

From DaaS to Cloud Desktops

Brian Madden Looks at What's Changed



vmware®



A few years ago, “Desktops as a Service” (or DaaS) was a popular topic. But more recently, people have been talking about “cloud desktops” more than DaaS. This change in terminology has led to confusion about what these terms mean and what’s actually going on. In this e-book, Brian Madden, industry veteran and the lead field technologist in VMware’s EUC Office of the CTO, looks at how the DaaS market has evolved over the years: what’s the same, what’s different, and why the term “cloud desktop” is replacing “DaaS.”

First, Some Definitions

Before we jump into the details, let's make sure we're all on the same page in terms of what DaaS is. For the purpose of this e-book, we'll define "DaaS" as the scenario where a customer pays for Windows desktops or Windows apps that are running somewhere else (e.g., off premises), with the user interface being delivered via some form of remoting protocol (such as VMware Blast).

This broad definition of DaaS means any of the following scenarios are "DaaS" for the purposes of this e-book:

- **Windows Server or Remote Desktop Session Host (RDSH)-based** ("session per user")
- **Windows client-based**, such as Windows 10 ("VM per user")
- **Windows 10 Enterprise multi-session** – Like RDSH but based on Windows 10 rather than Windows Server. This scenario is available only as part of Microsoft Windows Virtual Desktop running on Azure.

Each of these three options could be used to deliver full desktops or individual Windows apps, and they each could be persistent or non-persistent.

The key—again, for the sake of this e-book only—is that "DaaS" is Windows remoting technology where the Windows instances are running off premises.

What Was “DaaS” a Few Years Ago?

When the term “DaaS” was at its height a few years ago, most DaaS providers owned and managed their own hardware, which was either in their own data center or in some kind of co-location facility.

Typically, the customer entered into a long-term contract with the DaaS provider, often expressed in terms of cost per user, per month (e.g., something along the lines of “\$50 per user, per month”). Once the deal was signed, the provider would procure the physical hardware needed to run the environment, install the DaaS platform software to run it, work with the customer to migrate the users and data, and then the customer would be off and running.

The contracts were typically long-term because the DaaS provider had to invest in all the hardware needed for each customer environment, paying for infrastructure up front with the idea that they’d recoup their initial costs over the life of the contract. In many ways, DaaS providers back then were essentially managed service providers (MSPs), building and providing custom VDI or RDSH environments that they hosted.

From a customer standpoint, the DaaS of a few years ago didn’t offer a lot of flexibility. Here are just a few examples of the limitations of this type of DaaS scenario:




Inflexible contract length –

A customer who wanted to get out of their contract early would need to pay some hefty termination fees because the provider had to cover the cost of the hardware purchased for the project.



Limited alterations to the environment –

Customers were also limited in their ability to scale or change the DaaS environment once it was built. If they decided halfway through their contract that they wanted GPUs in their servers, but the physical servers purchased by their provider didn’t have the room for them, the customer was out of luck. No GPUs for them!



Inability to downsize –

Customers who decided that the initial specs per user were too big were similarly out of luck. Because the DaaS provider had already paid for a certain amount of memory per user and needed to recoup their costs, customers were unable to make reductions.

What Was “DaaS” a Few Years Ago? (continued)

Clearly, the DaaS of the last decade was *not* the uber-flexible cloud computing environment of today.

What customers were actually paying for back then was for someone else to build and run their VDI or RDSH infrastructure. It was *not* a cloud offering in the truest sense of the word, though it had many cloud-like characteristics, for example, customers didn't have to worry about the hardware or the infrastructure layer.

Another unfortunate reality of DaaS a few years ago was the significant pressure to get it right the first time and incorporate planning for future growth. Because the DaaS environment, once built, would be used for several years and difficult to change, it was crucial that customers did proper sizing and testing, just like for a traditional on-premises VDI environment. This meant that DaaS projects of the era were complex and long, as the teams performed server sizing tests, load testing, network testing, and more. Back then, moving to DaaS wasn't really any faster than moving to traditional on-premises VDI.



Enter the Cloud

What's changed since those early days? Simple: Now we have the cloud.

Over the past 5 years, we've seen the phrase "X as a service" evolve into "Cloud X." For example, Microsoft now describes Windows 10 management as a service as "Windows 10 cloud management." "SaaS" is becoming "cloud software." Continuing along this line of thinking, it makes sense that "desktops as a service" has evolved into "cloud desktops."

Most people leveraging cloud desktops today do so on existing public clouds, like Microsoft Azure and Amazon Web Services (AWS). These public clouds allow customers to leverage the true cloud promise of easily scaling up, scaling down, and changing hardware specs. The raw infrastructure services that cloud platforms offer are augmented by VDI and RDSH platforms (such as VMware Horizon® Cloud). This augmentation means that customers can use the same delivery and management infrastructure for their cloud desktops as they do for their on-premises desktops, as illustrated in the following three examples.



VMware Horizon Cloud on Microsoft Azure

Microsoft's Windows Virtual Desktop platform is native to Microsoft Azure. With Windows Virtual Desktop, customers are entitled to features like published apps and desktops on traditional RDSH servers, Windows 10 Enterprise multi-session, Windows 10 Enterprise VDI, and Windows 7 with free Extended Security Updates for up to 3 years. All a customer needs to pay for is the Azure consumption.

VMware Horizon Cloud on Microsoft Azure includes support for the features of Windows Virtual Desktop. The service runs natively on Azure, where Microsoft provides the underlying infrastructure and VMs, and VMware adds on Horizon protocols, management, brokering, a control plane, and more. This structure is similar to a traditional VDI or RDSH environment where you pay for the Horizon licenses separate from your delivery hardware, but in this case your hardware is monthly VM rentals from Azure instead of a rack full of physical servers.

The value is that you get the same Horizon experience you're used to (Blast protocol, management tools, etc.), but with the cost advantages and hands-off

Enter the Cloud (continued)



infrastructure that only the cloud can provide. In other words, if you want to change the size of a particular user's VM (larger or smaller) during your term, you can simply log in to your Azure control panel and make the change without talking to VMware or Microsoft. If you want to experiment with on-demand pricing versus reserved instance pricing, you can. And of course, as Azure pricing drops (or as Azure VMs get increased capabilities for the same price), you get those benefits too—as soon as they're available, even mid-contract. It truly is the power (and promise) of the cloud, delivered in desktop form.



VMware Cloud on AWS

On the AWS side, you can use VMware Cloud™ on AWS to essentially extend your on-premises VMware environment into AWS locations, which appear in your VMware management consoles just like any VMware vSphere® host. From the consoles, you can install and build whatever you want, including VMware Horizon® 7, which you can then hook into your existing Horizon environment (or keep standalone).

Even though the infrastructure is a bit different on AWS than Azure, many of the benefits are the same, especially when compared to old-school DaaS. With VMware Cloud on AWS, the customer is in control of the vSphere node they're renting. They can slice and dice it however they want, change VM sizes, and more, while still being able to add or remove capacity at any time, and while getting the current price as costs and offerings come down.



Hybrid Cloud

We've learned that while many companies are interested in cloud desktops, it's difficult to flip from on-premises virtual desktops to cloud-based virtual desktops in one motion. The reasons are often technical in nature, owing to the location of user data, applications that still live in the on-premises data center, or even the geographical location of the end users. In these situations, the concept of hybrid cloud has emerged.


Enter the Cloud (continued)



Hybrid cloud desktops mean running your desktop workloads in multiple data centers. This is typically a combination of on-premises and public cloud locations, but it could mean multiple public clouds or even multiple on-premises data centers. In the past, these deployments were completely independent of each other, with separate image, application, user, data, and security management. They often required different skill sets to manage, and each update or configuration change required twice the effort.

Today, VMware has solved the problems associated with hybrid cloud through the use of a cloud-based common control plane that can be used to manage any Horizon environment from a single pane of glass. As this common control plane unfolds, administrators will be able to monitor any environment, manage images across all platforms, broker users to any data center based on policies, deploy applications to all users in any location, and do so much more—in one motion, from a single interface.





Get Started

Today's cloud desktops are the logical evolution of the DaaS of the last decade. Most of the benefits (and challenges) of designing DaaS back then still apply to cloud desktops now. But today's cloud desktops are clearly better for most use cases. Organizations benefit from having someone else to run the infrastructure that powers their VDI or RDSH environment, while gaining the flexibility to change and evolve over time.

That flexibility is the key difference between old-school DaaS and today's cloud desktops. In the DaaS days, flexibility was trying to get your provider to use 3.5GB of RAM instead of 4GB. Today, flexibility is everywhere! You're free to take advantage of the strengths of each deployment platform. Applications and desktops that, for whatever reason, have to be in your data center can still remain there. If you have a large Horizon deployment that you want to move to the cloud untouched, vSphere and all, you can do it. And, if you want to take advantage of the latest Microsoft Azure-based capabilities like Windows Virtual Desktop, you can host desktops in Azure, too.

Today, you can do it all!

For more information, visit the

VMware Cloud website <https://cloud.vmware.com/>.



Join us online:



vmware[®]

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: VMW-EB-BRIANMDAASSCDKTP-USLTR-WEB-20200312 3/20