

LEARNING MADE EASY

Nutanix 3rd Special Edition

# End-User Computing

for  
**dummies**<sup>®</sup>  
A Wiley Brand



Discover  
virtual desktops

—  
Accelerate  
deployment

—  
Manage end-user  
economics

Brought to you  
by

**NUTANIX**<sup>™</sup>

Scott D. Lowe

# About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power the business. The Nutanix Enterprise Cloud Platform blends web-scale engineering and consumer-grade design to natively converge server, storage, virtualization, and networking into a resilient, software-defined solution with rich machine intelligence.

# End-User Computing

for  
**dummies**<sup>®</sup>  
A Wiley Brand





# End-User Computing

Nutanix 3rd Special Edition

**by Scott D. Lowe**

for  
**dummies**<sup>®</sup>  
A Wiley Brand

# End-User Computing For Dummies®, Nutanix 3rd Special Edition

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

Copyright © 2023 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: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY 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 ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

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](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN 978-1-394-17027-2 (pbk); ISBN 978-1-394-17028-9 (ebk)

## Publisher's Acknowledgments

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

**Project Editor:** Elizabeth Kuball;  
Rachael Chilvers (third edition)

**Acquisitions Editor:** Ashley Coffey;  
Traci Martin (third edition)

**Editorial Manager:** Rev Mengle

**Business Development  
Representative:** Matt Cox

**Production Editor:**  
Tamilmani Varadharaj;  
Mohammed Zafar  
(third edition)

# Table of Contents

INTRODUCTION .....	1
About This Book .....	1
Foolish Assumptions .....	2
Icons Used in This Book .....	2
Beyond the Book .....	2
<b>CHAPTER 1: An Introduction to Virtual Desktops and Applications.....</b>	<b>3</b>
Understanding the History of End-User Computing.....	3
Jettisoning the mainframe for something smaller.....	4
Discovering the emergence of x86 dominance.....	4
Virtualizing all the things.....	5
Understanding how end-user computing has evolved.....	5
Understanding the Critical Components in an End-User Computing Environment .....	6
Comparing Desktop Architectural Options.....	8
Physical desktops and laptops .....	9
Application virtualization .....	10
Session virtualization.....	12
Virtual desktop infrastructure .....	14
Desktop as a service .....	15
Defining Organizational and End-User Computing Needs and Outcomes .....	16
Operational.....	16
Financial .....	16
Security.....	16
Scalability .....	17
<b>CHAPTER 2: Discovering Virtual Desktop Infrastructure .....</b>	<b>19</b>
Understanding Key VDI Concepts .....	19
Taking Stock of VDI Software Components.....	20
Implementing required components .....	21
Considering other software components .....	22
Sizing a VDI Environment .....	24
Determining the factors that affect host utilization .....	24
Optimizing the vCPU-to-pCPU ratio.....	26
Enabling availability .....	26

	Choosing the Right Virtual Desktop Instance Type .....	27
	Individually provisioned .....	27
	Full clones .....	27
	Linked or fast clones.....	27
	Implementing the Right Infrastructure Solution .....	28
	Choosing VDI storage .....	28
	Calculating storage capacity .....	29
	Considering the network .....	30
	Determining the need for GPUs as accelerators .....	30
	Understanding the Pitfalls of VDI .....	31
<b>CHAPTER 3:</b>	<b>Accelerating Desktop Deployments with Desktop as a Service</b> .....	<b>33</b>
	Embracing Cloud Concepts in a Desktop Paradigm .....	33
	Implementing a done-for-you desktop environment.....	34
	Applying pay-as-you-go cloud services to desktops: Shifting CapEx to OpEx.....	34
	Enabling desktop agility .....	35
	Choosing between Cloud-Based DaaS and On-Premises DaaS .....	35
	Discovering the Potential Pitfalls in DaaS.....	37
<b>CHAPTER 4:</b>	<b>Deploying Operating Systems and Applications</b> .....	<b>39</b>
	Creating and Managing Master Images.....	39
	Understanding the Role of Agent Software .....	41
	Discovering Application Catalogs and Delivery Mechanisms .....	41
<b>CHAPTER 5:</b>	<b>Managing the Dollars and Sense of End-User Computing Options</b> .....	<b>43</b>
	Comparing Economics: Traditional versus Server-Based versus VDI versus DaaS .....	43
	Painting a Complete Economic Picture.....	45
	Calculating the total cost of a traditional desktop environment.....	45
	Calculating the total cost of a server-based desktop environment .....	47
	Calculating the total cost of a VDI environment .....	47
	Calculating the total cost of a DaaS environment .....	52
	Calculating the total cost of DaaS in a hybrid cloud environment.....	53
<b>CHAPTER 6:</b>	<b>Ten EUC Factors to Remember</b> .....	<b>55</b>

# Introduction

**A**lthough the world is one-fifth of the way through the 21st century and information technology has dramatically changed, one paradigm has remained stubbornly intact: the need to provide workers with a robust and complete end-user computing environment. This environment enables them to interact with the workloads that you've deployed across the on-premises data center and in the public cloud.

Today's workforce wants to work from any device, in any location, securely, with a great user experience. However, organizations are facing the same problems when it comes to supporting users.

Maintaining a fleet of computers running native applications isn't ideal, especially when you think about installing new apps, managing updates, and provisioning users.

Don't forget security is a nonnegotiable, and letting users save data natively to their devices is dangerous, to say the least. Lastly, end users want the freedom to work from anywhere — on any device. This requirement presents several challenges for organizations.

The foundational needs of end-user computing — providing access to enterprise-provided applications and workloads — haven't changed much since the beginning, but the methods by which this goal is accomplished have transformed over the years. User demands have evolved, and enterprise capabilities have matured to support myriad new opportunities for advancing this resource.

Today, the industry is decades into virtualization's conquering of the data center and at least a decade into the hybrid cloud era. Enterprises can also affordably leverage flash storage and new data-center architectural options to transform and improve the methods by which end users consume enterprise resources.

## About This Book

In this book, I introduce you to the history of end-user computing. You'll discover that, although the key outcome desired from end-user computing environments has stayed consistent since

the 1950s, all the technology that underpins such environments has changed.

You'll discover the methods by which you can deploy application and workload services to your employees and develop a budget to make sure that you don't spend overspend.

## Foolish Assumptions

In this book, I assume you have at least a basic understanding of virtualization, storage, cloud, and end-user computing. The general audience for this book is anyone in IT who wants to learn more about the various end-user computing options that are available on the market. The audience is intended to be technical staff, as well as managerial and executive staff.

## Icons Used in This Book

Throughout this book, you'll find a number of icons intended to help you better identify key concepts.



REMEMBER

I use the Remember icon when you need to stop for a second and make sure you recall a key concept before forging ahead in a chapter.



TIP

You should keep certain details in mind as you analyze your own data-center environment. When you see the Tip icon, put that information in your back pocket to save for later.



TECHNICAL  
STUFF

Although I don't go super-deep into technical details in this book, I provide some technical elements for you in various places. You can find these marked with the Technical Stuff icon. If you're short on time, you can safely skip anything marked with this icon.

## Beyond the Book

There's only so much I can cover here. To learn even more about end-user computing needs, go to [www.nutanix.com/EUC](http://www.nutanix.com/EUC).

## IN THIS CHAPTER

- » Discovering the historic end-user journey and critical context
- » Understanding the components of an end-user computing environment
- » Comparing desktop virtualization options

# Chapter 1

# An Introduction to Virtual Desktops and Applications

**T**raditionally, end-user computing (EUC), which encompasses the management of virtual desktops and apps, has been dominated by on-premises solutions like virtual desktop infrastructure (VDI). But the cloud has ushered in the emergence of desktop as a service (DaaS), which allows you to enjoy a software-as-a-service (SaaS) consumption model. With that said, this chapter discusses how digital workspace strategy has evolved over time.

## Understanding the History of End-User Computing

Looking back through history reveals a modern EUC environment that was born in the mainframe era. Today's EUC environments serve the same basic needs as those archaic mainframe-centric environments. The critical goal is to enable application and workload consumption, which is exactly what the earliest hardwired "green screen" terminals provided so many decades ago.

What's more interesting is to consider the methods by which organizations have opted to provide workload and application access to their users over the years. Those long-gone mainframe days were the very definition of centralized computing. Workers were assigned to access terminals that were hardwired into the mainframe. Workers had little to no control or influence over their environment. They simply used the tools that they were provided. They didn't have a way to get their work done with other tools or at other locations.

## **Jettisoning the mainframe for something smaller**

As the mainframe era began its slow wind-down, it gave way to the minicomputer era. Minicomputers were essentially scaled-down mainframes that could be more easily distributed and connected with one another because they were smaller and didn't require entire rooms or floors of buildings. They still used terminal-based end-user access methods, but with a twist that I explain shortly.

## **Discovering the emergence of x86 dominance**

Overlapping with the minicomputer era is the age of what we can now term *x86 distributed computing*, defined by the emergence of the local area network (Ethernet or Token Ring), with companies such as Banyan and Novell leading the charge, but the winner of which was eventually Microsoft. These environments marked the first major foray into real distribution of EUC devices, with organizations of all stripes deploying desktop and, eventually, laptop computers to their users.

These overlapping eras began to expose some new opportunities that were less accessible before. As users were issued x86 desktops and laptops to replace aging green-screen and orange-screen terminals, they didn't lose access to those minicomputer and mainframe-based applications. In fact, they gained new levels of accessibility even as they were exposed to emerging local area network (LAN)-centric client-server applications. Through the use of software constructs that emulated the single-use terminals that the users had grown accustomed to, these users suddenly had a more full-featured and complete EUC experience. Plus, with the increase in use of laptops and the meteoric rise of home-based

Internet connectivity, users found that there was more desire to work on their terms.

But there were challenges. Allowing people to access work systems from home wasn't technically difficult, but there were security and licensing challenges from companies that were still stuck in the past. Still, so-called *teleworking* was not nearly as widespread as it has become today.

And then, the 2000s hit and everything changed.

## Virtualizing all the things

In the early 2000s, virtualization came to the stage and, although initial forays of this technology turned the data center upside down, it wasn't long until someone asked the question, "Hmm . . . do you think we could do this for desktops, too?"

The answer was a resounding, "Yes, but. . . ."

And VDI was born! What's VDI, you ask?

In the mid to late 2000s, the server virtualization → desktop virtualization concept came to market in the form of what became known as VDI. Very early iterations of VDI worked, but they were fraught with challenges, including massive deployment complexity in software, hardware that was barely capable of keeping up, issues around handling graphics and video, and software application vendors that decided to try to use VDI as a way to drive more revenue. This conglomeration of issues made VDI difficult to deploy. Stories about such failures abounded.

Eventually, the VDI software evolved and new hardware options — most notably, hyperconverged infrastructure — came to market with the promise of making VDI deployments much simpler, a promise that was mostly fulfilled, at least on the hardware front. That led to VDI environments that started to make financial and operational sense.

## Understanding how end-user computing has evolved

As you review the history, a couple of key points become clear. First, the *reasons* behind providing EUC environments have stayed

static. They exist to allow workers to access workloads and applications that power the business.



REMEMBER

Second, two paradigms have shifted dramatically:

- » **The breadth of application locality:** Applications used to run in one place. That's far from the case today.
- » **User demand for flexibility in device and location:** No longer are many people willing to be tied to their desks all day, every day.

As the business landscape has transformed with new technologies becoming available, the EUC landscape has transformed to match. Today, companies exist with a physical presence that's nothing but a post office box, while others have sprawling campuses and employees scattered around the globe, all of whom access applications that run in the data center as well as the cloud.

Behind these companies are armies of employees demanding the ability to use any device they want from anywhere in the world. From smartphones to high-end laptops and from middle America to Antarctica, employees want "anytime, anywhere" EUC capability.

Today, at the current stage in the evolution of this critical environment, it's possible to provide just that.

## Understanding the Critical Components in an End-User Computing Environment

As you look around a typical organization today, the physical EUC environment may not look that dissimilar from the ones that sprung up in the '80s and '90s. Look a little deeper, however, and you notice some radical differences.

The basics are similar between the two eras. In most organizations, you see some kind of box on a desktop with a keyboard, mouse, and screen, but what that device actually *is* may be very different. That device, though, is the user's entry point into your application and workload environment, and it's at the nexus of their day-to-day needs.

The devices the users use — the *end points* — to do their work are necessary components in an EUC environment. The end point is critical component number one and exists in some form in *every* EUC environment. For some environments, the end point is all there is. The desktop or laptop has all the client software necessary to access all the applications that the user has access to.

For larger, more complex, and more secure environments, the end point is just the beginning. Myriad systems operate behind the scenes to make sure that the EUC environment remains available, secure, and able to meet business and user needs.



REMEMBER

The list of components that exist in an EUC environment depends on the overall architecture of that environment. But some commonalities exist across many of the architectures. Table 1-1 provides a snapshot of the components to consider in support of EUC environments.

**TABLE 1-1** Components Necessary for a Complete EUC Environment

Component	Description
End-point device	The end-point device is the user's entry point into your application and workload environment.
Operating system	Every device needs an operating system to be deployed to it.
Asset management system	The term <i>asset management</i> can be used to describe a system that maintains a physical inventory, but in this context it can also refer to the overarching system that manages the desktop environment, such as Microsoft Endpoint Manager (MEM).
Application delivery mechanism	In just about every desktop computing environment, different desktops get different applications, depending on the role of the user. Plus, as new applications are deployed and as updates are released, some centralized tool is needed to deploy these services. Microsoft Endpoint Manager can do this, but a number of other solutions focus on providing this service as well.
Security services	Your environment needs to be inoculated against viruses and other malware. End points are generally imbued with antivirus software, firewalls, and other tools to make sure that they stay safe. For laptops and other portable devices, a virtual private network (VPN) service may be installed to protect network communication.

(continued)

**TABLE 1-1 (continued)**

Component	Description
Identity management and access authorization	Interconnected desktops need to maintain connectivity to centralized authentication services, such as Active Directory, to track who has access to the desktop and what they're allowed to do in the rest of the environment.
Network connectivity	A long time ago, disconnected desktop computers ran local applications. That hasn't been the case for a long time and network connectivity of some kind today is essential.
Availability services	People often think of traditional EUC environments based on desktops and laptops as something less than critical in the event of a disaster. Nothing could be farther from the truth.

These are common components in EUC environments. As you consider different EUC solutions, the components that you need to deploy may change dramatically. For example, as you deploy VDI, you need to add all the components necessary to support that environment, including brokers, gateways, and much more. Chapter 2 goes into greater depth about these components.

The decisions you make about the type of EUC environment you'll operate dictate the components that you need. Decisions can get complicated quickly, and you need to look at options that can reduce the complexity.

## Comparing Desktop Architectural Options

None of the options at your disposal when choosing a desktop architecture for your organization are mutually exclusive. You can deploy multiple architectures if you have a variety of needs, although the more you deploy, the more complexity you add to your organization. The following sections provide an overview of what's available on the market. Before jumping into deep explanations, look at Table 1-2, which provides a high-level overview of the options.

**TABLE 1-2 Comparing Desktop Virtualization Options**

Characteristic	Terminal Services	Application Virtualization	VDI	DaaS
Capital expenditures (CapEx)	Moderate	Low	High	Cloud: None On-premises: Moderate
Operational expenditures (OpEx)	Moderate	Moderate	Moderate to high	High
End-user experience	Moderate	Moderate	Varies	Very positive
IT operational impact	Moderate	Moderate	Very high	Very low
Completeness of solution	Moderate	Low	High	High

## Physical desktops and laptops

Physical desktops and laptops are, for many organizations, the default method by which users access the applications and workloads necessary for them to do their jobs. In such environments, applications are installed on the end-user device, and the system’s local central processing unit (CPU), random access memory (RAM), and storage are used for ongoing processing.

Physical desktops are inflexible in terms of providing users with the ability to take their work elsewhere. Plus, in the event of a disaster that includes the physical desktop environment, getting operational again can be eminently difficult and time-consuming.

Laptops go a long way toward solving the locale issue from which desktops suffer, but they have their own risks and require administrators to take intentional steps to reduce risk. Perhaps most critical is the potential for data leakage associated with laptop computers. If an unencrypted laptop is lost or stolen, this can spell bad news for the company. The rise of local encryption for desktop and laptop computers has reduced the risk, but even today, not every organization actively enables such functionality.

Of course, local encryption applies only to data at rest. Additional security concerns come into play as users connect to the local coffee shop's open Wi-Fi network to connect to your customer database and do their work. If they don't use a VPN, every key-stroke is potentially available for interception.

Beyond the security concerns is the administrative overhead associated with desktops and laptops. It starts as soon as a device is ordered. You have to either manually install all the tools that your company uses or implement some kind of imaging system to automate it. Every time something changes, you need to make sure to update all your devices, including the ones being used by, for example, your remote salesforce.



REMEMBER

Administratively — that is, on the business side — physical systems have to be included in some kind of ongoing replacement strategy. Typically, organizations turn over their desktop and laptop fleets every three to five years, meaning that the constant purchase → image → deploy → manage cycle never ends.

Between management cost and complexity, the desire for users to be able to work from anywhere they want and to use their own devices (see the sidebar “The rise of BYOD”), and the constant replacement cycle, the traditional desktop/laptop computing model began to face serious challenges. A number of new solutions were born.

## Application virtualization

For organizations that deploy applications right to the desktop, application compatibility can become a problem. Different versions of Windows or even different operating systems altogether — Linux, macOS, and so on — are all problems that administrators encounter and are exacerbated by things like BYOD. As consistent as administrators strive to make the desktop environment, achieving perfect uniformity is difficult, if not impossible.

*Application virtualization (or process virtualization)* is one method by which to address this problem. Under this paradigm, instead of running directly on an end point, an application runs inside an abstracted, but local, environment on that device. The abstracted environment contains enough of the underlying operating system, including libraries, the registry, executables, and the like, to run that application.

## THE RISE OF BYOD

A number of years ago, as end users became more tech savvy and as they increased their acquisition of personal devices, such as high-end laptops and tablets, many started to wonder why they couldn't just use their own devices instead of relying on the ones provided by their employers. This was the start of the business-centric bring-your-own-device (BYOD) movement. Some verticals, such as higher education, had been operating this way for students for quite some time, but the phenomenon eventually extended to their workforces as well.

At first glance, BYOD would seem to be just a replacement for the physical desktop device, but supporting this initiative raised a number of concerns. Foremost were concerns around data security. Because a user, not the employer, owned the device, what steps could reasonably be taken to lock a device down and encrypt local storage? Second, how would that user access applications? Would applications be installed locally, or would some other access method be required?

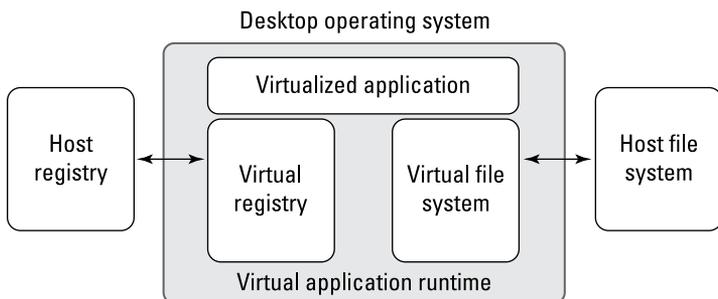
A number of concerns were ultimately resolved — or at least addressed — by the implementation of BYOD policies, but the technical access concerns remained. A number of the access methods described in the following sections were implemented in some companies in order to support BYOD and to allow users to enable what amounts to a ubiquitous computing model. In this model, users could access their authorized resources anytime, from anywhere, and from any device.

Consider how difficult it can be to upgrade to newer versions of Windows, particularly when you're running mission-critical software that hasn't been modified to work on newer releases of Windows. Most organizations want to stay as current as possible with Windows releases to ensure that security is as complete as possible and to get new features, but they need to be careful about leaving their old software unsupported.



An application virtualization solution, shown in Figure 1-1, can help these companies get the best of both worlds. To make this solution work, an administrator installs a clean version of Windows that supports the application target and is then installed while the application virtualization solution keeps a watch on how that application is interacting with the operating system. It

then pulls the necessary components and rolls them into a package that can be directly deployed to appropriate workstations or streamed from a centralized server. The benefit of streaming from a server is that the virtualization application package can be more easily updated as new versions are released.



**FIGURE 1-1:** An application virtualization architecture.

With application virtualization, a heterogeneous desktop environment appears homogenous to the application. This approach allows greater flexibility in the desktop environment for administrators.

## Session virtualization

You may not be familiar with the term *session virtualization*, but I'm willing to bet that you've heard of Microsoft Terminal Server, Remote Desktop Services, and the like. These are examples of session virtualization tools on the market; others are available from companies like Citrix and Parallels.

## NOTHING IS PERFECT

Not all applications can operate in an application virtualization environment and even for applications that *do* work, some functionality may not work as expected. In Windows application development, there are right ways and wrong ways to do this. If a developer chooses the wrong way to, for example, store application data, an application virtualization solution may have a difficult time fully supporting that application. Although I've had more successes than failures with application virtualization tools, the times I've had failures were generally with applications that did "weird" things.



REMEMBER

In this section, when I mention Remote Desktop Services, it's in the context of Remote Desktop Session Host (RDSH) mode and not a full VDI deployment. Microsoft's Remote Desktop Services is widely available, but other solutions are on the market, too. Citrix, for example, is synonymous with this technology and pioneered terminal services functionality in Windows.

Session virtualization is a completely different approach to some of the other solutions because it relies on application software running on centralized servers and being accessed by client computers. The only data that is transferred over the network is the screen redraw information that's necessary to show users what they're doing. The local client otherwise has no role in processing or storing data.

For these reasons, some organizations turned to session virtualization tools as a way to ensure that organizational data stayed in the data center and to democratize the desktop/laptop fleet and allow users to use any device they want. Most enterprise client-server applications are Windows-only affairs, but session virtualization means that any desktop operating system that has an appropriate client can consume these applications.

With tools such as Terminal Services and Remote Desktop Services (in session host mode), administrators only have to install an application once — on the server — and then all authorized users can get immediate remote access to that application.

Session virtualization is good, but it isn't always feasible. Some desktop applications don't run well on server operating systems, although this is less of a problem than it used to be. Enabling this capability on Windows Server can sometimes add some serious licensing headaches, too. Microsoft hasn't always catered to clients that really want to use Remote Desktop Services in their data centers.

Many organizations use session virtualization, but it isn't a fit for every application. Some applications won't run on a server operating system and, from a user experience perspective, they interact with the Windows Server operating system, which may not be what the user is accustomed to. The user doesn't have carte blanche access, though. You can still limit what they can do. In essence, every user shares the underlying server hardware, as well as the operating system. Individual sessions are isolated from one another, but everything else is shared.

One additional complicating factor with session virtualization is that it requires always-on network connectivity to work, limiting its usefulness in low-bandwidth or no-bandwidth scenarios.



TIP

When it comes to data security and the ability to quickly deploy an application to many users, session virtualization is often turned to as a viable option. All the data that the user interacts with during a Terminal Services session stays on the server. Only screen updates are sent to the client. This practice reduces the potential for data leakage and eliminates risk if a user's end point is lost or stolen. In terms of end-point device flexibility, session virtualization is widely supported, with clients available for just about any device you can image. For Windows systems, as long as the end-point software client can support the Remote Desktop Protocol, you'll have no trouble connecting to a server environment.

In a downside for the user experience, because the user is only sharing a slice of the underlying server operating system, the environment isn't as customizable as one in which the user gets to "own" the entire operating system. Plus, if you ever need to allow users to install their own software, this can be difficult in a session virtualization environment and you'll need other options.

## Virtual desktop infrastructure

One of those other options is the next item up for discussion. After server virtualization took hold, x86 virtualization was in full swing as administrators sought to virtualize everything in their path. At some point, some people started wondering if the desktop environment could be virtualized as well.

The answer was a resounding, "Yes!," and the results were often a resounding "thud." However, the initial roadblocks to success in VDI were ultimately removed — thanks in large part to the advent of flash storage and the introduction of hyperconverged infrastructure — which, together, removed much of the hardware complexity inherent in the solution.

In a VDI environment, users interact with a complete Windows desktop environment, dedicated to them, depending on the configuration. From an interface perspective, VDI doesn't look any different from a physical desktop interface, so the transition is smooth and application compatibility issues largely go away.

However, the complexity in the underlying VDI software has remained. That's probably the biggest downside to VDI. Getting it absolutely right can be difficult, even with the easiest hardware available.



Even though VDI looks a lot like server virtualization, there are some critical nuances. First, whereas server workloads are generally very different from one another, desktop workloads are generally very uniform. This may sound like a positive, but it isn't, at least from a performance perspective. VDI administrators need to be careful about how many desktops they load into a single host server, and they need to keep peak workloads in mind. VDI gave rise to the terms *boot storm* and *login storm* because these activities caused VDI hosts to buckle and there were stories about users waiting minutes for login processes to complete. These problems don't exist, for the most part, in traditional desktop environments.

At inception, VDI was looked at as a potential cost-saving measure, too. After all, if you could just buy terminals rather than complete desktops, that had to save money, right? Those early hopes were quickly dashed as any potential savings were eaten up by the need for new storage and the software licensing associated with VDI. (You can learn all about VDI in Chapter 2.)

## Desktop as a service

With the rise of the cloud also came the rise of a complete spectrum of “as-a-service” offerings. These welcome additions to the enterprise IT landscape have made it possible to very quickly and very easily deploy most new services without a lot of heavy lifting needed from IT.

The most recent addition to the EUC landscape comes thanks to the cloud. DaaS was born in the cloud but is also starting to make its way into on-premises data centers.



Where DaaS shines is in the elimination of the complexity overhead associated with VDI. It's a kind of desktop vending machine — you insert coins, and you get desktops. With DaaS, you truly buy ready-to-run desktops and layer your applications atop them. All the server-side hardware and software is accounted for in a monthly fee that you pay to the provider.

As is the case with VDI, all you need is a remote connection client to access the virtual desktop, and you're on your way to productivity. (You can learn much more about DaaS in Chapter 3.)

## Defining Organizational and End-User Computing Needs and Outcomes

This chapter wraps up with a strategic look at what decision makers desire from an EUC environment. There are four areas of interest.

### Operational

This is the most fundamental need. Decision makers want and need a desktop environment that allows the organization to operate. It's that simple. Anything that gets in the way of that goal is a nonstarter.

The user experience is part of this area. Modern workers demand flexibility in their workplace. They want to use their own devices to work from any time and from any place. Enterprises need to accommodate these demands with tools that provide necessary levels of flexibility.

### Financial

That doesn't mean that an enterprise wants to spend millions to provide this flexibility, though. The EUC environment is still an expensive budget item. Chief financial officers (CFOs) are willing to spend what's necessary to optimize worker productivity, but no more. So, the EUC environment must be cost-effective, and CFOs increasingly want to operationalize this expense.

### Security

As a boardroom issue, security needs permeate all aspects of the organization, including the EUC environment. Whatever is deployed must be safe and secure and not increase the organization's exposure to data theft or leakage.

## **Scalability**

No one wants to buy something they can't grow into. This holds true for the EUC environment. As new employees are brought on board, the environment must be able to grow as the organization grows. Scalability should be inherent in the solution, and such scalability should be a core feature that is handled simply and elegantly.

- » Understanding key VDI concepts
- » Identifying the factors that affect VDI performance
- » Sizing VDI hardware resources

# Chapter 2

## Discovering Virtual Desktop Infrastructure

Virtual desktop infrastructure (VDI) isn't a new concept at all, but it's one of those services that has just enough cost and complexity that it still confounds people who don't always look at it for the right reasons. It used to be seen as a cost-savings opportunity. Thanks to hardware complexity at the time and the introduction of software specific to supporting VDI workloads, direct cost savings didn't always emerge.

Instead, VDI enables other outcomes, including providing inherent flexibility for users and, in many cases, increasing the overall security posture of the organization.

### Understanding Key VDI Concepts

The two most popular VDI solutions come from Citrix and VMware. Citrix was a pioneer in server-based computing, so it's no surprise that it continues to lead with Citrix Desktop as a Service (DaaS). VMware found great success in server virtualization and translated this to the desktop world with VMware Horizon. Table 2-1 outlines the main components present in both Citrix and VMware VDI environments.

**TABLE 2-1 Citrix versus VMware VDI Components Overview**

Component	Description	Citrix	VMware
Administration	The tool that administrators use to configure and maintain the environment	Citrix Management Interface: Quick Deploy or Full Configuration	VMware Horizon 7 Administrator Console
Hypervisor	The abstraction layer that allows virtual machines (VMs) to share host hardware	AHV Hyper-V vSphere XenServer	vSphere
Provisioning	Software that allows administrators to create new virtual desktops for users	Citrix PVS and MCS	VMware View Composer
Connection broker	A tool that directs inbound desktop access requests to authorized resources	Delivery Controller	VMware Horizon Connection Server
User portal	The front end for the user experience	Citrix WorkSpace	VMware Identity Manager User Portal
Connection protocols	The protocol that client devices use to connect to hosted desktops	RDP HDX	RDP VMware Blast Extreme PCoIP

## Taking Stock of VDI Software Components

The biggest downside of VDI is the complexity. It doesn't *have* to be too insanely complex, but to get things just right requires alignment of a lot of moving parts.

## Implementing required components

There are some components that you simply can't do without in a VDI scenario. They're required and, without them, you'll be left with little more than a pile of parts and no functionality.

### Choosing a hypervisor

It should come as no surprise that, in a solution that's based on a virtualized operating system, a hypervisor is a required component. Some VDI solutions support a variety of hypervisors; others don't. For example, if you choose Citrix for VDI, it supports VM creation in a variety of environments, including the following:

- » Citrix XenServer
- » Microsoft Hyper-V
- » Nutanix AHV
- » VMware vSphere

VMware Horizon, on the other hand, unsurprisingly supports just vSphere. View includes vSphere for Desktops, which is rebadged vSphere Enterprise Plus.

Your choice comes down to making sure that the combination of the hypervisor and the VDI software you choose has all the features you need to make your VDI project a success.

### Selecting client software/agents

Technically, you don't need anything special to connect to a desktop VM hosted in your data center. You can simply enable remote access to that desktop and use any one of a variety of remote desktop clients to connect to that desktop. However, that leaves a lot to be desired in terms of user experience, as well as operations.

I cover connection protocols in this chapter under "Considering Other Software Components." For now, bear in mind that the Remote Desktop Protocol (RDP) is good enough for high-level use, but it isn't great as you start to push desktop a little harder, particularly with graphics, video, and audio.

The client software or agent you'll use is determined by which vendor's VDI stack you choose. If you go with VMware Horizon, you'll use the VMware Horizon Client, which is available for a wide

variety of platforms, including Windows, macOS, iOS, Android, Linux, and Chrome.

In Citrix, you'll use the Citrix Workspace app (formerly Citrix Receiver). The Workspace app is also available for Windows, macOS, iOS, Android, Linux, and Chrome.

Plus, each VDI solution has an agent that is installed inside every managed virtual desktop, physical system, or RDS host to allow it to be managed as a part of the VDI environment overall. The agent handles tasks such as printing, USB support, and more.

## Understanding the role of connection broker software

Technically, you don't need anything except a simple client and RDP to connect to a Windows desktop running on a virtual host. However, for organizations of any size, the use of a connection broker to handle such activity is all but required.



TECHNICAL  
STUFF

A connection broker helps a user's VDI client connect that user to an appropriate VM instance. The broker handles the task of authenticating users as they try to connect. The broker also tracks which desktops a user should have access to.

The connection broker may also make associations between assigned virtualized applications and a logged-in user. The broker handles responsibility for ensuring that any user policies are applied at login.

In a Citrix environment, this service is handled by the Delivery Controller component. In a VMware Horizon environment, brokering is handled by the Horizon Connection Server.

## Considering other software components

The previous section discusses components that are, more or less, required in a VDI deployment. This section covers some additional components that you should consider.

### Connection protocols

Connection protocols have a direct impact on the user experience, central processing unit (CPU) consumption on the host and end

point, and the battery life of mobile end points, such as smartphones and tablets. If you're on a super-low-latency and ultra-fast connection, you may be able to use a protocol that doesn't need to compress as much information. If, on the other hand, you're in a bandwidth-constrained environment, you may need a more efficient protocol.

Bear in mind that the more bandwidth-efficient the protocol, the more CPU or graphics processing unit (GPU) cycles it will likely require at the host to compress information. There is always a trade-off.



TECHNICAL  
STUFF

A number of protocols are available to you, depending on which VDI solution you choose:

- » **RDP:** Assuming that most VDI environments are based on Windows, RDP is the default protocol, and it's available regardless of whether you go with Citrix or VMware. It's the lowest common denominator, but it's still good for many users.
- » **PCoIP:** Developed by Teradici and licensed by VMware, PC-over-IP was the solution to solve the need for an efficient protocol that also provided a closer-to-native experience for end users. PCoIP provides an encrypted experience with granular bandwidth controls, support for 32-bit color, audio redirection, real-time audio/video, multiple monitor support, shared clipboard, and USB redirection. PCoIP is a lossless protocol that perfectly replicates the desktop environment.
- » **VMware Blast Extreme:** In Horizon 7, VMware introduced Blast Extreme, an H.264-based protocol that requires less bandwidth than PCoIP, making it more suitable for mobile devices. Blast Extreme works well and can leverage servers that have been equipped with GPUs for 3D work. Blast Extreme is based on a video codec that includes compression, so the image displayed on the end point may not be a perfect match for what's on the host. It provides a number of benefits, including bandwidth friendliness, GPU support, lower CPU consumption, multi-monitor support, and hardware decoding.

» **Citrix HDX:** Citrix has done well over the years with HDX, its high-fidelity, bandwidth-friendly VDI client. HDX, like PCoIP, aims to replicate, as closely as possible, a native-quality user experience in a virtualized environment. HDX also provides multi-monitor support, as well as support for 3D graphics and GPU acceleration.

## Security servers and gateways

Not every user will connect to your VDI environment from within your corporate network, and they may not be using a company-issued laptop. For them, you may consider the deployment of security servers and gateways intended to provide access to external clients.

In VMware Horizon, this service is provided by the VMware Unified Access gateway, which handles direction of authentication requests from external clients, as well as direction of clients to connection servers.

For Citrix VDI deployments, Citrix Gateway handles these services and provides security and access control policies for end points trying to connect to the environment.

## Sizing a VDI Environment

Perhaps one of the biggest challenges in a VDI scenario is sizing it appropriately. Early forays into VDI often failed due to either a miscalculation of certain resource needs or a failure to fully understand all the variables involved.



TIP

Nutanix makes available a sizing tool that helps you determine the number of nodes, the power draw, amount of rack space, and more, for your VDI deployment. If you'd like to try it for yourself, visit [www.nutanix.com/size-your-datacenter](http://www.nutanix.com/size-your-datacenter).

## Determining the factors that affect host utilization

A number of factors and variables affect how much capacity you need in terms of CPU, random access memory (RAM), storage capacity, storage performance, and network bandwidth. The goal is to optimize the desktop density per host.

## Concurrent desktop instances

The number of desktops that you need to run simultaneously has the most significant impact on your host utilization. This number may differ from the total number of virtual desktops that you need. For example, if your company runs in three shifts and has 3,000 employees, each with a dedicated desktop, you don't have to size the environment to support 3,000 simultaneous desktops. You need to size it to support the maximum number of desktops that will run at the same time.

## Size of desktops

Sizing virtual desktops is an activity that asks you to specify the number of virtual CPUs (vCPUs) that each desktop will have, the amount of RAM, and the amount of disk space.



But the aggregate math isn't quite as simple. For example, if every desktop will be configured with 80GB of disk capacity, you can't assume that 3,000 desktops will require 240,000GB (240TB) of capacity. Technologies such as deduplication, thin provisioning, and linked clones (explained later) may allow you to dramatically reduce your storage capacity need. The only way to get close is to do a small pilot and see what happens.

The same goes for RAM. Some hypervisors do tricks with RAM to reduce RAM consumption and increase VM density. Again, you need to experiment.

## Application types

The kinds of applications you run are critical. If you're only supporting email clients, you can increase the number of desktops per host. Desktops with more-demanding applications also require more resources, which reduce the overall density of the VDI environment.

## Graphics needs

GPUs are becoming all the rage, and they're popular in VDI environments. GPUs can allow hosts to offload some of the graphics processing needs (which can be significant) to the GPU, thereby increasing the density of the desktops on the host. This solution only works to a point. Eventually, the GPU itself gets saturated, so you need to watch overall performance.



TIP

The beauty of a VDI environment is that adding hosts is a simple matter if you begin experiencing performance problems. You can just add a node and reduce the density of the existing hosts to free up capacity.

## Optimizing the vCPU-to-pCPU ratio

With server workloads, you have to be careful about oversubscribing physical CPU cores to a point where you have CPU ready/CPU availability concerns. The same holds true in VDI environments. If you add too many virtual CPUs per physical core, your desktop environment eventually suffers as the host struggles to make physical resources available.

In most scenarios, except those in which organizations need to deploy monster desktop VMs, you should be able to get a higher vCPU-to-physical CPU (pCPU) ratio than you can with server workloads. General guidance suggests that a 6:1 to 10:1 ratio is safe. You should be able to plan on using six to ten virtual CPUs for every core in your system.

If you have 500 dual core desktops that have to run concurrently, you need 1,000 cores of capacity. Assuming that each physical processor has ten cores and you have two cores per host, you can support about 100 desktops per host and should plan for about five hosts. Of course, this estimate is general and doesn't consider leaving overhead for spikes and availability.

## Enabling availability

Your VDI environment is likely more critical than many people give it credit for, and it needs some level of availability. The first step is to design your VDI architecture with an “N+1” mindset. Make sure that it's built so that it can support 100 percent of your desktop environment even if a single host fails.

From there, consider the capabilities of the VDI architecture. Hypervisors often have built-in availability mechanisms that can be leveraged by the VDI deployment. For example, if you're using VMware Horizon, you can use vSphere's high-availability feature to ensure automatic failover for particularly critical desktops.

In some environments, failure is truly not an option. Consider health care. In such environments, look for additional availability guidance from your VDI vendor. For example, VMware Horizon

provides what the company calls an AlwaysOn architecture for these types of environments and makes available comprehensive guidance for all aspects of the environment.

## Choosing the Right Virtual Desktop Instance Type

After you've deployed the underlying hardware and the overarching VDI software, it's time to start creating virtual desktops! You have several options.

### Individually provisioned

Just as you can individually deploy physical desktops, you can individually deploy virtual desktops. This approach is feasible only for small deployments or in deployments where every virtual desktop will be unique.

### Full clones

The more common way that virtual desktops are deployed is by cloning an existing virtual desktop that is created as a master image or *parent VM*. The first cloning method is a full clone. Under this instance type, when you clone the master image, an exact replica of that image is created, completely detached from the master image. The cloned desktop stands fully separate. It's a way to get a copy of your master image up and running, though each one may take a few minutes to create.

### Linked or fast clones

The next cloning method is called either *linked cloning* (VMware) or *fast cloning* (Citrix). These types of clones share virtual disks with the parent, with the result being much less need for disk capacity. After all, the entire operating system in a virtual desktop is replicated, so with linked clones, you get to jettison that capacity overhead and just use the parent VM's virtual disks. Because there's no need to replicate all the storage contents, linked clones are created quickly.

One downside of a linked/fast clone is a bit of a performance dip because all child VMs share the disks of the parent. Plus, if the parent VM becomes unavailable for some reason, none of the linked VMs can operate.



REMEMBER

Linked clones keep a *delta disk* available that holds all the changes and additions that have been made to the clone since its creation. Bear in mind that the delta disk will grow in size over time and may grow up to the size of the parent's disk.

## Implementing the Right Infrastructure Solution

Hardware in a VDI world is important, and there are decision points all along the way.

### Choosing VDI storage

All kinds of storage architectures are available on the market, and any of them is suitable for use with VDI, provided it has the performance characteristics to support your VDI needs. You can choose from hybrid storage arrays, all-flash storage arrays, and hyperconverged infrastructure. For some VDI systems, even cloud storage may be an option.

Hyperconverged infrastructure remains a top choice today for supporting VDI. VDI is complex, and hyperconverged infrastructure's hardware simplicity can bring simplicity to your VDI deployment. As you need to add more nodes, you just add compute, RAM, and storage all at the same time, and you don't need to worry about downtime.

The key guidance on storage, besides capacity, is to make sure it can support the input/output operations per second (IOPS) load that your VDI environments will impose. Modern flash-centric architectures have largely eliminated early IOPS challenges associated with VDI, but especially in larger deployments, make sure that the storage solution you choose can handle even the most intense boot and login storms.

Table 2-2 provides an overview of your common storage options.

**TABLE 2-2 VDI Storage Options Compared**

Factor	Direct/Local	NAS/SAN	HCI
Cost	Low	Medium to high	Medium
Complexity	Low	Medium	Low
Performance	High	Medium to high	Medium to high
Scalability	Low	High	Medium to high

## Calculating storage capacity

Calculating storage capacity requirements also has some variance that you need to understand. If you use linked clones rather than full clones, the amount of capacity you need is cut dramatically. And if you're using storage that provides deduplication and compression services, you'll find that VDI workloads tend to reduce by high ratios thanks to the uniformity of the environment. In some cases, you may see a 50 percent to 90 percent capacity reduction.



TECHNICAL  
STUFF

Keep in mind the RAID level of the storage you intend to use. For example, if you plan to deploy VDI into an environment with a RAID 1 configuration, you instantly lose 50 percent of your raw storage capacity (and take a performance hit — read this: [www.techrepublic.com/blog/the-enterprise-cloud/calculate-iops-in-a-storage-array](http://www.techrepublic.com/blog/the-enterprise-cloud/calculate-iops-in-a-storage-array)). In this section, the capacity figures refer to post-RAID usable capacity.

As you calculate storage capacity needs for your VDI cluster, consider these variables:

- » **The cloning method:** Full clones need disk space immediately whereas linked clones need only a fraction of the raw capacity. Capacity needs will grow over time.
- » **The RAID level in use:** I mention this in the preceding Technical Stuff paragraph. Don't forget it!
- » **The storage's data reduction capabilities:** VDI workloads reduce very well, so if your storage includes data reduction, plan to see a high level of efficiency.

» **Ongoing growth needs for linked clones:** Linked clones' storage capacity needs grow over time as they drift further from the baseline configuration of their parent VM. VMware provides formulas (also applicable to Citrix) to help you estimate the storage capacity (<https://docs.vmware.com/en/VMware-Horizon/index.html>).

## Considering the network

The network plays a critical role in a VDI environment. All the VDI software components are interconnected over network and all your users access their desktop environment from devices that rely on network connectivity to function.

As in everything with IT, your main concern with networking should be latency. Latency is especially challenging in a VDI environment because its impact is immediately visible.



REMEMBER

Different kinds of workloads require different levels of bandwidth. For example, according to VMware, a typical connection for a knowledge worker consumes 150K to 200K of continuous bandwidth. Those who work on video or intensive graphics may need as much as 2Mbps. Now, multiply that by the number of users in your environments.

## Determining the need for GPUs as accelerators

GPUs have recently emerged as processing powerhouses thanks to the number of processing cores packed into each unit. With VDI, you don't have to use GPUs in your architecture, but doing so may offer benefits. Modern productivity apps such as Chrome browser, Outlook, and O365 look to GPU availability first and, failing to find a GPU, could potentially over-utilize the available CPU.

First, because GPUs offload the CPU, additional compression can take place without affecting other workloads. More data can get pushed over the wire to end points, thereby potentially improving the user experience and getting closer to full physical desktop fidelity.

This offloading operation also means that you may be able to increase the vCPU-to-pCPU ratio because the server's physical CPUs are relieved from some intensive operations. This capability can have a positive impact on desktop density per server.

# ENABLING UBIQUITOUS COMPUTING

As your users demand access to more and more devices, VDI helps you answer the call and implement what amounts to a device-agnostic environment — anything works from anywhere, at any time of day. You can place limits on what people are allowed to do, but your users are provided with a robust and secure method by which to access the applications that they rely on to get their work done.

GPUs also affect the kinds of applications you can run. With GPUs in your servers, you can unlock the potential in 3D applications, which is a boon for organizations with high-end needs that also want to leverage VDI, such as engineering firms. The same goes for colleges and universities that want to support a wide array of applications, including computer-aided design (CAD) and engineering tools.



TECHNICAL  
STUFF

Using specialized hardware in virtual environments used to be difficult, but thanks to services such as single-root I/O virtualization (SR-IOV), VMs can directly access the raw power of GPUs and other hardware installed in servers.

Whether you need GPUs in your VDI cluster depends on what you're running and the user experience you want to deliver. The common protocols described in this book can leverage GPUs to increase performance and quality.

GPUs also add complexity. Given the already complex nature of VDI, make sure you want to assume the added burden. On the hardware front, if you choose to go the hyperconverged route with something like Nutanix, the hardware will arrive ready to go.

## Understanding the Pitfalls of VDI

VDI isn't for everyone. Perhaps the most important shortcoming is one of skills and dedication. If you don't have someone on your team with the right skills who can dedicate time to make a VDI

deployment work, you're better off without one. That's the big takeaway, but you should also understand some other potential pitfalls:

- » **Not setting the right success metrics:** A VDI project will fail if you don't put in place achievable metrics for success. Perhaps the most common metric that likely won't materialize is a significant cost savings.
- » **Huge capital outlay:** A VDI environment may require a substantial outlay to get underway.
- » **Complexity:** VDI can be a complex undertaking. VDI is constantly evolving, and the documentation is evolving along with it. This technology has so many variables that it's difficult to dial into what you need until you're already underway on a project.
- » **Too much variance:** If you have too much variance in your desktop environment, VDI is doable, but it likely won't save you much in the way of operational efficiencies. Your users may enjoy the additional flexibility but you may find better ways to achieve it.
- » **VDI → bring your own device (BYOD) → new headaches:** Many BYOD policies have VDI at the core, which may require you to revamp your policies around such activities.

## IN THIS CHAPTER

- » Discovering the intersection of the cloud and the desktop environment
- » Understanding the difference between cloud DaaS and on-premises DaaS
- » Avoiding DaaS pitfalls

# Chapter 3

# Accelerating Desktop Deployments with Desktop as a Service

Where virtual desktop infrastructure (VDI) imposes complexity on an organization, emerging desktop-as-a-service (DaaS) environments remove it. DaaS is the next iteration of VDI. Although it doesn't currently replace VDI, as DaaS evolves, it may be able to supplant VDI in the not-too-distant future.

## Embracing Cloud Concepts in a Desktop Paradigm

The cloud has given rise to instant-on, done-for-you software-as-a-service (SaaS) tools. These SaaS tools are super-easy to adopt, requiring only a credit card to get started.

What if you could apply this methodology to desktops, eliminate the complexity inherent in VDI, and get the same kinds of outcomes you expect from VDI?

## Implementing a done-for-you desktop environment

That's where DaaS comes in. Born in the cloud, DaaS solutions require only that you log in to the DaaS service, provide a credit card number, and for some providers, download a client.

Getting a VDI environment fully operational can take weeks or months. With a DaaS environment, it can happen in minutes thanks to the cloud-centric nature of the service.

You should consider desktops in the cloud for several reasons:

- » **No more local VDI:** You don't need to deal with all those moving parts.
- » **On demand:** Hiring a new employee? Their desktop is provisioned and they can use the laptop that they brought with them.
- » **Pay as you go:** Pay for just what you're using and nothing more.
- » **Reduced capital expenditures (CapEx):** You don't need to buy stuff up front to make this happen.
- » **A strategic focus:** Instead of spending your time on tactical desktop operations, you can focus on company strategy.

This section dives into two of the most important elements.

## Applying pay-as-you-go cloud services to desktops: Shifting CapEx to OpEx

When possible, organizations prefer to avoid expending large amounts of cash all at once, particularly for services that are consumed at different rates or may not be 100 percent utilized. This is true for traditional approaches to IT. When IT decision makers embark on the next replacement cycle, they try to predict what usage they'll experience in three to five years.

Rarely do these estimates end up entirely correct. To avoid running out of capacity at mid-cycle, organizations often over-provision environments, thus leading to overspending.

Decision makers have turned to the cloud for help. Doing so enables organizations to pay for what they use and not a penny more. Whether the cloud is cheaper is an open debate, but it enables an operational expenditures (OpEx)-centric method for financing. Traditional infrastructure is challenged in this regard.

The desktop environment is in need of financial attention on a regular basis. It's also an area that every employee sees every day. Employees often clamor for different kinds of devices, work-from-home scenarios, and the ability to work more while traveling.

Organizations have turned to VDI to address these needs, though they often see VDI as cumbersome and expensive. VDI often requires capital outlay.

DaaS morphs a capital-intensive service into one that merely sips operational expenses and only requires a financial spigot big enough to match current needs.

## Enabling desktop agility

*Bring your own device* (BYOD) was formerly a buzz-term that described a trend in which employers were expected to allow employees to use their own phones, tablets, and laptops.



REMEMBER

Although it isn't in the news as much anymore, BYOD is alive and well. Organizations are working to allow it, creating policy frameworks to authorize it, and implementing technology frameworks to enable it.

Like VDI, DaaS brings inherent agility to the desktop paradigm, allowing users to work from anywhere and with any device that can provide an HTML5 connection.

## Choosing between Cloud-Based DaaS and On-Premises DaaS

When you think of “as a service,” you likely think of “cloud.” That's generally been true. DaaS is a born-in-the-cloud service, but that doesn't mean that it's *restricted* to the public cloud. You can operate DaaS using an on-premises deployment instead. For example, you may work for an organization that requires

end-to-end control of the computing experience for your users because of security or compliance concerns, or you may work for a company that prefers to house its own equipment.

As you consider on-premises DaaS versus cloud DaaS, what differences emerge? Table 3-1 has the answers.

**TABLE 3-1 On-Premises versus Cloud DaaS**

Characteristic	On-Premises DaaS	Cloud DaaS
Who supports it?	Provider or IT	Provider
Cost model	Pay monthly per desktop	Pay monthly per desktop
Hardware purchase needed	Underlying networking	None
Scaling latency	Moderate (new hardware must be shipped)	None
Reliability	Dependent on your environment	Guaranteed via a service-level agreement (SLA)

From a user perspective and from the perspective of the person in IT who manages the DaaS environment, an on-premises DaaS solution operates identically or very close to the cloud solution. Both use the same management interface and tools.

However, the on-premises DaaS solution may have challenges to overcome. Most important, your IT department will manage some aspects of the environment that typically would be handled by the service provider. Even though the service provider likely will provide support for the service, IT will have to ensure that power, networking, cooling, and other underlying data-center services are provided to the DaaS servers. Even if the provider sends an all-in-one cabinet that contains all the networking needed for the solution, IT still has to connect that to the network. The process involves making sure there is appropriate IP addressing, firewalling, and more.



TIP

When the time comes to add capacity to the DaaS environment, you need to wait for the provider to ship you new hardware. However, it's more than likely that the provider implemented some level of capacity management, so they may be able to provide hardware before you even know you need it. I recommend that you work with the provider to see what they can do for you.

When you buy cloud-based DaaS, you're essentially buying a service that's ready to go as soon as you provide a credit card number. With an on-premises DaaS environment, a little more work will be involved for IT.



REMEMBER

There's one huge *but* here. With either on-premises DaaS or cloud DaaS, you don't need an IT skill set that involves VDI. The provider will handle all aspects of that for you.

## Discovering the Potential Pitfalls in DaaS

With cloud-based DaaS, you don't get as much control over the environment as you do with VDI. With VDI, you get to control everything, including storage. With DaaS, you're at the mercy of the provider. For some, this is a pitfall, but for others, it may be a benefit. For those who refuse to store confidential information in the cloud, DaaS will not be an option for now unless they choose something like a hybrid DaaS model and don't run cloud-based desktops.

As is the case with VDI versus traditional desktops, DaaS may not save money. It will reduce capital expenditures in favor of operational ones, but the total sum for DaaS over time may be the same as doing VDI. As you embark on your desktop modernization journey, make sure you look at the total cost of ownership and don't forget to consider staffing costs in the VDI portion.

### WE WROTE THE BOOK ON DAAS!

If you want to learn much, much more about DaaS, jump over to [www.nutanix.com/content/nutanix/en/go/daas-for-dummies](http://www.nutanix.com/content/nutanix/en/go/daas-for-dummies) to download your very own copy of *Desktop as a Service (DaaS) For Dummies*, compliments of Nutanix.

## IN THIS CHAPTER

- » Creating and managing master desktop images
- » Understanding the role of agent software in a VDI environment
- » Discovering application catalogs and delivery mechanisms

# Chapter **4**

# Deploying Operating Systems and Applications

**N**o matter how much you might want to, you can't get away from operating systems in your desktop environments. Whether you're deploying physical desktops, virtual desktop infrastructure (VDI)-based desktops, or you've made the jump to desktop as a service (DaaS), you still need an operating system and application deployment plan in place.

## Creating and Managing Master Images

A long time ago, there was Ghost. No, not the movie. It was one of the original methods by which organizations cloned desktops to avoid having to manually install Windows, Office, and a slew of other applications manually on every new PC that rolled into the environment.

Ghost is still around for physical machines, but for VDI environments, the software-centric nature of a virtual machine (VM) enables some new opportunities.

With every VDI solution come tools to help administrators create and manage master images for each type of virtual desktop in their VDI environment.

For VMware, that's the job of either Horizon Composer or Horizon Enterprise Instant Clones. If you're using Instant Clones only, you don't need to deploy a Composer server or database because everything is handled via the Instant Clones tool.



TECHNICAL  
STUFF

In Citrix, two tools are used for image management:

- » **Machine Creation Services (MCS):** MCS is managed via the Citrix Studio and builds linked clones from a master image that you create. Each clone includes a differencing disk, which is akin to the delta disk used in a VMware Horizon deployment. MCS also supports the creation of full clones for times when you don't want to or can't use a differencing disk.
- » **Citrix Provisioning (PVS):** PVS is a streaming technology that streams desktop images to supported devices instead of simply copying images to them. PVS is highly scalable and reduces the overall storage footprint needed for disk images. You can stream PVS images to virtual desktops and even to physical desktops, which is how you can get away with not having to create images for all the desktops in your environment.

MCS supports dedicated and pooled virtual desktops, as well as desktops hosted in Amazon Web Services (AWS) and Microsoft Azure. PVS supports streaming to both virtual desktops and physical desktop devices but not to cloud-based virtual desktops.



REMEMBER

A lot of additional nuance surrounding the MCS versus PVS issue is beyond the scope of this book. If you're deploying Citrix VDI, make sure to read Citrix's detailed documentation on this topic to make sure you know all the use cases for both technologies.

Regardless of which VDI solution you choose, the first step in your deployment process consists of creating a master, golden, or parent image that includes agent software.

# Understanding the Role of Agent Software

What exactly does agent software do in a VDI architecture? Begin with a look at VMware Horizon. In that environment, the Horizon Agent must be installed inside each of your virtual desktops to allow that desktop to communicate with vCenter Server and the Horizon Connection server.

In VMware, the agent installs a number of services, including communications protocols such as VMware Blast, VMware Tools, Snapshot Provider (for cloning), and USB redirection services.

In Citrix, the agent is known as a Virtual Delivery Agent (VDA) and is installed on each physical or virtual machine that delivers desktops or applications. VDAs manage connections between that machine and user devices and ensure that policies are also deployed.



REMEMBER

Agent software is a critical part of your VDI environment. Without it, your environment won't be complete.

## Discovering Application Catalogs and Delivery Mechanisms

Application catalogs and automated self-service delivery aren't unique to VDI and DaaS environments. These kinds of mechanisms have been around forever in tools such as Microsoft Endpoint Configuration Manager (MEM). The purpose of these capabilities is to streamline the delivery of applications and reduce calls to the IT help desk requesting installation of a specialized piece of software. As the world becomes much more oriented to on-demand self-service, expect users to demand these services in your organization.



REMEMBER

The fact is that you don't need VDI or DaaS to create an application catalog. You can use stand-alone tools such as VMware Horizon Apps to publish applications and deliver them using session virtualization technologies such as Microsoft RDS. When deployed in this way, the delivery mechanism is imbued with the VMware

Blast Extreme protocol to smooth out the user experience, which isn't always the greatest with native RDP.

VMware makes available App Volumes, a complete application and user management solution that provides application delivery, isolation, and monitoring for VMware Horizon, Citrix VDI, and Remote Desktop Session Host (RDSH) environments, depending on the product edition you select.

App Volumes allows IT to very quickly and efficiently create new virtual applications and attach them to virtual desktops and physical desktops by deploying them as VMDKs or VHDs.

Finally, VMware also makes available its JMP (pronounced jump) platform, which stands for Just-in-Time Management Platform. JMP wraps together real-time application delivery, fast desktop provisioning, and policy management to create an automated and complete user experience.

In the world of Citrix, enterprise applications are managed via Citrix StoreFront, which is an enterprise app store that manages access to and delivery of new applications to users. It's a self-service app store that allows users to interact on their own with applications they're allowed to deploy.

## **USING EXISTING ENTERPRISE DESKTOP MANAGEMENT SERVICES**

Just because you deploy VMware Horizon or Citrix VDI, you don't have to instantly jettison your existing investment in tools like MEM. If you're comfortable with that tool, you don't need to use any of the application delivery mechanisms available with your VDI solution. You can keep using MEM's application catalog and other features. Enterprises have devoted a lot of time and effort to perfecting MEM deployments. You can extend this investment into your VDI environment simply by deploying the MEM agent to each of your virtual desktops. This adds a bit of complexity at the beginning of your project, but if you already have a staff that knows how to deploy software and patches and is using MEM's built-in malware scanner, keep using it as a part of your VDI deployment.

## IN THIS CHAPTER

- » Comparing economic factors between solutions
- » Assessing desktops and laptops
- » Diving deeper into each EUC architecture

# Chapter 5

## Managing the Dollars and Sense of End-User Computing Options

**A**t the end of the day, everything comes down to money. Dollars and cents ultimately drive most business decisions. In this chapter, I provide high-level guidance about what you need to consider in each of the end-user computing (EUC) models that this book discusses.

### Comparing Economics: Traditional versus Server-Based versus VDI versus DaaS

Table 5-1 provides a high-level overview of the expenses that go into each of the different EUC models. It's extraordinarily difficult to include the necessary nuance that this chart deserves, so I'm making some assumptions. First, I assume that there is no bring your own device (BYOD). All end points are provided by the organization. Plus, you'll notice that there is a significant spectrum in

some of the cells. For example, in the “End-point hardware” row, “Low to high” indicates that you can choose low-cost end points, such as terminals, or high-cost end points, such as business-class laptops with discrete graphical processing units (GPUs).

**TABLE 5-1 Comparing EUC Economics**

Category	Traditional	Server-Based	VDI	DaaS
Overall simplicity	Low to high	Medium	High	Low
Time to value	Low to high	Medium	High	Low
Operational expenditures (OpEx)	Low to medium	Medium	Medium to high	Medium
Capital expenditures (CapEx)	Medium to high	Medium	High	Low to high
Servers	N/A	Medium	High	N/A
Storage	N/A	Medium	High	N/A
Networking	N/A	Medium	Medium to high	N/A
End-point hardware	Medium to high	Low	Low to high	Low
Software licensing and maintenance	N/A	Medium	High	N/A
Cloud desktop licensing (includes management)	N/A	N/A	N/A	Medium
Management tools	N/A to high	N/A	High	N/A
Staffing and skills development	Medium	Medium	High	Low
Bandwidth	N/A	Low to medium	Medium to high	Medium
Data-center space, power, cooling	N/A	Medium	High	N/A to low

## DOLLARS AND CENTS: PHYSICAL DESKTOPS AND LAPTOPS

The big question today is: What kinds of desktops and laptops do you need to buy for your users? It comes down to the type of environment you intend to operate and the level of flexibility you want to offer your employees.

Laptops are more expensive than desktops, though both have come down in price over the years. When you consider the portability of a laptop, you have to ask yourself if the additional flexibility provided by such devices will have a return in terms of employee productivity.

This information also doesn't include operating system licensing because those models vary widely for different types of organizations. You ultimately need it in all environments anyway, so it doesn't always affect the cost. The first four rows of data in the table provide a high-level outcomes comparison and the remaining rows get into more detail.

## Painting a Complete Economic Picture

The following sections provide guidance about what to consider as you develop a total cost of ownership (TCO) for different types of EUC environments. These sections don't provide dollar figures, but they indicate orders of magnitude and the options within each category.

For each of these sections, I use a variation of Table 5-1 and list only the details for that category with an explanation for what to look for.

### Calculating the total cost of a traditional desktop environment

A traditional desktop environment can vary widely from cheap to expensive, depending on what you're trying to accomplish, as explained in Table 5-2.

**TABLE 5-2 Traditional Desktop Environment**

Category	Traditional	Details
Overall simplicity	Low to high	A traditional desktop environment can run the gamut from super-easy to manage to super complex. If you buy desktops or laptops preimaged from a vendor and you don't do anything else, the process is easy. On the other hand, if you've decided to deploy your own imaging, deployment, and management solutions, such as Microsoft Endpoint Configuration Management (MEM), complexity can ramp up rather quickly.
Time to value	Low to high	This is a function of the simplicity in this case. If you buy devices that are ready to go, time to value is very low. If you buy devices and need to manually install everything, time to value may be high. If you use a tool such as MEM, time to value will likely fall in the middle.
OpEx	Low to medium	On the low side, if you manually manage your environment, OpEx expenses stay relatively low. If you use automation and deployment tools, your management costs ramp up.
CapEx	Medium to high	In a traditional environment, you spend a good chunk of money on an annual basis refreshing your desktop fleet. If you have high-end laptops in the mix, your costs may be high.
Servers	N/A	You don't need servers to directly support desktops, unless you choose to deploy something like MEM. In this table, the hardware categories reflect only what's needed to directly host the desktop fleet.
Storage	N/A	N/A
Networking	N/A	N/A
End-point hardware	Medium to high	Your fleet needs refreshing on a regular cycle. To calculate the cost per year, determine the average cost of acquiring an end point, multiply by the number of end points, and divide by the length of your replacement cycle. This assumes that you turn over an even percentage of your fleet annually.

Category	Traditional	Details
Software licensing and maintenance	N/A	N/A
Cloud desktop licensing (includes management)	N/A	N/A
Management tools	N/A to high	If you choose to deploy MEM or a similar tool, you need to procure the software as well as the skill set necessary to manage it. These tools can get complex and expensive, and the skills can be difficult to find. This bucket includes all the licensing and software maintenance associated with these tools as well as any servers — virtual or physical — needed to host them.
Staffing and skills development	Medium	To support a desktop environment takes a reasonable amount of skill, even in moderately sized organizations that don't use management tools.
Bandwidth	N/A	N/A
Data-center space, power, cooling	N/A	N/A

## Calculating the total cost of a server-based desktop environment

A server-based desktop environment can operate with a fleet of inexpensive thin clients, or you can simply use a software client that's built into Windows on your existing desktops. The calculations in Table 5-3 assume you're going to use all thin clients.

## Calculating the total cost of a VDI environment

This is where a one-size-fits-all TCO calculator is impossible to develop. There are dozens of variables. As such, the information in Table 5-4 is general in nature, but it provides guidance about where you need to probe more deeply for more information.

**TABLE 5-3 Server-Based Desktop Environment**

Category	Server-Based	Details
Overall simplicity	Medium	You need to deploy all your thin clients, as well as all the server-side services.
Time to value	Medium	Assuming that the desktop environment is ready to go, you need to deploy the necessary components to host a session virtualization option such as Terminal Services/Remote Desktop Services. In general, this isn't difficult, but you may need to deploy gateway servers if you intend to operate an environment of any size. This increases the complexity.
OpEx	Medium	You need to procure appropriate licensing from Microsoft to enable this option. This is an ongoing expense.
CapEx	Low to medium	The environment consists of inexpensive thin clients, as well as a series of physical or virtual servers running the various components to enable the solution. These costs are not significant compared to other options.
Servers	Medium	You need servers to host the terminal sessions. Because terminal sessions share the underlying operating system between sessions, you can achieve a greater density of sessions compared to trying to run virtual desktops, each with its own operating system. To calculate this portion, determine the number of physical servers you need.
Storage	Medium	The server farm needs some type of shared storage. To calculate this portion, determine the storage capacity you need, as well as the performance, and procure a solution that aligns with both.
Networking	Medium	All the terminals and the storage must be able to communicate with each of the other components, including the servers. To calculate this portion, add the amount of additional networking equipment you need to buy to support the environment.

Category	Server-Based	Details
End-point hardware	Low	The assumption is that you're running a fleet of thin clients. These do not need to be replaced nearly as often as traditional desktops. Therefore, you're paying less per device <i>and</i> you get to keep them for a longer period of time. It's a win-win. To calculate this portion, determine the number of terminals you need and multiply by the price per unit.
Software licensing and maintenance	Medium	Microsoft Remote Desktop Session Host (RDSH) licensing and any end-point operating system licensing. Include the cost from Microsoft for this component.
Cloud desktop licensing (includes management)	N/A	N/A
Management tools	N/A	N/A; included with RDSH
Staffing and skills development	Medium	To support a desktop environment takes a reasonable amount of skill, even in moderately sized organizations that don't use management tools. For this component, include the annual salary of all the people managing the environment, and multiply by the percentage of time, on average, that they spend on this function. Also include the cost of any training courses they need in order to gain the necessary skills.
Bandwidth	Low to medium	Your network must be reasonably capable to support this deployment.
Data-center space, power, cooling	Medium	The servers, storage, and networking equipment require rack space, power, and cooling. Determine the cost per unit (U) of rack space and the power draw from each of the servers you have to deploy. You can also get cooling figures by reviewing the technical specifications for your servers to determine their heat output and get a rough calculation for what it costs to remove that much heat.

**TABLE 5-4 VDI Environment**

Category	VDI	Details
Overall simplicity	High	VDI requires that you deploy and configure a series of servers, build images, and implement a deployment strategy for your virtual desktops.
Time to value	High	Most VDI deployments take a while to build. A period of time is needed for testing to help determine the performance characteristics of your environment's desktops.
OpEx	Medium to high	You need to procure appropriate licensing from Citrix, VMware, and possibly Microsoft to enable this option. This is an ongoing expense.
CapEx	High	The initial hardware includes a series of servers, robust storage, and the networking components to bring it all together. You may also need to include thin clients for certain users.
Servers	High	You need servers to host the VDI environment. These systems can't support the density of terminal environments because each desktop virtual machine (VM) includes a complete operating system, which imposes significant overhead. To calculate this portion, you need to determine the workload characteristics that match your environment. These dictate the number of servers you need. You need to maintain a supportable vCPU-to-pCPU ratio to ensure consistent performance.
Storage	High	Your VDI environment requires robust storage that can survive intense periods of activity and has sufficient capacity to host all your virtual desktops. You almost certainly need a hybrid storage solution, an all-flash solution, or a hyperconverged infrastructure solution. <b>Note:</b> A hyperconverged infrastructure solution more than likely will bring down the overall cost and will replace the servers and storage components in this matrix. To calculate this portion, determine the storage capacity you need, as well as the performance, and procure a solution that aligns with both.
Networking	Medium to high	All the server and storage components and the end points must be able to communicate with each of the other components. To calculate this portion, add in the amount of additional networking equipment you need to buy to support the environment.

Category	VDI	Details
End-point hardware	Low to high	The end points in this case can be anything from a phone to a thin client to a desktop or even a laptop when you're using a streaming VDI technology. You may already own hardware that you can reuse, or you may augment an existing environment with a new VDI deployment. Regardless, you need to determine the types of end points you want to buy and the quantity of each, and then do some multiplication to arrive at a final figure for this component.
Software licensing and maintenance	High	Software licensing costs in a VDI scenario can be significant. You need to factor in the VDI software as well as any optional components you may choose to deploy. You also need to pay serious attention to your Microsoft license agreement and make adjustments to ensure that you stay in compliance with licensing rules for desktop operating systems running in a virtual environment. You may be able to reduce this cost somewhat by deploying on a hypervisor that doesn't require separate licensing, such as Nutanix AHV running on Nutanix hyperconverged appliances.
Cloud desktop licensing (includes management)	N/A	N/A
Management tools	N/A	This is included in the software licensing section.
Staffing and skills development	High	Supporting VDI takes a lot of skill. The administrator has to be familiar with virtualization, storage, deployment, and desktop technology. For this component, include the annual salary of all the people managing the environment, and multiply by the percentage of time, on average, that they spend on this function. Also include the cost of any training courses they need in order to gain the necessary skills.
Bandwidth	Medium to high	Your network must be very capable to support this deployment.

*(continued)*

**TABLE 5-4 (continued)**

Category	VDI	Details
Data-center space, power, cooling	High	The servers, storage, and networking equipment require rack space, power, and cooling. Determine the cost per U of rack space and the power draw from each of the servers you have to deploy. You can get cooling figures by reviewing the technical specifications for your servers to determine their heat output and get a rough calculation for what it costs to remove that much heat.

## Calculating the total cost of a DaaS environment

DaaS starts to look pretty good in a number of ways as you compare it to VDI. Take a look at Table 5-5.

**TABLE 5-5 DaaS Environment**

Category	DaaS	Details
Overall simplicity	Low	DaaS is dead simple to deploy. All you have to do is provide a credit card.
Time to value	Low	Immediate time to value.
OpEx	Medium	DaaS exchanges some CapEx for OpEx, so it may have higher OpEx over time, but the CapEx side stays low.
CapEx	Low	Any end point that supports HTML5 works.
Servers	N/A	N/A
Storage	N/A	N/A
Networking	N/A	N/A, except for basic network connectivity for end points.
End-point hardware	Low	The end points in this case can be anything that supports HTML5.
Software licensing and maintenance	N/A	N/A; included in the cloud desktop licensing fee.

Cloud desktop licensing (includes management)	Medium	This is the main area of cost. You pay a monthly fee that includes all aspects of the service, except for your application licensing and your end-point devices.
Management tools	N/A	N/A; included in the cloud desktop licensing fee.
Staffing and skills development	Low	An IT generalist can deploy the service. If you want to enable some advanced capability, such as integration with Azure AD or local Active Directory or with single sign-on solutions, you may need more skill.
Bandwidth	Medium	You need moderate bandwidth to support H.264 streaming.
Data-center space, power, cooling	N/A	N/A

## Calculating the total cost of DaaS in a hybrid cloud environment

As you begin to move DaaS into a hybrid cloud environment, the economics may shift a little, but you still won't have to procure hardware separately, at least initially.

Table 5-6 offers guidance on the total cost of DaaS in a hybrid cloud environment.



TIP

If you need some ballpark figures on hardware acquisition and power needs, take a look at the Nutanix Sizing Calculator (<https://sizer.nutanix.com/#/>).

**TABLE 5-6** DaaS in a Hybrid Cloud Environment

Category	DaaS	Details
Overall simplicity	Low to medium-low	Hybrid DaaS adds a little complexity because of the need to combine two environments, but it isn't onerous.
Time to value	Low	Pretty quick time to value, although you'll have to deploy a local hyperconverged infrastructure solution to support the on-premises portion.

(continued)

**TABLE 5-6 (continued)**

Category	DaaS	Details
OpEx	Medium	DaaS exchanges some CapEx for OpEx, so it may have higher OpEx over time, but the CapEx side stays low compared to other options.
CapEx	Low	Any end point that supports HTML5 works.
Servers	TBD	A hyperconverged solution that combines servers and storage is necessary. If that solution is acquired as part of your monthly DaaS fee rather than purchased outright, your calculations here are zero.
Storage	TBD	See “Servers.”
Networking	Low	The local appliances must be able to talk to one another. Account for sufficient networking for this purpose.
End-point hardware	Low	The end points in this case can be anything that supports HTML5.
Software licensing and maintenance	N/A	N/A; included in the cloud desktop licensing fee.
Cloud desktop licensing (includes management)	Medium	This is the main area of cost. You pay a monthly fee that includes all aspects of the service, except for your application licensing and your end-point devices.
Management tools	N/A	N/A; included in the cloud desktop licensing fee.
Staffing and skills development	Low to medium	An IT generalist can deploy the service. You may need some additional networking knowledge.
Bandwidth	Medium	You need moderate bandwidth to support H.264 streaming.
Data-center space, power, cooling	Low	There is no service overhead in hybrid DaaS. You don’t need to deploy myriad support servers to enable the solution. Just calculate what the individual appliances need.

# Chapter 6

## Ten EUC Factors to Remember

**K**ee these ten important factors in mind as you embark on your end-user computing (EUC) transformation:

- » **Outcomes and methods are different things.** The intended outcomes for the EUC environment haven't changed since the earliest mainframe days, but the methods by which those outcomes are achieved have evolved. Organizations still expect users to interact with enterprise applications using the EUC environment, but that environment looks vastly different than it used to. There are myriad methods by which you can achieve your EUC goals.
- » **No two environments are alike.** As you're considering your needs and taking stock of your costs, make sure that you run a pilot program before you fully commit to a solution. Your individual needs may not align with what you expect and you may need to change course. Consider using one of the many virtual desktop infrastructure (VDI) sizing calculators available on the web.

- » **Connection protocols aren't created equal.** You have your choice of a number of connection protocols, from the RDP standby to PCoIP to HTML5. Your protocol choice will be dictated by the solution you choose and the types of applications you need to support.
- » **EUC is an original use case for hyperconverged infrastructure (HCI).** HCI revolutionized VDI and it continues to do so today, but it can also support your terminal-based workloads as well as your hybrid desktop-as-a-service (DaaS) environment.
- » **DaaS is rising.** If your EUC plans don't currently include DaaS, start adjusting them, particularly as hybrid and on-premises DaaS solutions continue to grow in popularity.
- » **Bring your own device (BYOD) isn't dead.** No matter what anyone says, BYOD is here to stay. Your EUC environments must take into consideration that users are becoming more mobile and more remote. As the market sees the rise of a mobile workforce, a combination of BYOD and DaaS is likely to become the new normal.
- » **Total cost of ownership (TCO) is notoriously difficult to nail down in EUC.** All kinds of costs may or may not be included in your TCO. Do you include staff costs that are considered "sunk" costs or do you leave them out? Do you include a portion of a storage array you plan to use or do you skip that part? Does your TCO calculation include *all* of the licensing components you need? It's not just VDI that can be tough to size.
- » **You may not save money with VDI.** Then again, the point may not be to save money directly, but to enable flexibility and agility in the desktop environment, which may translate to increased user productivity.
- » **End points are whatever you want them to be.** Today, smartphones, tablets, desktops, and laptops can be viable end points in an EUC environment. Make sure your organization can support them all.
- » **Graphics processing units (GPUs) are also rising.** The massive core count in GPUs makes them a natural fit for offloading graphics-intensive tasks in a server-based EUC environment. Consider the role of GPUs in your environment and determine if the user experience needs a boost from this hardware.

# Provide excellent end-user experience wherever you deploy



Explore Nutanix solutions for End User Computing at  
[nutanix.com/euc](https://nutanix.com/euc)

**NUTANIX**<sup>TM</sup>

# Explore new options for end-user computing

The foundational needs of end-user computing—providing access to enterprise-provided applications and workloads—haven't changed much since the beginning, but the methods by which this goal is accomplished have transformed over the years. User demands have evolved and enterprise capabilities have matured to support myriad new opportunities for advancing this resource. This book is your guide to the latest methods for deploying application and workload services to your employees.

## Inside...

- Trace the history of end-user computing
- Explore virtual desktop infrastructure
- Accelerate deployment with desktop-as-a-service (DaaS)
- Apply cloud advantages to desktops
- Integrate third-party services
- Compare total cost options

## NUTANIX™

**Scott D. Lowe** is a former CIO and co-founder of ActualTech Media, a content creation and demand generation firm focused on creating content, conducting market research, and connecting technology companies with the right audiences.

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

for  
**dummies**®  
A Wiley Brand

ISBN: 978-1-394-17027-2  
Not For Resale



# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.