

TOP 5 CONSIDERATIONS FOR BUILDING A SUCCESSFUL VDI STRATEGY

GET STARTED >





Virtual desktop infrastructure (VDI) allows users to work anywhere, anytime, and on anything with an internet connection, without compromising workstation performance or security. With many great platform options and service models available, there are solid benefits for IT workers too, including centralized device maintenance, simplified scaling (up or down) as needed, and individual security control.

But where there are options, there must be decisions. Should you choose a cloud-based VDI solution? Desktop as a service (DaaS)? Maybe a managed solution? From a platform perspective, should you choose hyperconverged infrastructure (HCI), composable, or a rack solution? Here are the top five considerations you need to bring into your VDI decision-making process.





CONSIDERATION 1: BUSINESS DRIVERS AND VDI USE CASES

Identifying the use cases for your business is critical, as it lays the groundwork for the other four considerations. You must first analyze your business drivers.

Business drivers

Business drivers are the key factors, inputs, and activities that drive operational and financial results and help you achieve your business goals. They can vary drastically from industry to industry depending on the goods and/or services your business provides.

The question you need to answer is, “What are the key elements that drive my business profits?” These could be tangible assets—salespeople, the number of stores, web ads, manufacturing, managers, applications, designers, and developers. They could be technologies—artificial intelligence (AI), cybersecurity, automation, data analytics, or the Internet of Things (IoT). They could even be relationships, maintaining the trust and confidence of existing customers and business partners.



What are my business drivers?

If you are not sure what the key drivers are for your business, begin with a profitable line item from a recent profit and loss (P&L) report and ask yourself, “What factors contribute to this number?” Dismiss any factors you don’t directly control, like the cost of materials or economic factors.

For each factor you uncover, keep asking, “What drives this?” until you reach an item that has no precedent, typically one based on an operational or capital decision. That last item is your key business driver behind that profit line.

From the business drivers, create a list of the number and types of users that support each of those drivers. Ideally, you want to allocate users into the three traditional user types (task workers, knowledge workers, and power users), which can be challenging because users are rarely created equal. Use the list of user types below to help you discern.

Know your user types



Task workers

- Telemedicine
- Database entry
- Market research
- Client-server and browser-based applications



Knowledge workers

- Office work
- Remote learning
- Customer service
- Collaborative groups
- Electronic health records (EHR)
- Back office applications



Power users

- Computer-aided engineering/design
- Financial trading
- Geophysical analysis
- Research
- Software development
- Media and entertainment

Use cases

Your use cases connect your key business drivers to the type of infrastructure required to support them. It is common to find a number of different use cases across an organization, based on the number and the type of users.

Now that you have a rough idea of the number and types of users you support, build your use cases by revisiting that list with a focus on the specific technology needs behind each use case, such as virtual desktop, virtual application, GPU accelerated, persistent, or non-persistent. From this you can characterize the expected workload compute and storage relative to the end-user requirements.

Be aware that use cases, provisioning models, and virtual machine (VM) profiles all require different resources from a compute node and storage perspective. Understanding these requirements lays the groundwork and sets expectations for the overall scope, utilization, and performance.

Some users may require persistence. Others may not. Others still may benefit from application virtualization, GPU, or a combination. Determining use cases and characterizing the workload are critical for understanding the hardware/software stack necessary to support it.



An example

A VM consisting of two vCPUs and 8 GB of RAM might achieve a density of 100 users per compute node relative to vCPU-to-core overcommit. However, a VM with four vCPUs and 8 GB of RAM may only achieve 50 users per compute node relative to vCPU-to-core overcommit.

Note: Never overcommit RAM.

CONSIDERATION 2: WORKLOADS

With your business drivers, use cases, and provisioning models established, the workload will begin to take shape. Virtual workloads are defined by VM profiles which encompass vCPU, RAM, the application stack, image size, persistent versus non-persistent, and/or vGPU requirements. An accurate VM profile is the basis for all node density projections and is critical for establishing node equilibrium with respect to hardware requirements, utilization, and performance.

Your VM profile can be determined in two ways.

1. **Do a full physical, virtual desktop, and application assessment.** This process usually requires a consulting engagement and trended data collection using workload assessment tools. This is a great path if you have the time. The process will:
 - Assess physical and virtual requirements
 - Define use cases
 - Determine feasibility
 - Set density expectations
2. **Run a proof of concept (POC).** Develop a model based upon real-world users with specific user images and use cases relative to the defined VM profile expectations. This will:
 - Define test/success criteria
 - Test real-world images across subsets of users specific to individual use cases across the organization





CONSIDERATION 3: SCOPE AND SCALE

Understanding scope and scale in conjunction with your use cases and workload requirements establishes the ground rules for platform selection. Some platforms and architectures perform and scale better than others relative to workload or use case, which impacts hardware selection.

Scope considerations

The scope determines how the workload will run across use cases and how many users the solution will be required to scale to. For instance, the economics begin to shift in favor of more traditional compute/storage architectures at roughly 2000 users and above, so the number of concurrent users dictate platform selection.

The factors that play into scope include:

- Single or multiple organizations
- Initial footprint versus end footprint
- Global or regional users
- Centralized or decentralized (distributed) resources



Scale considerations

Scale goes hand-in-hand with scope because you need a solid understanding of how the deployment will roll out. Will it be a single site, multiple sites, or global? This determines factors such as WAN versus LAN access, and local compute resources versus regional compute resources.

Beyond the initial deployment, it's also essential to look ahead. Try to define what the solution is going to look like 6–18 months from now. It's important not to paint your solution into a corner because the platform doesn't scale.

The factors that affect scale include:

- Initial and end footprint
- How will scaling occur?
 - Phased user
 - Phased app
 - Phased site

CONSIDERATION 4: PLATFORM-TO-WORKLOAD REQUIREMENTS

Your platform is dictated by the use case, workload, and scope information. Once these variables are firmly understood, the right platform choice tends to present itself.

Many VDI platforms are available to match with specific use case, workload, scope, and deployment models. If you have already decided upon a form factor—traditional rack-based or blade-based servers, converged infrastructure, HCI, disaggregated HCI, or hosted desktop infrastructure (HDI)—then you have already narrowed the options down to a good starting point.

If you have specific use cases where resource contention is a concern, consider an HDI solution. HDI can provide a physical 1:1 resource for cases requiring high levels of system responsiveness, such as financial services, traders, light/medium engineering applications, and design and content creation.

For data-intensive workloads like AI or deep learning, consider GPU-accelerated solutions such as those offered by NVIDIA®. Be aware of any centralized storage requirements.



Private cloud

- Traditional VDI
- Hardware and software on-site
- Purchase-based/CAPEX

Public or managed cloud

- DaaS
- Hardware and software on- or off-site
- Subscription-based/OPEX
- VDI as a service

Hybrid cloud

- Combination of both private and public cloud
- Hardware and software both on- and off-site
- Both subscription- and purchase-based

CONSIDERATION 5: DEPLOYMENT MODEL OPTIONS

The final step is determining which VDI deployment model will ultimately work best for your business, organization, and budget. There are three deployment models to consider.



Private cloud

Hardware and software are located on-premises. This is a traditional model that can be CAPEX or OPEX based. It is of particular interest for highly regulated industries like government, healthcare, financial, or any workloads that require extra emphasis on governance, security, or intellectual property concerns.



Public or managed cloud

Public or managed cloud models provide desktop as a service (DaaS) through a subscription model or a pay-as-you-go model based on consumption of resources. It is suitable when you have predictable computing needs—such as communication services for a specific number of users—but also offers flexibility to scale instantly in response to varying peak demands.



Hybrid cloud model

In the hybrid cloud model, deployment spans private, public, or managed cloud environments so you can get the benefits of both public and private clouds and still take advantage of the existing architecture in a data center. The inherent flexibility of the hybrid cloud model makes it a good choice for mixed workloads. You can allocate dynamic, frequently changing workloads to public cloud for easy scalability and the more predictable or sensitive workloads to private cloud or on-premises data center.



CLOSING THOUGHTS

To make your VDI strategy a reality, desktops, applications, servers, storage, and networking all have to come together. Making the right decision begins with these five considerations. You may be partial to one platform or another based solely on the number of employees you need to support or the structure of your business, but without a well-thought map that identifies actual workloads, you could end up over-allocating resources and then overspending to make it right. To get it right the first time, remember this:

- Start with a thorough and accurate assessment of your organization's operations and identify the key elements (i.e., business drivers) that produce revenue.
- Next, consider all the employees behind each business driver and identify their user type and any specific technology needs they require. Collectively, this information translates into workloads from which you can assess your overall physical and virtual requirements.
- Don't forget to account for the scope—how many simultaneous users need to be supported—and the scale, which anticipates the user load in the coming months and years.

When you're armed with this information, you will be able to match your actual needs to the optimal VDI platform and deploy to a cloud model that securely and affordably serves your needs for years to come.





NEED MORE HELP?

Planning out your VDI strategy can be a daunting task for sure. And in some cases, it is best to call in the professionals. HPE is here to help you, providing expertise and experience, technology and partnerships, as-a-service delivery, and cloud economics, as well as financial services to help you build the ideal desktop virtualization solution for your business.

Let us help you build your agile workplace.

LEARN MORE AT

hpe.com/solutions/desktopvirtualization

**To find out more
information contact:**
sales@strawberrygt.com
02089731414
www.strawberrygt.com



© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All third-party marks are property of their respective owners.

a50002220ENW, July 2020