

Realtime  
publishers

# Desktop Virtualization and Evolving Strategies for IT Service Delivery

Jaime Halscott

sponsored by



Chapter 2: Smart Approaches for Merging Desktop Virtualization Today and Tomorrow .. 16

    Lesson 1: How Legacy Solutions Integrate Into Service Delivery ..... 17

        Take Inventory..... 17

        Analyze Supporting Components and Processes..... 18

        Adapting Server Solutions to Desktop Virtualization..... 20

        How Management Solutions Fit ..... 21

    Lesson 2: How the Cloud Affects Service Delivery to Users..... 21

    Lesson 3: Technologies We Never Saw Coming ..... 23

    Lesson 4: Finding a Home for Traditional VDI..... 25

    Lesson 5: When It Doesn't Fit, Cut It Out..... 27

    What's Next ..... 28

## **Copyright Statement**

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

[**Editor's Note:** This book was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology books from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

## Chapter 2: Smart Approaches for Merging Desktop Virtualization Today and Tomorrow

---

Although the promises of desktop virtualization both in terms of scalability and cost savings don't always work the way they do with server virtualization, many of the technological similarities can allow solutions to serve double duty. The previous chapter looked at the key differentiators between server and desktop virtualization in general, including cost models and project justifications, problems with supporting different types of users, and the reasons for a desktop virtualization project in general.

The trend towards desktop virtualization, as a natural extension of the trend in server virtualization, shows no signs of slowing. Even if there are doubts about whether desktop virtualization is going to address all the needs of a business, it is certain that businesses are going to try to leverage desktop virtualization for at least some portion of the infrastructure. The challenge then becomes determining where to leverage and where not to, what to keep, what to reuse, and what to retire.

This chapter looks at which legacy solutions fit into virtual desktop service delivery and how to best use them. We will explore how the cloud is affecting the desktop virtualization trend. In addition, there is a whole series of emerging technologies that will make you rethink where and how to leverage virtual desktops. Based on this information, you'll be able to determine where desktop virtualization best fits into a service delivery model.

To have a successful IT strategy, a business must look at not only what exists currently but also what is emerging. A smart business has to understand what has taken the market by storm and how it has become an embedded part of the IT landscape. Desktop virtualization is going to be affected by trends not just related to virtualization but technology in general. Lesson 1 will assist in determining how legacy solutions will fit in overall goals of service delivery, from hardware to management software and even support processes.

## Lesson 1: How Legacy Solutions Integrate Into Service Delivery

When I talk about legacy, I don't want there to be any confusion as to what I mean. In the realm of desktop virtualization, I'm speaking about technologies that are typically used either in the server space exclusively or have not traditionally been a part of a desktop service discussion.

It may seem a little odd that I'm referring to legacy solutions as part of a virtual desktop discussion, but the fact is, almost no one is going to be able to rip out all of their existing infrastructure just to start from scratch with VDI. Honestly, why would any business decide that it is appropriate to redesign everything just to support an emerging technology? That technology would have to be so revolutionary and give that business such a competitive edge that there would be no other option than to adopt it completely and suddenly. Desktop virtualization isn't that type of technology.

Chapter 1 established that VDI wasn't the final destination but merely a part of a transformative journey that improves the end user experience and enables IT to more cost-effectively support users' needs. With that focus, consider the following questions for your business:

- If we are going to have VDI in some form, what do we do about all of the legacy desktops that are still in production?
- What about the servers and networking equipment in the data center and corporate offices?
- What about the management platforms and software we use to support it all? What do we do with it? Do we reuse it? Do we repurpose it?

You should reuse and repurpose as much of it as possible. This makes natural sense, as we all have to work within limited budgets. Then how do you determine what you can reuse and repurpose?

### Take Inventory

Let's start by taking an inventory of the current business infrastructure. We need to determine which components can be reused and which have to be replaced. In addition to hardware, this inventory process will necessarily include software and even support processes too.

Hardware is a good place to start with the inventory process. You already know the requirements for virtualization platforms in general. This is one area where the differences between server and desktop virtualization aren't all that varied. I struggle to think of a server that has shipped in the past few years that doesn't fully support the technologies required for virtualization.

From a server hardware standpoint, if your servers are capable of supporting a virtual server infrastructure, they are going to be capable of supporting a VDI. Vendors don't sell different servers for each purpose. These servers are just as good for virtual workloads as they are for physical ones. It doesn't matter whether those workloads are virtual server or virtual desktop in nature.

Hardware support for virtualization has been around for several years. Although this hardware support has matured and brought speed and scalability gains, there is one key difference between now and when server virtualization started to become mainstream. When server virtualization started to come into the forefront of IT strategies, servers didn't always support hardware-accelerated virtualization. Virtualization platforms required direct hardware support for their hypervisors that many servers could not provide. As a result, only the most current generation of servers at that time were viable candidates for becoming virtual hosts.

Fast forward 5 years and we see that all servers currently shipping feature some form of hardware virtualization support. The newer the server, the better the support will be with technologies that allow much more efficient direct access to hardware resources. Nonetheless, all servers in the past few generations will have the required support for desktop virtualization efforts. To move a physical server from a server virtualization role to a desktop virtualization role, not much is required. Certainly, the hardware does not have to be replaced, creating the first chance to reuse infrastructure components.

In the first chapter, I identified the importance of creating an effective pilot to determine performance needs and ensure that any solution will fit the users' expectations for usability. Being able to employ servers that still meet the requirements for virtualization but might be aging to the point where they are targeted for replacement will be a great way to test desktop virtualization without adding new servers. To complete this step, take an inventory of all of your servers and indicate which include hardware virtualization support.

### Analyze Supporting Components and Processes

Let's move from the server hardware itself to the other supporting components in the data center. In this area, upgrades that were performed to support server virtualization projects will translate into real-world use with desktop virtualization. There are no special requirements for networking hardware or storage systems to support desktop virtualization. You can leverage all the switches, routers, power distribution, and such that you've already put in place without additional cost.

With that being said, there is a caveat to reusing all of the supporting components without planning for any upgrades. Implementing VDI can be a massive tax on the existing network infrastructure. Being able to reuse a good portion of the networking infrastructure doesn't mean that networking infrastructure upgrades and optimization won't need to occur. An entire industry of specialized networking equipment designed just for desktop virtualization scenarios has emerged.

Let's turn our attention to support. At this point, you will have supporting software that is used for desktop and server support. You also have management and support processes and methodologies that will have to be analyzed. Because the desktop operating systems (OSs) are the same regardless of whether they are on a physical desktop or presented virtually, you can leverage all your existing investments in your management software platforms.

You would think that support processes and methodologies wouldn't change appreciably either; however, support changes are one of the primary drivers in going to desktop virtualization. The support team will still have to respond to hardware problems occurring with a virtual desktop, such as failures with the desktops or thin clients that are accessing the virtual desktops. The team responding to that will have a different set of diagnostic and repair steps to perform than they would with legacy solution support.

Looking at how much can be reused, it seems that this VDI transition is pretty much a slam dunk, right? Not so fast! There are some problems. In order to understand what these problems are, let's step back a minute and look at traditional desktop hardware. Business desktops aren't designed with high availability in mind. They typically don't have redundant hard drives. They don't have redundant power supplies and data center-class components. They are designed to be cost effective and border on being commodities.

It may be fine that these business desktops don't typically include redundancy features. When a new PC can be imaged and put in place when a critical component fails, that downtime affects only a single user. It isn't acceptable when that critical component fails on a server that must support 10, 20, 50, or more virtual desktops. That type of outage would critically affect a business and illustrates the problem: VDI takes all of our resources for multiple desktops and puts them into a single point of failure.

One of the tricks I've found in creating successful VDI projects is to determine how much exposure the failure of a single component will cause in terms of downtime and lost productivity to all the virtual desktops that rely on that component. It sounds simple enough, but it really isn't. To better illustrate this point, let's use an example where you decide that you want to put your company's accounting department of 20 users onto virtual desktops. Here is how it looks in the physical desktop world in terms of cost and supportability:

Each user gets a PC. The cost of each of these accounting PCs is \$1000 and there are two spares of the same type and cost. The total cost of this physical desktop infrastructure is \$22,000. When a user PC fails, the IT support team images a spare with all the standard applications and customizations that the user needs and replaces the PC for the user. This image, deployment, and setup process takes about 1.5 hours. A failure affects one user only. The failure of a single user isn't critical because there is more than one of each type of user and the department continues to operate while this single user is down.

Now let's take a look at this scenario through a post-VDI implementation lens:

To make your accounting VDI rollout as cost effective as possible while supporting all 20 users, you put all the accounting users on a single server using shared storage. This server has some redundant components such as power supplies and hard drives so that it can continue to operate with the loss of one component here and there. The cost of the base server is \$6000 with the memory and processor required to support the workload of all 20 users. The cost of the shared storage, almost certainly from a SAN, is the wildcard here, but can be \$50,000 or more.

I understand that these numbers are somewhat difficult to quantify exactly. The costs of server components are generally going to be more expensive than those of a desktop even before any kind of redundancy is added to the equation.

This rough approximation of the two technology approaches reveals that you will get roughly the same results; however, VDI, because of reliance on high-availability hardware and SAN storage, is going to be much more expensive. The nature of the shared resources in the virtualization route means that the failure of a critical component will hinder productivity for everyone, not just a single user. In a pilot, this may be acceptable, but it certainly won't be in a production environment.

### Adapting Server Solutions to Desktop Virtualization

One of the most common ways to improve the reliability of a virtual infrastructure is to implement high availability in the form of a SAN and clustering. Neither of these high-availability solutions is new, but their use in supporting a desktop infrastructure is. Clustering and SANs are common for servers; a natural extension of that is for use with server virtualization. However, it is a slightly different paradigm for VDI. These technologies will bring the needed reliability and scalability to desktop virtualization but not without costs.

SAN storage is considerably more expensive than direct attached storage (DAS) primarily because of the redundancy and scalability that is part of the SAN package. The move to VDI means you are trading inexpensive disk space on a single physical desktop for expensive virtual desktop space on a SAN simply because so many virtual desktops rely on the shared storage.

Technologies like thin provisioning allow you to save a considerable amount of disk space up front by allowing the virtual desktop disk usage to start with only the minimum requirements and grow as needed (as with server virtualization). However, the concern still exists that the nature of a virtual desktop will consume a fair amount of space on the SAN.



Remember that the goal is end user service delivery. You are creating a user workspace so that the user can perform their particular job role. Virtual desktops aren't the only means to do so. There are alternatives such as application virtualization where the application is encapsulated and isolated from the OS. These virtual applications can be deployed to users without the extra support requirements of an entire virtual desktop.

Wouldn't it be much more cost effective to use application virtualization instead of simply giving users a complete desktop? Certainly there would still need to be the investment in redundancy, but the storage costs would be far less. Letting the users have their own physical desktop but giving them on-demand access to applications has other advantages that are outside the scope of this guide. It is important to note that application virtualization is a technology that can fit better than VDI for many uses. In addition, application virtualization works well in a VDI, so these two virtualization technologies can happily coexist.

### How Management Solutions Fit

Your business is going to have an array of software that is used to manage and maintain your physical desktop infrastructure, including security software, patch management solutions, and troubleshooting and diagnostic tools among other point solutions. Most of this is going to be reusable with one caveat. The nature of VDI as a shared infrastructure means that you are going to want to minimize the impact on users from actions like virus scans and patch deployments. These can rob performance from individual desktops and that effect is going to be magnified on a shared VDI. There are vendor recommendations that antivirus solutions not even be run on top of VDI instances. Likewise, a best practice is to patch only VDI virtual machine parent images rather than the individual instances.

Many vendors include some form of management toolset for their virtual desktop management platform—partly to ease the cost burdens of buying additional third-party components for patch management. You might choose to leverage the vendor-supplied management framework, but I recommend that you test your existing management solutions to assess impact. If there is no impact or the impact is negligible, you can re-evaluate the ROI of your management solution at subscription or support renewal time.

## Lesson 2: How the Cloud Affects Service Delivery to Users

We just looked at what it would take to build a VDI in a private data center, but what about the buzz about the emergence of the cloud as a home for VDI? Can desktops be hosted in the cloud and still be a part of a VDI rollout? How does this fit with service delivery?

There is a natural temptation to leverage a cloud service for hosted desktops, as cloud-based solutions are maturing rapidly and vendors have put a lot of effort into figuring out how to support VDI challenges such as scalability on demand. Thus, the cloud computing service provider has put the investment into designing and testing the solution; a private company with budget constraints might not be able to afford to develop a VDI from scratch.

Investing in a VDI has been a barrier to adoption for many businesses. As discussed in Chapter 1, there are several costs that might not be obvious from the onset ranging from infrastructure upgrades to software licensing. Using hosted desktops in the cloud will replace the in-house backend server and infrastructure needs of a VDI but won't address the client, or access side, of the equation. There will still have to be some sort of means of accessing the virtual desktops, whether via a physical desktop, thin client, or some other device like a tablet. Turning to the cloud for a hosted VDI solution will only affect the hardware capital expenditures for the data center-side of the equation. There will still be the costs for the desktops or thin clients that will be used to access the hosted VDI.

In a hosted-desktop setup, bandwidth and latency are still going to be important factors in the overall end user experience. Ideally, a provider that offers a cloud-based solution for hosted virtual desktops will have a well-developed means for analyzing and allocating bandwidth for access to ensure that the presentation will be as indistinguishable as possible from a local VDI. The provider should also have a means for supporting users with low latency access at multiple locations around the globe. The problem is that no such service exists. Supporting geographically-distributed users will require planning to make sure that bandwidth and latency problems don't adversely affect users. No hosted VDI vendor can accommodate these needs for seamlessness.

Perhaps the most difficult challenge to overcome with a cloud-hosted VDI solution is access to server-based resources. The earlier discussion assumed that the data center that was going to be hosting the virtual desktops was also the same one that had the file servers, email servers, application servers, and so on. In that scenario, the access to server resources would be very fast and the problems with bandwidth and latency virtually disappear.

Going to a cloud-hosted VDI introduces a big problem: Your users working on the LAN leave the network to go out to the cloud to access virtual desktops which in turn come back to the LAN to access server resources. This "in and out" network traffic is going to seriously impact the user experience because of the increased latency and reduced bandwidth. It could possibly expose your network to security problems inherent with opening inbound and outbound access to these resources. Additionally, it places increased demand on the speed and reliability of Internet access. One can appreciate the irony of having perfectly good desktops on the LAN and perfectly good servers on the LAN, neither of which can "talk" because there is no Internet access to the virtual desktops hosted in the cloud.

Where the cloud-based VDI solution makes sense is in the limited demands of a particular class of user. Let's use an example to see where this works: Your company has a large sales team that accesses a single application for customer management and order entry. This desktop image is standardized and highly locked down. The application does not access any data on your network but instead accesses a third-party, Internet-based provider such as Salesforce.com.

This scenario best illustrates a fit for cloud-hosted VDI. Keep in mind that the more basic the functions that need to be performed and the less the volume of data that must be transferred, the better the cloud VDI solution works. So this begs the question, How is this better than a physical desktop with a locked down configuration? For many businesses, the answer is that it's not really better. Where a fit exists is for field-based personnel that are outside of the typical management reach of the business or are accessing the virtual desktop from personal systems but still needs a locked-down desktop for their role.

Once again we run into limitations with where resources are located based on how much network speed is available and how reliable it is. Costs for Internet access have been decreasing, but not as quickly as VDI technology adoption has been increasing. This is perhaps why so much effort has gone into the development of protocol optimizations by virtualization vendors. The smallest increases are still critical to providing acceptable user experiences. Virtually all major vendors have made great strides in reducing protocol overhead, but this isn't going to be a panacea to the problems associated with accessing remote resources over slow or high latency links.

The cloud does have a role in VDI, but for many businesses, that role is limited at best. Until cloud service providers can effectively tackle the responsiveness issues that are going to be critical to service delivery, the role of the cloud will be for very specific workers and roles.

### Lesson 3: Technologies We Never Saw Coming

Thus far we have been looking exclusively at a Windows desktop served virtually to a physical desktop. This is almost certainly going to be a large portion of the enterprise VDI adoption. What was unexpected is the emergence of a whole host of different form factor devices that users have gravitated to. This wasn't a concern with server virtualization because the presentation of resources to users didn't really come into play. However, in a VDI, the way the user interacts is critical to service delivery.

Two years ago, server virtualization was already mature and desktop virtualization was coming into its own. However, no one was quite sure how the forms of access to virtual desktops were really going to take shape. At that time, Android was just a glimmer in the eye of a Google developer. The iPad didn't exist. The iPhone, Blackberry, and Windows Mobile phones were primarily collaborative devices. Things have rapidly changed. Apple's iOS took the world by storm as part of a successful tablet and phone family of products with a wide array of applications. Android came to power a variety of devices including phones and tablets. Microsoft introduced a whole new mobile OS and announced major plans for tablet functionality in the next major release of Windows. How does this all fit together in the context of desktop virtualization and service delivery?

One of the first technologies to emerge in the business world that wasn't designed for the business world was the iPhone and the resulting iPhone OS and later iOS. Users that had these phones for personal use almost immediately demanded that businesses support them. Apple responded by building in more business-oriented support, including support for Microsoft Exchange Server. The resulting explosion of iPhone adoption lead vendors to develop additional applications for business needs.

Tablets flooded onto the scene with the introduction of the iPad. Although tablet form factor PCs have been on the market for almost a decade, the iPad propelled the true tablet form factor into the mainstream. The iPad was able to capitalize on the success of the iPhone and the inroads that had been made with application development for the iPhone. The result is a unified ecosystem and a desire for service delivery on these new platforms.

Not long after the success of the iPhone, Google introduced a competitive OS in Android. Android similarly serves as the platform for a whole host of devices including smartphones and tablets. Android uses a similar application purchase and installation method to the one that Apple uses with the iPhone and iPad's underlying iOS. Likewise, rumor has it that Microsoft has spent considerable energy in developing a unified tablet and smartphone OS. Both the Google and Microsoft offerings have further pushed the desire for users to be able to work anywhere at any time with the same access to data and applications that they have inside their corporate networks.

This discussion will treat these OSs the same and these devices as a single class of devices. We aren't concerned with who created them so much as with the impact they have had as a whole on the ability for users to work outside of a traditional Windows PC environment.

With regard to desktop virtualization, the widespread adoption of these devices wasn't predicted simply because they were the first time we saw mobile devices that weren't PCs that brought high-resolution screens with always-on Internet access over cellular networks, support for LAN connectivity via WiFi, and application support for access to virtual desktops. The functions that these devices can perform were previously the exclusive domain of laptops—and not just laptops but rather laptops with some sort of always-on network access. Laptops with always-on networking are not as common as these new smartphone and tablet devices that ship with WiFi and 3G or faster Internet access. Additionally, these smartphones and tablets have considerably longer battery life than their laptop equivalents.

Having always-on network access is critical to virtual desktop delivery in the traditional VDI sense. Once these devices emerged with this capability, it was only a matter of time before someone saw the potential for service delivery. The major desktop virtualization vendors immediately began making versions of applications to access both their desktop and application virtualization solutions.

This is great news for the most part, but there is a catch: now there is a whole class of devices that has to be supported and secured. IT departments began having problems immediately with just providing support for secure email and contact management. Providing secure access to corporate networks for desktop access became a real concern.

These devices started out as consumer devices and weren't designed to be securely or centrally managed like desktops or thin clients. Mobile device management tools have always been behind the curve compared with desktop and server management solutions. Regardless, the reality is that these devices are becoming part of the VDI ecosystem.

The problem that the user needs to have uninterrupted service if they are realistically going to be able to access a virtual desktop with a smartphone or tablet. Unless the smartphones and tablets are connected to the LAN at your corporate office, there can be no guarantee that the user experience will be acceptable, much less preferred. But users are endeared to these devices because of their travelability, which means virtual desktop access is going to be over a cellular network or external WiFi network. This level of reliability and speed won't consistently meet service delivery goals.

A potentially better solution is to have a mechanism for secure application delivery instead of delivery of an entire virtual desktop. Regardless of whether this application delivery takes traditional or virtual application forms, this setup will allow the user a better means of working outside the corporate LAN. Traditional VDI simply isn't the best fit for this type of device and use case. So then where does a traditional VDI fit?

## Lesson 4: Finding a Home for Traditional VDI

Let's take a small step back and look at desktop virtualization in the most basic form and see where it best fits. We previously examined the requirements for a successful VDI. The biggest challenge is often providing end-to-end connectivity that is suitable for the experience that the end user is expecting. We also saw that providing virtual desktops becomes more difficult with remote and offline users because of the nature of how a desktop is accessed.

The end-to-end connectivity is affected by speed, in the form of total bandwidth/latency and reliability. To provide a VDI, this problem must be addressed. So where are both speed and reliability the highest? Earlier, we isolated this to the LAN, but it is more complex than that. Here we are looking for the ideal conditions to use traditional VDI—you don't want any kind of connectivity issues to impact how you deliver a virtual desktop. To ensure ideal conditions, you must also consider the device being used to access the virtual desktop and what sort of LAN connectivity it should use.

Leveraging existing desktops for the purpose of accessing your VDI is going to be the most common means for several reasons, including reducing hardware replacement costs. Being able to take existing physical desktops and either convert them into some form of thin client or using them as they are with RDP, ICA, or some other protocol will keep costs down on the hardware front. However, this "double desktop" scenario increases costs in management. Which one of these methods for accessing virtual desktops is going to be the best fit?

Many companies are going to be sensitive to reducing the cost of supporting a physical desktop. Even if it is converted to act like a thin client, there will be some sort of support needed to keep it running. This could be in the form of hardware support or management of the underlying software used. The level of support depends on the way the physical desktop is converted. There are several solutions out there that will accomplish the thin client conversion goal. Each has a set of benefits that will make it attractive.

The products that sit on top of Windows will still require patch management and security software. This will have to factor into the ongoing management costs. One very big positive is that there won't be problems with drivers and peripheral support for the existing hardware. Products that replace Windows will tend to reduce the security vulnerabilities because these are usually stripped down, having a small footprint and reduced attack surface. This small footprint even allows some of the thin client OSs to be loaded fresh at each boot and can even run diskless.

Existing management tools and methodologies are going to make the physical desktop as an access method for a virtual desktop the quickest to get up and running. Often this requires nothing more than a client connection. Conversely, security and ease of ongoing support may make the reinstallation of physical desktops as thin clients the most attractive means to provide virtual desktops. There needs to be a thorough analysis of the needs of the organization in order to determine whether using the desktop as an access method or as a thin client will work best.

Not all customers will choose to leverage the existing desktops in a thin client manner. Many will choose to continue to use physical desktops for their current purpose and retire them as they fail. Retired desktops will then be replaced by "real" thin clients. These vendor-supported thin clients will require upfront capital expense to purchase and deploy, but offer benefits in security and supportability from the vendors themselves.

The current generation of thin clients also offer support for the rich multimedia presentation technologies provided by the latest desktop virtualization platforms and provide offloading of some of the tasks that would otherwise be intensive on the servers hosting the virtual desktops. Some of these thin clients are even referred to as "zero clients" where there is no setup ahead of deployment. All you need to do is unbox, connect cables, power on, and go. The thin client detects the virtual desktop server and gets everything it needs to allow the user to login and start being productive.

Going the thin client route will eliminate a large portion of the support costs related to desktop hardware support. A business just needs to keep enough spare thin clients so that a failed thin client can be replaced in the field with a minimal amount of downtime. Thin clients for use with VDI offers a number of benefits: The appliance-like reliability of thin client hardware makes them inexpensive to support and maintain. The proprietary, stripped-down OS makes them very secure, and the capital expenditure is typically less up front than a traditional desktop, although the savings are realized in the reduced total cost of ownership.

The thin client in general has also proven itself to be a tried and true technology for use with terminal services and Citrix over the past decade plus. Using a thin client solution for VDI is a natural extension. Because of the static nature of a thin client, meaning it is just going to sit in a fixed location like a desk, it is going to have fast, always-on LAN connectivity. Thus, it addresses the key concerns for end user service delivery with a virtual desktop. Also, the modern thin client and, more important, the transport protocol used, has support for peripherals over USB.



Although repurposing desktops as thin clients is an attractive mechanism to save on some of the capital expenditures of transitioning to a VDI, there is still a phased path to migration away from those legacy desktops in the long run to “true” thin clients. This type of LAN-connected, fast and highly reliable connection for basic productivity users makes it the ideal place for traditional VDI. Now let’s take a look at what doesn’t work as well with VDI.

## Lesson 5: When It Doesn’t Fit, Cut It Out

A lot of this guide has been dedicated to reusing as much infrastructure as possible. The current economic conditions have stretched budgets farther than just about any time in modern IT history, so it is going to be difficult to say that something doesn’t fit and has to be removed and replaced. The good news is that it doesn’t have to be all at once. Using a phased removal process may work perfectly well for many infrastructure components, particularly existing physical desktops.

The first step is to determine the criteria for what does and does not work. Earlier we examined how existing desktops can be repurposed to serve either as thin clients or leveraged in a hybrid approach where there is still the core OS in use but also virtual desktop access from the desktop platform. With some exceptions, these desktops will make an effective part of a transition to a VDI. The exceptions are desktops that either do not meet the requirements to run the software you select to convert them to thin clients or are too old to remain in service. You may have a policy in place that mandates retirement of desktops based on age or particular specifications for which these desktops no longer fit. It is tempting to keep these desktops around, but the near commodity nature of desktops and reduced costs of thin clients may not make this the best option.

Networking infrastructure is the next main area that must be examined. The first chapter spent time discussing the merits of high bandwidth and low latency to the overall user experience. Again anything that can be used to increase speed and decrease latency in the network connectivity should be implemented if within the budget. At a minimum, Gigabit Ethernet end to end is required, and realistically 10 Gigabit Ethernet is going to be required at the core for the servers hosting the virtual desktops and communicating between each other and backend storage. Peripherals such as printers and specialized input devices must all be scrutinized. Although the majority of businesses are going to have shared printing, you’d be amazed how much printer sprawl I see. I’ve encountered many businesses where there are as many printers as there are employees. This is a terrible waste of efficiency and costs a fortune to support and supply. Once you have determined the route you will use to have users access their virtual desktops, you must test all of these peripherals and be prepared to retire those that don’t work. Don’t be afraid to use this time to correct the printer sprawl in your organization. You can even use this correction as part of the cost justification of transitioning to a VDI. The important takeaway point here is that you must be prepared to remove items that have already been heavily invested in.

Now comes the trickiest part to figure out: the virtualization platform itself and the management tools. You may have invested heavily in a server virtualization platform and corresponding set of management tools only to find out that those really don't work exactly the same when applied to desktop virtualization. This is OK. As we have already seen, the management challenges are different for virtual desktops than they are for virtual servers. Don't feel like you need to extend your virtual server platform and management tools to the VDI. There are specific management tools that will work best for each scenario.

There is good news here overall. As we have already seen, VDI isn't a perfect fit for all scenarios, so the impact of removing parts of the infrastructure that don't fit isn't going to be something done all at once. Using a phased approach to removal can be just as important to success of the project as a phased introduction of new equipment.

## What's Next

By taking a look at what you can keep and what you must change, this chapter expanded upon Chapter 1's foundation of how to use VDI and where there will be hidden challenges. In addition, it is important to step back and take a look at where the industry in general is trending so that any investments in a VDI are reusable.

This chapter determined that existing desktops can be a powerful part of a VDI used either as they are for access to a virtual desktop or repurposed as dedicated thin clients. We took a look at the cloud to examine whether it will be a realistic part of your virtual desktop future. We analyzed technologies that were never expected to be a part of desktop virtualization, focusing the discussion around service delivery. We even tried to find the best home for traditional VDI and learned how to identify where to avoid the temptation to keep components that aren't going to be a great fit for your virtual desktop strategy.

The next chapter will tie it all together. So far, we've looked at VDI and service delivery from a high-level overview, which was meant to help you determine if and how much VDI is going to be a part of your business. Chapter 3 explores what works well together and how you can fit what you have into what you want and where you expect to be with desktop virtualization. This chapter also starts to look into the near future for solutions that don't yet exist but that you might need to capitalize on any investment in VDI. Most importantly, we will determine what metrics are required to decide whether you made the right investments at the right times to maximize your return on investment.

## Download Additional Books from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this book to be informative, we encourage you to download more of our industry-leading technology books and video guides at Realtime Nexus. Please visit

<http://nexus.realtimepublishers.com>.