



White Paper

AEC Centralized Desktop Solution

When making a decision on what is the best solution for your environment, the end result is very dependent on user acceptance. In the realm of architecture, design and engineering, efficient and streamlined workflows are essential and a predictable and performant system is critical. A centralized desktop deployment approach offers significant advantages for firms, enabling them to optimize resource utilization, enhance security, improve collaboration and most importantly give the user the expected and required experience to effectively get the job done. Let's delve into the key aspects of this strategy.

What Is Centralized Desktop Deployment?

Centralized desktop deployment involves migrating physical workstation computer resources from the deskside to the datacenter. With this, managing and distributing software applications and operating systems from a central and controlled location allows for greater flexibility in both management and user experience. Instead of having machines across a locked office, across a cube farm, or at a user's house - installing software, updating software, and maintenance are centralized, ensuring consistency and ease of maintenance.

For users, access to data is quicker since it resides next to the servers and with the appropriate connecting software to the end user's desk, using a remote access protocol, access to the full resources of the system can be given to users at their desk, at their home, on a jobsite, or a customer facility. A flexible work-from-anywhere environment has been created.

The Problem

The issue can be stated in two ways.

1. I have been spending too much time trying and failing to get a stable virtual GPU system. The only answer I have is to lower the density of my VDI system and increase the cost per user. Everything works well, until it doesn't, and I have to go through the process again.

2. I don't like the workstations scattered around the office or worse yet in people's homes, I don't want use VDI, but I don't want to face the challenge of setting up architects, designers, or engineers on a virtual machine.

VDI systems are great and if you have the time and resources to correctly balance them, the experience for users can be excellent. The issue is getting to that point and staying there when you introduce high-end users into the system. Often, users' acceptance is an uphill battle. And in that fight when it comes between satisfying the user and the needs of IT, it is easier just to declare victory and give the user workstation just to keep business moving.

The Solution

Although the VDI portion of the question is one that is constantly evolving, for the AEC world, the preferred answer still a dedicated workstation. A fixed, predictable experience that is not affected by another user demanding more CPU, GPU, or disk resources. Where there is no bottleneck on being able to complete a task other than where the user has to sit and their network connection.

The one item the VDI world does very well is the remoting protocols – being able to take what is being rendered on the backend and delivering that screen to a user a hundred feet away or a thousand miles away. The rest of this document will quickly describe a validated setup and provide some comparison

The Configuration

The Machine

For the AEC machines, the important part of this is processor speed (single core) and to a much lesser extent the GPU. Many of the workstations we see running are focused on Xeon based processors that offer a lot of cores, some speed, and are good at data pushing, however, what we've seen is the best performance is still on the Core based Intel processors that give you fewer cores, but higher performance per core. For ClearCube we focus on our A-Series Blade PC form factor to allow for i9 processors and NVIDIA GPUs up to the RTX 4000ADA. This system gives us density and consolidation benefits in the data center when the high-end user count grows beyond 15 or so users.


- ClearCube A6112
- Intel i9 13900 CPU
- NVIDIA T1000 GPU (** NO vGPU subscription – standard NVIDIA drivers)
- 64GB RAM
- 1TB NVME storage
- 2.5GbE network connection

The Software Stack

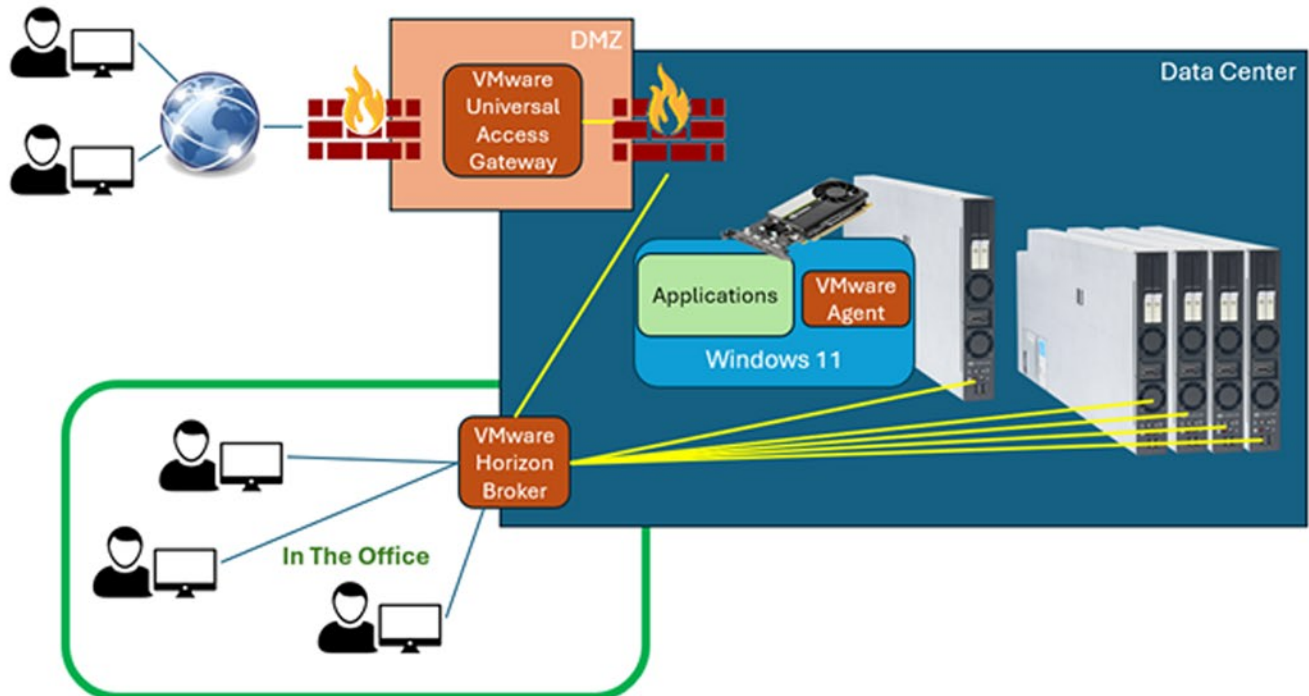
Typically, customers will provide their own Windows stacks for the system. Since this is standard Intel chipsets and NVIDIA cards, there typically is no special configuration that needs to be done to support the hardware.

For the remoting protocol, a variety can be used, but document covers using VMware Horizon licenses just for the protocol and brokering components. Although this can be placed alongside the VDI stack for other users, this will only cover the connections to these physical boxes.

VMware Horizon Apps Standard or VMware Horizon Standard Plus offers the needed components: VMware Horizon Agent, VMware Horizon Connection Broker, and the Universal Access Gateway (if connecting across a firewall).



The Architecture



Typically, customers will provide their own Windows stacks for the system. Since this is standard Intel chipsets and NVIDIA cards, there typically is no special configuration that needs to be done to support the hardware.

For the remoting protocol, a variety can be used, but document covers using VMware Horizon licenses just for the protocol and brokering components. Although this can be placed alongside the VDI stack for other users, this will only cover the connections to these physical boxes.

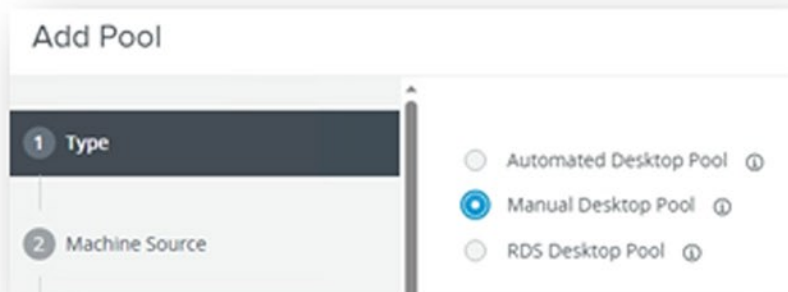
VMware Horizon Apps Standard or VMware Horizon Standard Plus offers the needed components: VMware Horizon Agent, VMware Horizon Connection Broker, and the Universal Access Gateway (if connecting across a firewall).

The Setup


Connection Broker

There is little difference in setting up a BladePC VDI VM versus a standard ESXi configured VM from the connection broker's perspective.

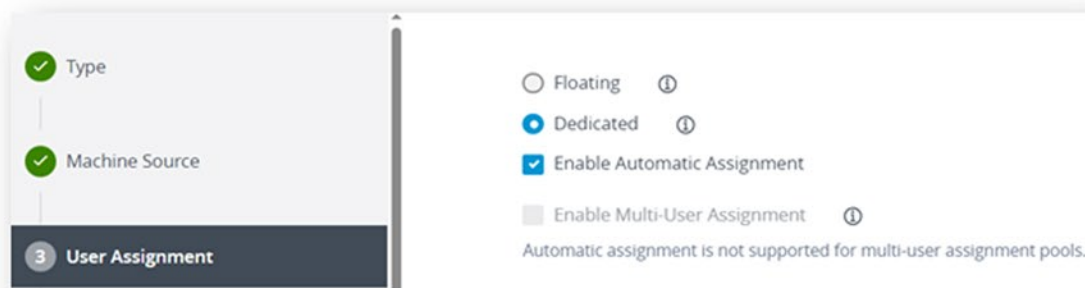
1. The pool will be a Manual Desktop Pool



2. The source is from Other Sources machines.



3. The type of user assignments will allow for setting up a pool of machines to allow for first come, first served, as dedicated for a traditional 1:1 relationship, or as Multi-User where multiple users are dedicated to specific machines – but use the machine one at a time. In most cases, the idea is to have a machine dedicated to a worker. However, there are cases of expensive and underutilized software; this may be an oversubscribed machine or simply floating.





4. For the Desktop Pool Identification, set the ID and Display Name to something significant. If you are using Access Groups, you can specify that here as well.
5. In the Desktop Pool Settings, the Session Types will be desktop and you can determine power control, after disconnecting, and restart policies.
6. The Remote Display Protocol should be set to VMware Blast and the 3D Renderer should be set to Manage using vSphere Client. Since Blast is being used, you can Allow Session Collaboration if needed.
7. On the Machines choice, simply add wanted machines from the list presented. At least one machine must be available – others can be added later.

Multiple pools can be set up for different types of users. Virtual Center machines can be added to include VMware managed virtual machines as well.

After assigning the machines and adding the appropriate entitlements, the system is ready for use. This can be folded into your existing VDI strategy or run separately. Considerations must be made on update and backup strategies as well as maintenance.

Physical Management

With Intel vPRO, the out of band management is complete. All the basic functions are available by default and can be accessed with standard AMT management tools. The Intel Manageability Commander does require the TLS certificates to be signed, but in this example, we can show the same functionality as the previous product – the Intel Mesh Commander.

On the 12th and 13th Gen board, the Intel AMT management is done straight through the BIOS screens. Setup is under the Advanced menus under AMT Configuration then choose MBEX. From there the standard AMT configuration instructions apply of setting a password and configuring the network. The one function you do want to make sure is disabled is set User Consent to None since this is a headless machine.

With AMT you have access to the following and more:

- detailed hardware information,
- system status,
- remote desktop which gives access to interactive desktop control with the on-board graphics controller (not PCIe graphics),



- remote power control including
 - reset, boot to bios,
 - boot to remote drive with IDE redirection
 - boot to PXE or boot to HTTPS
 - power state check
 - graceful shutdown, hard shutdown, system reset

Testing Results

The comparison of show here is between the BladePC in standard configuration accessing it via VMware Horizon and a configuration for a tuned vGPU enabled VDI configuration.

The testing is broken into two parts. The first is a general comparison with the Passmark Performance Test using its Standard Test Suites. This will give a good overall feel of the system doing generic office type task, single threaded numbers and comparing usage of the GPU. The configurations are from the BladePC will all cores enabled, the BladePC with only the Performance cores enabled, and the Dedicated BladePC VDI configuration.

The second test is a Revit File open test under both Revit 2019 and Revit 2023.

Machine Differences

Naturally the same base for each test would make the most sense, but since the original tests for this was to get a feel for the improvements, the data gathering was to make an apple to apples comparison and not Fuji apple to Fuji apple comparison. To note where some of the differences may impact performance, they are listed below. The reader can draw their own conclusion from the differences.

The base system is as stated originally:

Blade PC

- ClearCube A6112 Blade PC
- 2.5Gbs NIC
- Intel i9 12600 processor (3.3GHz base, 4.8GHz turbo)
- NVMe Gen 3 512GB M.2
- 32GB DDR4 2400MHz RAM
- Quadro RTX 4000 GPU

Traditional vGPU VM

- AMD server with AMD EPYC 74F3 processor (3.2GHz base, 4GHz turbo)
- 40 Gbps NIC
- SAN storage connected to 40Gbps networking
- No over provisioning
- Virtual Configuration
 - 6 Cores
 - 1 Gbps network
 - 32GB RAM
 - NVIDIA A40-2Q vGPU

With previous test, we have seen that the bottle neck is not in the GPU, network, or storage once you get to a certain point. The CPU is where the current bottleneck exists. With that in mind, results can clearly be compared in the tests below.

Passmark Performance

The performance test runs the system through what a typical set of task users may do and tries to do a test across the different parts of a system. This is a good comparison of two systems and in general the percentage increase or decrease in the score number directly relates to relative speed of two systems.

With the direct BladePC we see a total Passmark score of 6662. This is a decent number with respect to the processor installed in the system. We can see the CPU is at 36335 which is at the time of this paper in the 92nd percentile. For the CAD systems we are more focused on the CPU Single Threaded number which is 4512. Other numbers can be adjusted with faster drives or newer GPU but for a CAD system the single threaded number is good.



Taking the same test and running it with on a virtualized configuration we see there is significant difference between the server and a dedicated machine. The Core processor from the blade is just simply faster. This allows for better processing of single threaded actions that is critical in AEC CAD work.



Revit Open Performance – Real world testing

One of the intensive moments of Revit is the file open. With that there are two independent times. The first open which builds a cache, and subsequent opens for your normal daily operations. This test compared the open times of a medium size Revit project. This model is of a US middle school that is two stories with 25 education rooms plus the standard school amenities. The first test is opening the two versions of this – one in Revit 2019, the second version is Revit 2023.

Across a sample of 11 opens (one initial and 10 subsequent) we have generated the following chart.



	Revit 2023	Revit 2019
First Open BladePC	23.0s	77.0s
First Open Traditional VDI with vGPU	39.0s	140.0s
BladePC	17.1s	57.2s
Traditional VDI with vGPU	38.1s	125.0s

With this we can see a difference does exist in favor of the dedicated Blade PC. But the difference in is about 44% for the opens. Compared to the Traditional VDI configuration there is substantial savings.

Revit Project Upgrade

Although within a project, converting from one version of Revit to another is not common, it is a time-consuming task that gives a feel of what Revit's needs are in other parts of the application.

In this test, the project was to go from Revit 2019 to Revit 2023 and compare the times of the systems.

Whereas the test for the opens were in seconds, the upgrade is registered in minutes.

Instance	Total
BladePC	21:08
Traditional VDI with vGPU	37:41

Like the Revit Open test, the Blade PC is substantially faster. The Traditional VDI with vGPU pushes the user outside of what would be an acceptable amount of time. In other words, the traditional VDI would force the use consider calling it a day verse continuing to work for the day

Conclusion

With this method and easy to find/use tools, a side by side comparison can be made when researching protocol, devices, and changes to your environment to help determine real numbers your VDI or CDI environment is delivering.