



CLEARCUBE



ClearCube White Paper

Best Practices Pairing Virtualization and Centralization

Increasing Performance for Power Users with Zero Client End Points



Introduction

Centralization and virtualization initiatives are being adopted at a phenomenal pace. The tantalizing results of reducing overall total cost of operation, decreasing desktop management overhead, and improving worker efficiencies are driving technology adoptions to put centralization of computing resources and virtualization into practice.

But like anything new, there may be painful learning curves and anxiety-producing implementation agendas that prove daunting to those tasked with getting the initiatives into day-to-day operational modes. Moving from a distributed PC infrastructure, which has been the technology standard for the past two decades, to a fully virtualized infrastructure begs the question that goes through the mind of everyone given the assignment – “how do you do it with minimal risk”?

Do you jump in the “virtualization pool” with both feet, risking a “cold water” shock to the system? Do you worry that your power users may not be satisfied with the new performance profiles? Do you consider leaving power users completely out of the equation and just put noisy, un-managed, large footprint PCs under their desks?

Or is there a way to enter the virtualization pool a little at a time? Transitioning toes first, then legs, then torso to lessen the shock? Is there a way to migrate from distributed PC infrastructures to centralized PC infrastructures to fully virtualized infrastructures smoothly?

Purpose Statement

This white paper’s purpose is to inform those considering virtualization adoption of the potential for using ClearCube’s hybrid VDI architecture for use cases that need virtual machine (VM) flexibility with performance on par with standalone PCs.

More specifically, this white paper explains the complementary role that Blade PCs perform in virtualization projects.

Hybrid VDI is a half-step to full virtualization. Hybrid VDI uses zero clients at the desktop just as in full virtualization projects, but instead of connecting to virtualized servers all the time, users can provision Blade PCs running fat client Windows applications.

Achieving the best performance for power users in a virtualized environment using zero clients is based on the following premises:

1. Dedicated processing resources will perform better than shared processing resources
2. Software aided by specifically designed optimization hardware will perform better than software alone



Benefits of Dedicated Processing Resources

The phenomenal hardware reduction efficiencies offered by virtualization also introduce performance challenges for power users that run 3D and GPU-intensive applications. Virtual servers are shared devices that must schedule their resources to service requests. This means that the response is dependent upon elements that fluctuate and are not steady-state reliable. Therefore, multiple virtual desktop users all compete for resources that need to be scheduled. The larger the community of users, the higher the level of competition for limited resources. In some instances, those resource requests over potentially constricted bandwidth connections are enormous, so the end result is that the end user experience or the responsiveness of the application that the user is trying to use can degrade as the number of users scale upward.

Distributed networks use standalone fat client PCs that have dedicated CPUs and GPUs with a higher total cost of ownership than VDI shared resources. However, dedicated processing resources outperform shared processing resources. Logically, that is easy to explain and justify, even to a non-technical audience. Most people understand that when one person is using a resource, whether it is a printer or a processor, that person gets 100 percent performance return on that resource. Share the resource between two, three, twenty, or fifty people and the return on that resource can only diminish. In the case of shared printers, and processors for that matter, people have to wait longer for the result than if they did not have to share – whether it is waiting for a printed page to queue up or for a delayed screen response. It is a fact of life.

Elements such as local processing, dedicated graphics processors, and purpose-built CODECs, taken for granted under distributed processing environments using standalone networked PCs, are removed from the equation when using virtual desktops – replaced by issues and concerns involving bandwidth constraints, protocol efficiencies, and resource sharing metrics.

So given the agenda to virtualize, what is a network administrator to do if he has power users?

Connection Brokering Flexibility Coupled with PCoIP Performance

VMware's desktop client View brings a unique capability to address performance requirements. View has the ability to broker connections to designated resources. Typically that resource is a server in the cloud, but that resource also can be a dedicated computer.

At the core of VMware is the PCoIP protocol. Understanding how this highly efficient protocol factors into the performance equation will be critical to understanding how the solution can be best architected.

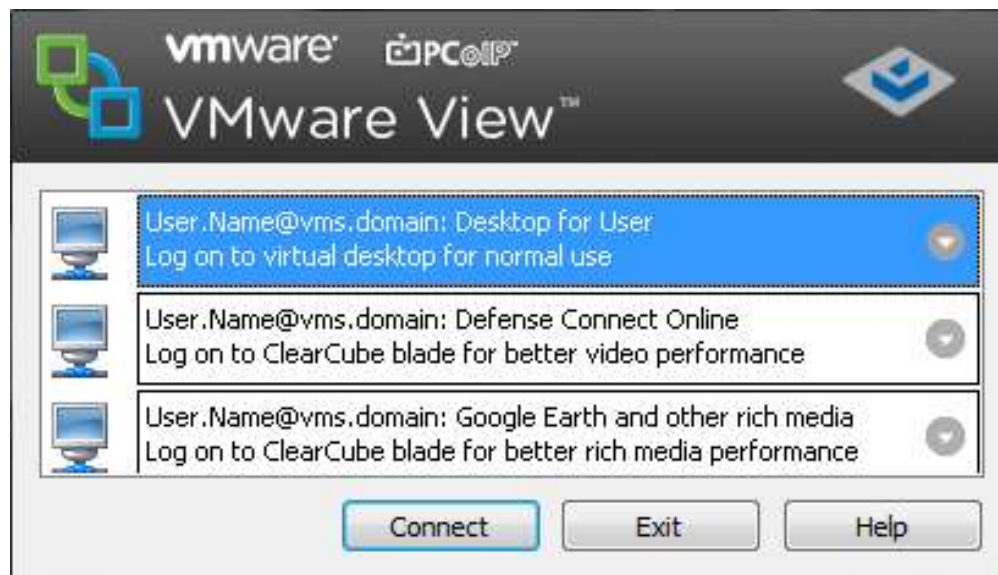


PCoIP is UDP based rather than TCP based. TCP is highly reliable but what makes it so is that it re-tries to send dropped Ethernet packets. This is very important with data (dropping a few zeros will cause heart attacks if you are a bank manager). With UDP, data are not transmitted to the end point; pixel changes are sent, along with USB signaling. Power users typically demand high performance sound and video, so TCP protocols do not add value and just slow everything down by re-transmitting dropped packets. UDP does not perform re-tries because your mind does not perceive the small number of dropped frames and packets. But not having to perform transmission re-tries means UDP is much faster. And fast is what is needed to create a satisfactory user experience for power users.

Hybrid VDI Satisfies Power Users (I:I)

To summarize, PCoIP's underlying core code is native to VMware and is as good as it gets in efficient transmissions across the network. If a use case emerges that requires the equivalent of standalone PC performance, VMware, as previously stated, has a unique View client that ClearCube Technology, Austin, Texas, has put to good use to add a boost using zero clients in virtualization environments.

ClearCube calls it “Hybrid VDI” architecture. The user selects his route to a VM (in the first box) or to dedicated blade PCs (in the middle and bottom boxes). This is referred to as “connection brokering”.

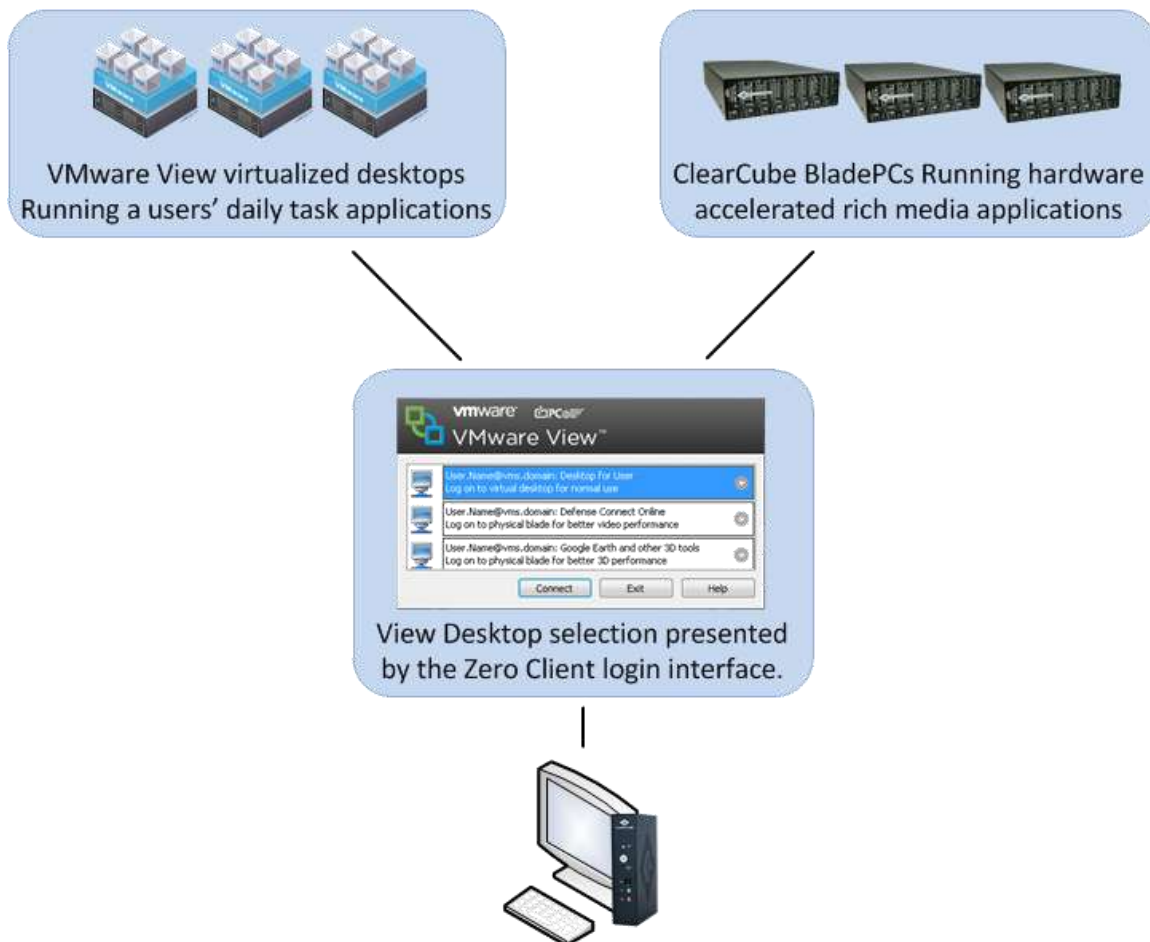




Blade PCs in the Datacenter

ClearCube's flagship hybrid VDI architecture is designed to satisfy network administrators' requirements to centralize computing resources yet still deliver high performance over distance using zero client desktops. Zero clients are stateless with no local memory, storage, or operating system so they are highly rated for their security. Those desktop advantages include small footprint, low power consumption, fan-less/noiseless operation, no resident operating system, memory or storage, and elimination of software patch management at the desktop.

In Hybrid VDI architectures, VMware View may broker the zero client connection on a 1:1 basis to a dedicated computer to run 3D and other robust applications that power users may require such as Google Earth, ArcGIS, ESRI Suite, and FalconView. View also allows the zero client to connect to virtual machines (VMs) running on networked servers. The parameters are defined and determined by the network administrator.





Users that want robust performance utilize their zero client desktop devices to route to fat client blade PCs that deliver the same high level performance profile as the standalone PCs on distributed networks. What sets the ClearCube blade PC apart from the standalone PC is that the blade PC contains a PCoIP hardware compression adapter closely coupled to a dedicated graphics processor unit (GPU) per user that is specially designed to deliver optimized performance over the connection.

In addition, the blade PC utilizes PCoIP hardware based CODECs that are located in the zero clients and on the PC blade host compression adapters. Hardware based CODECs combined with a dedicated GPU deliver far superior performance over software based CODECs because there are dedicated high performance purpose-built chips per user that are devoted to running the I/O instruction sets. Rather than having to schedule demanding I/O requests through a shared server resource, a dedicated blade PC is used to provide dedicated I/O processing for the duration of the connection.

Summary

ClearCube's hybrid VDI capability serves as an investment protection safety net for virtualization adoption with high performance use cases. Like an insurance policy, this capability exists ready to be used when required. It will be reassuring to VMware View adopters that the capability can be applied to the equation at any time to overcome performance slowdowns caused by over-utilized shared hardware resources.

Hybrid VDI provides the compatibility and performance comfort levels associated with distributed PC architecture, along with the many management and ergonomic benefits that centralization and virtualization offer. With this approach, the transition path to virtualization is straightforward for high performance use cases.

Author: Jim Zakzeski

For more information, please contact:

ClearCube Sales

sales@clearcube.com

512-652-3500