depths²



Up to 2.8x more Total Bytes
Written at similar capacity³



Up to 40x faster response
time under workload¹



Up to 60x better
at 99% QoS²

Learn more at intel.com/optane

1. Responsiveness defined as average read latency measured at Queue Depth 1 during 4k random write workload. Measured using FIO² 2.15. Common configuration - Intel 2U Server System, OS: CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU: 2x Intel® Xeon™ E5-2699 v4 @ 2.2GHz (22 cores), RAM: 396GB DDR @ 2133MHz. Intel drives evaluated - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Samsung® drives evaluated: Samsung SSD PM1725a, PM1725, PM963, and PM953. Micron® drives evaluated: Micron 9100. Toshiba® drives evaluated: Toshiba ZD6300. Test: QD1 Random Read 4K latency, QD1 Random RW 4K 70% Read latency, QD1 Random Write 4K latency using FIO 2.15.

2. Common configuration - Intel 2U PCSD Server ("Wildcat Pass"), OS: CentOS 7.2, kernel 3.10.0-327.el7.x86_64, CPU: 2x Intel® Xeon™ E5-2699 v4 @ 2.2GHz (22 cores), RAM: 396GB DDR @ 2133MHz. Drives evaluated - Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600GB. Performance measured under 4K 70-30 R/W, QD1-16 using FIO 2.15.

3. Comparing Intel® Optane™ SSD DC P4800X 750GB spec to Intel® SSD DC P3700 Series 800GB spec. Total Bytes Written (TBW) calculated by multiplying DWPD x warranty duration x 365/year.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.

© 2018 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and Optane are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

The journey to hybrid IT for the SMB begins here!

Companies small and large are adopting cloud services. 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 SMB. Discover how HPE can help guide your journey to this new normal.

Inside...

- Get comprehensive overviews of your Hybrid IT options
- Identify the unique challenges facing SMB IT departments
- See how you can have the industry's most secure hybrid IT environment
- Consider the HPE product portfolio and how it drives your hybrid IT success



Scott D. Lowe is an IT veteran, with 10+ years spent as a CIO and more than ten in the trenches. Today, Scott is the CEO of Actual-Tech Media, a recognized speaker and author, and a consultant.

Go to **Dummies.com**[®]
for videos, step-by-step photos,
how-to articles, or to shop!

for
dummies[®]
A Wiley Brand

 Also available
as an e-book

ISBN: 978-1-119-52104-4
Not for resale

WILEY END USER LICENSE AGREEMENT

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