

Why VDI Makes Sense for Hard-Core UNIX and Linux Engineering Applications

By Brian Madden, BrianMadden.com

Contents

| | |
|---|---|
| <i>Desktop Virtualization? Seriously?</i> | 2 |
| <i>Wait: Is This Stuff Even Real Today?</i> | 3 |
| <i>Benefits of Data Center-Hosted Desktops for Engineering Apps</i> | 4 |
| <i>Downsides of Data Center-Hosted Desktops and Apps</i> | 8 |
| <i>The Bottom Line</i> | 9 |

Brought to you compliments of

OPENTEXT

By now, everyone has heard the terms desktop virtualization and VDI. It seems like we've been pounded with marketing propaganda about these technologies from companies like VMware, Citrix and Microsoft since 2006. But if desktop virtualization is so good, why are most of the world's desktops still "old-style" physical desktops and laptops?

And that's particularly true in the context of high-end, graphically intensive UNIX and Linux engineering applications. Many people view desktop virtualization and virtual desktop infrastructure (VDI) as an edge solution geared more toward simple Windows-based, Office-like applications. But many hard-core engineering apps run on UNIX, with the incorrect perception being that they require "real" engineering workstations and aren't appropriate for VDI.

Desktop Virtualization? Seriously?

Let's take a step back and make sure we're all on the same page: *Desktop virtualization* is a somewhat generic umbrella term that describes any of the technologies that allow IT organizations to separate the management of desktops and applications from the devices users use.

In the old days (you know, like 2003), everything was physical. If a user needed a workstation or laptop, IT physically installed a local operating system and physically loaded the user's applications onto it. Users were essentially locked to their devices — their apps were there, their data was there, and the environment was there. If a workstation or laptop was broken, lost, stolen or locked in the office when the user was at home, the user was out of luck: No work was getting done.

That was the physical world of 10 years ago. It applied to servers, data centers, workstations and laptops.

But then, everything changed. The idea of virtualization evolved from mainframes to x86-based computers, and suddenly we realized that servers were nothing more than disk images that we could run on any physical piece of hardware. This enabled us to do all sorts of cool things, like instantly provisioning new virtual servers without buying new physical ones. We could move our virtual servers on demand so they would always run on physical hardware with plenty of resources. We could even move them while they were running in order to perform maintenance activities on physical servers without affecting service to users.

In a few short years, we quickly realized that virtual servers were better than physical servers in every way.

As such, it didn't take long for people to ask, "Hey, since this virtualization thing is so good for servers, why can't we do the same thing for desktops?" And in a flash, the concept of desktop virtualization was born.

Of course, virtualizing servers was easy because they already lived in the data center. So converting a physical server to a virtual server was fairly trivial, and users didn't even notice (or care). But if IT wanted to virtualize a desktop, what exactly would that look like?

One method is to run the user's desktop OS in a virtual machine on his workstation or laptop. Doing so gives IT organizations many of the benefits of virtualization — in that IT doesn't have to worry about hardware compatibility and management of the client device — but it's really not too different from the traditional model in use for the past 20 years.

Another and much more impactful way to virtualize users' desktops is to physically "move" each user's desktop OS and applications off the physical desktops and into the data center. Then, users access their desktops via remote protocol, essentially allowing them to remotely control own data center-based personal desktops from anywhere in the world.

Data center-based desktop virtualization comes in many flavors, broadly divided into either VDI, where each user has a full virtual machine in the data center, or session virtualization, where each user connects to a single user session on a shared server in a data center. (Session virtualization has been around since the 1990s in various forms, including Microsoft Terminal Server, Microsoft Remote Desktop Session Host, Citrix MetaFrame and Citrix XenApp.)

This paper focuses on data center-based desktop virtualization, including its advantages, disadvantages and how it can be used to deliver high-end UNIX and Linux engineering environments to users anywhere in the world.

Wait: Is This Stuff Even Real Today?

When I talk about data center-based desktop virtualization and VDI, I often hear objections like: “Hey, I tried Citrix [or VDI] a few years ago, and it sucked!” Or people say, “Yeah, that VDI thing is fine for simple applications like Web browsing and office suites, but I have graphically intense, high-performance applications that require OpenGL, DirectX and great user performance.” Still others believe that VDI is just about Microsoft Windows applications and that there’s no option for UNIX or Linux environments.

So let’s clear the air right now: Hosting desktops in a data center is vastly different today compared with even a few years ago, and it’s like night and day from five to 10 years ago.

I also find that when I talk about remote desktops and applications, it turns out that most users have tried only the basic access protocols — typically Remote Desktop Protocol (RDP) for Microsoft Windows desktops or virtual network computing (VNC) for UNIX or Linux desktops. But today’s remoting protocols are not RDP or VNC.

We now have many high-end protocol options, including Citrix HDX, Teradici PC-over-IP and Microsoft RemoteFX for Windows environments; Thin X Protocol, used by OpenText Exceed VA TurboX for UNIX and Linux. We also have extremely powerful options for putting high-end hardware in the servers that power our virtual desktops. We can install GPU card that offer OpenGL and DirectX support to the desktop virtual machines they run, and we have options to support “blade workstations,” where we can install workstation-class blade hardware in the data center that is every bit as powerful as anything we can install on a user’s desk. And, again, thanks to advanced remoting protocols, we can actually deliver a high-performance experience to connect users into their high-end remote workstations.

So back in 2008, I would have agreed with someone who said, “Desktop virtualization is fine for office apps, but not for engineering apps or UNIX or Linux environments.” But that is simply not true now. The industry has created technologies that make this possible for all apps — whether they’re Windows, UNIX or Linux, or whether they need multiple CPU cores or GPUs or intensive storage.

Benefits of Data Center-Hosted Desktops for Engineering Apps

There are many benefits of data center-hosted desktops. Some are obvious, while others are a bit more subtle. The good news is that there are benefits that apply to all involved — the overall business, IT management and end users. Let's take a look at each of these.

Access critical engineering applications from anywhere

The fundamental advantage of data center-hosted virtual desktops and applications is that your users can access your engineering applications from anywhere. They do not need to be in the same building, city or country as the data center. They can access high-end UNIX and Linux engineering applications with specialized requirements (such as OpenGL) from client devices running different platforms and with different form factors via one of the solutions mentioned earlier that support high-end remoting protocols or a Web plug-in, rather than requiring the specific engineering environment to be installed locally.

Hire better talent, have happier workers

With data center-based desktops, the IT engineering needed to support the applications can be located near the data center rather than near the users. It also means users don't have to live near the data center.

This unlocks several benefits not previously available. What if your business didn't have to hire your talent in the same city as your IT staff or your data centers? It means your talent can live anywhere. You really want to hire someone from rural Ohio who doesn't want to move? No problem.

It also means you can hire talent wherever it's cheapest. You're a hot new biotech firm in Silicon Valley, but you can't match Valley salaries? No problem. Hire a biologist from Omaha and pay an Omaha salary while your office stays in California.

This can also benefit you on a micro-level: Users dealing with snowstorms or sick kids, or who prefer to work at night or, say, get stuck in traffic court — all can be fully productive because they don't have to be physically chained to their desks to get their work done. And they can do this on a moment's notice. For instance, they can dash out of the office without having to copy all of their files to their laptop. No more phone calls like, "What? You've been in an accident? I'll be there as soon as I copy the latest engineering files to my thumb drive!"

Provide better support

If you consolidate your data centers from dozens to two or three, you can afford to invest in better support while still saving money overall. Rather than paying dozens of mediocre engineers \$50,000 each all over the world, you can invest in a few high-quality support folks. You also avoid having dozens of IT "cowboys" around the world, inventing their own ways of doing things.

This all means you save money while getting better talent, your overall environment is more consistent and stable, and your users always have access to high-end support.

Engineering apps in the field (and on iPads)

Engineering and technical apps can be beasts when it comes to the data they require, and it's not uncommon for project files to grow to several gigabytes in size. In the old days of physical computers, a user in the field would need a laptop that had to download and synchronize all

of these multi-gigabyte data sets. Not only was this slow, but it also meant that (by definition) the user didn't have access to last-minute changes. More typically, we'd find that field users would save download time by downloading only the bare minimum — and, inevitably, they'd find they needed access to something while in the field that they didn't think about ahead of time.

In contrast, running the engineering applications in the data center and connecting remotely from the field means users always have access to the full and latest versions of their data.

Another benefit is that, since data center-based desktops and applications can be accessed from non-native devices like iPads, your users can have full access to their applications and data sets from portable devices. Imagine how this would affect an architect's tour of an unfinished sports stadium, for example. In the old world, the architect would have to pre-render some videos and load those on an iPad to take on the tour, but if clients wanted to see something that hadn't been anticipated, the architect would be unable to show it to them. In contrast, if the architect's iPad can connect to the actual design application, he can do walkthroughs at the stadium on demand and make changes right from the dirt field.

Faster time to deploy and patch new applications

Everyone knows that keeping users' workstations up to date is a never-ending struggle. Applications are constantly changed, updated, patched and reconfigured. If you're in charge of keeping your users' applications and desktops up to date, what do you think is easier: updating workstations spread all over the world, or updating workstations in a single data center?

Consolidating your users' desktops into a data center means that they're patched, updated and fixed much faster than in the traditional world, where they're sitting on someone's desk.

Reliability

With data center-hosted desktops, your critical applications run in a data center that you control. If a user's computer breaks, critical business apps aren't broken with it. The user can just move to a different desk, grab a computer from home, or go to Best Buy and pick up a laptop and be back up and running within minutes.

This also plays into the "user happiness" benefit mentioned above: If you have die-hard Mac fans, let them use Macs. You don't care!

Combine different application platforms into one overall environment

One of the challenges in the engineering space is that many engineering applications tend to run on UNIX or Linux, while users typically prefer Windows or Mac operating systems for their everyday desktop needs. By moving the engineering applications into a data center, you can provide remote access to the special UNIX- or Linux-based engineering apps while still allowing users to do their other work in Windows or the Mac.

Reduced risk of theft

By definition, data center-based desktops live in your data center. This means your applications and data also live in your data center, rather than on the hard drives of the workstations in user land. This has major benefits in terms of security.

First, when it comes to intellectual property theft, you don't have to worry about your expensive designs walking out the door on someone's USB drive (be it a rogue employee or contractor).

Second, if someone literally steals the client device, that person doesn't have a hard drive full of millions of dollars of IP.

Third, since the high-end hardware is in the data center, the client devices can be pretty "dumb," typically nothing more than a thin client or a low-end desktop computer.

Instant scalability

It's easier to scale virtual machines in a data center than physical workstations on users' desks. If you move all of your engineering applications to your data center, adding 10 more engineers during crunch time is simple — you just give them the URL to your engineering environment, and they're off and running. When the project is over, you can deprovision them just as easily, without having to send an IT grunt out to collect a \$15,000 desktop workstation.

Consistent performance

One of the classic benefits of data center-hosted desktops is that the desktop applications in the data center always run at full speed, and they always have LAN-speed connectivity to the data sets they require, regardless of what's happening on the client device. So a user's junkware-laden laptop no longer negatively affects the performance of critical business applications.

Better resource utilization

One of the downsides to running high-end engineering apps in the physical desktop world is that each user requires a high-end workstation on his desk. So you might have one user rendering some intense visualization that will max out his workstation for six hours. Meanwhile, the person sitting next to him is using only 2% of his workstation to check email. Wouldn't it be great if the excess capacity of the second user could help the first user's rendering happen faster? Not only does that make the best use of resources (since you already spent the money to buy all those workstations), but it also means that the first user's rendering might finish in three hours instead of six. Talk about a productivity boost!

Data center-hosted virtual desktops make this a reality. Since each desktop is just a virtual machine running on big hardware, resources are shifted as needed, automatically, without anyone's input or knowledge. Users are happy because everything they do happens faster, and the business is happy because the organization is getting its money's worth from its hardware.

Allow users to suspend and resume their desktops

The key to data center-hosted desktops is that the desktops and applications run in a data center, not on the users' client devices. So if you have a user using a laptop, he can close the lid and go offline while the remote desktop session stays running. The user can log back in from another location to check the progress from anywhere.

Engage in “follow the sun” roaming

If you have engineers working around the world, moving your desktops and applications into your data center means multiple users can share the same resources. If you have 30 users each in North America, Europe and Asia whose desktops are all in the same data center, you can get away with buying only enough capacity to cover 30 users at a time. Compare that with buying 90 \$15,000 workstations!

This is even better if you have applications that are licensed based on concurrent users. If your application allows it, you have to buy only enough licenses for those actually using the application at any given time. And with the prices of some of these engineering apps, you might pay for your entire project right there.

Fewer data centers, thanks to centralization

Today's enhanced remoting protocol capabilities combined with recent advances in network bandwidth mean we can now significantly increase the distance between our users and the data centers that host their desktops.

In the early days of data center-based desktop hosting, we could only get decent performance over a LAN. This essentially meant our users had to be in the same building as our servers. Then, as networking bandwidth increased and high-speed Internet connections became more common, we found we could extend our radius to metropolitan areas. Our users could work from home if they were in the same city as our servers.

But now, again thanks to improvements in remoting protocols, we have technology that can deal with ever-higher latencies. We have the ability to deliver remote desktops across countries and even oceans, with a user experience that is very acceptable.

Taken together, these advances mean we're now able to increase the radius of the individual data centers that host our users' desktops. Imagine an environment with 40 offices in 25 cities in 10 countries. Fifteen years ago, we would have had to build desktop-hosting servers in all 40 offices. Ten years ago, we would have been able to consolidate those down to 25 locations (one for each city), and five years ago, we might have been able to consolidate them down again to 10 locations (one for each country). But with today's technologies, we can serve all of our users in the entire world from perhaps two or three data centers.

This is what we mean when we say we're “extending the radius” of our data centers that host desktops. Fifteen years ago, we might not have saved too much, since we would have had to build 40 data centers around the world. But if we can do that today with two or three data centers, we're talking about something compelling!

Downsides of Data Center-Hosted Desktops and Apps

Like every design option in IT, the decision to move your desktops to the data center involves some tradeoffs and downsides, including:

Users need an Internet connection to use their apps

The biggest downside to data center-hosted desktops is your users will need an Internet connection to get any work done. While it seems like people are always predicting that Internet connections will be ubiquitous, there are still scenarios where people are offline, such as when they're on airplanes or when their home cable modems go out at random times.

Internet connection quality

It's true that Internet connections are always getting faster and more reliable, and that today's remoting protocols can do amazing things to deal with low bandwidth, high latency and packet loss. We shouldn't forget, however, that remoting engineering applications is still going to require a decent connection.

Exactly how "decent" depends on lots of things, like the application's requirements and the location of the user. But the reality is, if you have a user with only one bar of signal strength on a 3G MiFi connection, that user probably isn't getting any work done if all his engineering applications are remote.

Might be a hard to sell to nonbelievers

Another reality is that a lot of folks in IT have been burned by bad experiences with remote desktops, whether they tried Citrix MetaFrame in the 1990s or more general VDI in the 2000s. There's a whole slew of IT pros out there who flat out don't believe that data center-based desktops can perform well with graphically intensive engineering apps. (Of course, they're wrong, but that doesn't help the cause.)

What do you do with your existing workstation investment?

You might decide that moving your engineering desktops to your data center makes sense. But if you do that, what do you do with all of your existing engineering workstations? Sure, you could use them as simple access clients for the new remote environment, but if you just spent \$300,000 on 20 new engineering workstations six months ago, trying to build out a huge server environment now is probably going to be a hard sell. Unfortunately, it's not like you can trade in your desktop workstations for server hardware. So you might have to think about how you would justify the new environment and the timing of your purchase based on what your users have on their desks today.

The Bottom Line

Given the many benefits of moving desktops into data centers, doing so makes a lot of sense in many scenarios — and engineering applications are no exception. Even with OpenGL, DirectX, Linux and UNIX applications, the technologies of today allow us to provide an acceptable experience to users connecting from laptops, workstations, home computers or iPads.

Today's environments are heterogeneous and will continue to be so. It's perfectly acceptable for users to use whatever client device and platform they want while IT provides the critical engineering applications they need from a centralized location, irrespective of the client platform or device.

Data center-hosted desktops and applications are now in use at thousands of companies around the world, and everyone should be encouraged to take a hard look at their own infrastructure to see whether they can join those ranks.

Brought to you by OpenText

When it comes to UNIX desktop virtualization, engineers need secure and reliable remote access to centralized applications and data. In order to shorten their design cycles and reduce time to market, they require an advanced enterprise-wide remoting protocol to enable high performance and productivity. With these enhanced protocol capabilities, companies can consolidate their data centers and extend their radius to all engineers, anywhere in the world.

OpenText Exceed VA TurboX lets your users work and collaborate virtually on UNIX applications from Windows, and Linux desktops anywhere in the world. By removing the limitations and the complexity of traditional remote access solutions, Exceed VA TurboX offers the fastest connection to your business with a uniquely intuitive user experience. Designed for the Enterprise data center, Exceed VA TurboX improves the security, manageability and availability of your UNIX applications.

In recent years, 70% of the world's top 20 semiconductor companies have put their trust in OpenText. OpenText helps companies deliver design tools access to engineering staff all over the globe with speed and accuracy that is second-to-none. It offers an enterprise-class management platform that meets industry security standards, out-of-the-box, while maintaining a highly scalable and fully managed framework to empower administrators with full control of remote sessions.

Exceed VA TurboX can provide remote access application solutions to power your desktop virtualization project that reduce IT infrastructure costs, shorten time-to-market and provide end-to-end secured access to your critical data.

OPENTEXT