



VMware Horizon 6 with View Performance and Best Practices

TECHNICAL WHITE PAPER

Table of Contents

Introduction 3

Horizon 6 Feature and Performance Enhancements 4

 Remote Desktop Session Host Applications4

 Improved PCoIP Performance4

 Hardware-Accelerated 3D Graphics5

 VMware Virtual SAN6

Performance Results 7

 RDSH Sizing8

 Calculating the Number of vCPUs and RDSH Virtual Machines8

 Calculating the Number of Users per Core9

 Calculating the Number of Application Sessions per Core 10

 Comparing Display Protocol Performance 10

 CPU Usage 11

 Bandwidth Usage 11

 Performance Results of PCoIP Default Changes 12

 Hardware-Accelerated 3D Graphics Performance 13

 Light 3D Workload 15

 CAD Workload 15

 VDI Characterization on Virtual SAN 18

Horizon 6 Best Practices 20

 RDSH Virtual Machine Sizing 20

 RDSH Session Sizing 20

 RDSH Server Virtual Machine Optimization 20

 Guest Best Practices for Bandwidth and Storage 21

 PCoIP Settings 22

 3D Graphics Best Practices 22

 Virtual SAN Best Practices 22

Conclusion 23

References 24

Authors and Contributors 25

Introduction

The VMware Horizon® with View centralized desktop infrastructure offers advantages for both end users and IT staff. End users are no longer locked to a particular machine and can access their system and files from anywhere, anytime. Horizon with View transforms IT by simplifying and automating desktop and applications management. IT administrators can quickly create virtual desktops on demand based on locations and profiles. By centrally maintaining desktops, applications, and data, Horizon with View reduces costs, improves security, and increases availability and flexibility for end users.

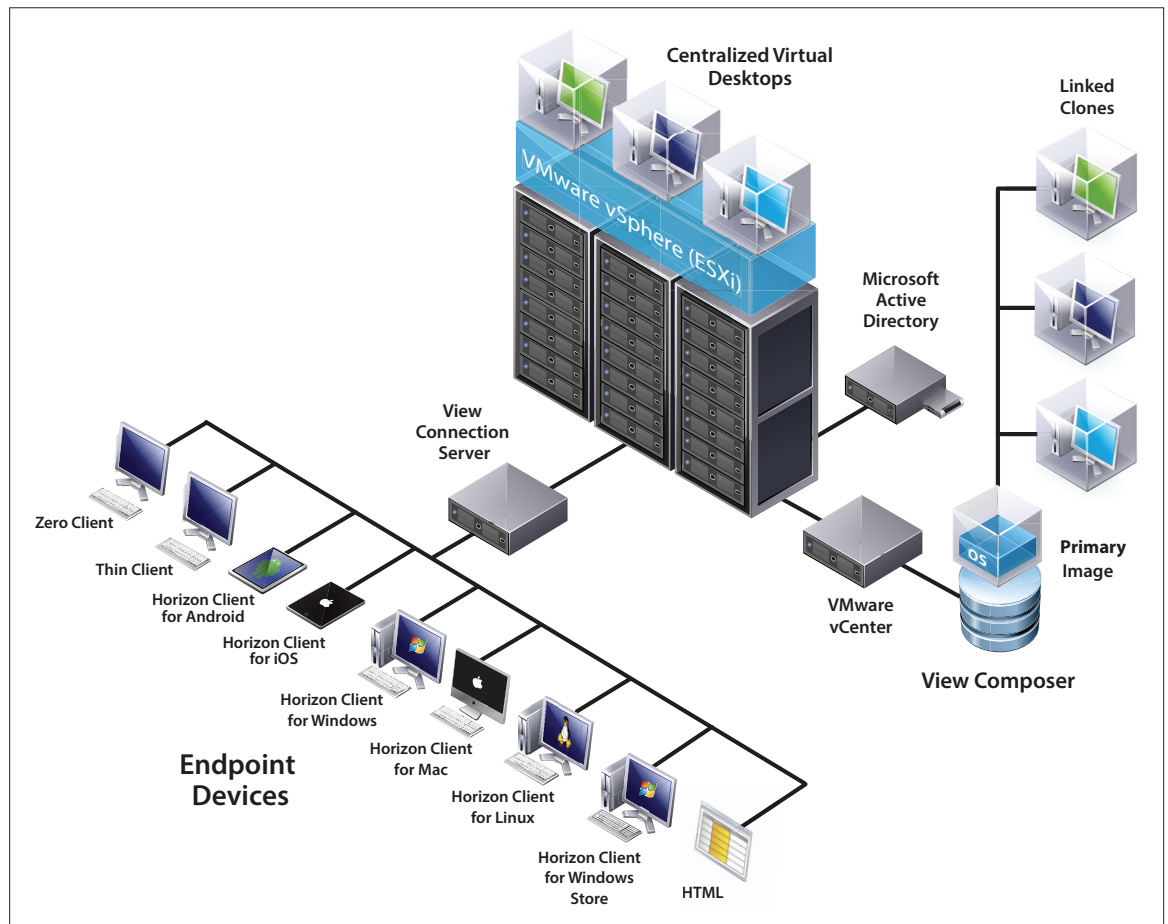


Figure 1: Horizon with View Architecture

Horizon 6 Feature and Performance Enhancements

Horizon 6, which includes Horizon with View in each edition, introduces some new features, such as Remote Desktop Session Host (RDSH) applications and application publishing with seamless Windows support, and enhancements that heighten and improve performance, including adjusting the PCoIP defaults, expanding 3D graphics support, and adding support for VMware Virtual SAN™.

This white paper describes the performance gains achieved with the following Horizon 6 enhancements, as indicated by the testing results. It details the architecture systems used for testing the features and recommends best practices for configuring your system.

Remote Desktop Session Host Applications

Horizon 6 with View introduces RDSH-hosted apps and extended capabilities for RDSH-based desktops. Users can connect to the RDSH server to get a full desktop session or use the published applications seamlessly on the client side.

Improved PCoIP Performance

The PC over IP (PCoIP) display protocol's adaptive technology provides an optimized virtual desktop delivery on both LAN and WAN. Horizon 6 with View improves the end-user experience by introducing new bandwidth management algorithms that increase the frame rate and reduce its standard deviation.

Horizon 6 also has new PCoIP default settings that improve performance. You can change the settings if needed, but in most use cases, the new defaults are suitable and use less bandwidth.

The changed default settings are

- **Build to Lossless (BTL) is set to off** – A fully lossless image is usually not required, and disabling BTL reduces the bandwidth by up to 20 percent. In some cases, however, you might need to enable BTL, for example, when a doctor needs to examine a precise image.
- **Maximum Initial Image Quality changed from 90 percent to 80 percent** – This setting controls the quality and compression level of pixels when the screen changes for pictures, animations, or videos. The previous default setting of 90 percent provided a perceptually lossless level of compression, meaning that the human eye could not perceive any distortion in the image. Because the screen is constantly changing for animation and video, people see no perceptible change at 80 percent of lossless. The new default reduces bandwidth usage by up to 30 percent.
- **Minimum Image Quality changed from 50 percent to 40 percent** – This setting specifies the lowest quality and compression level of changed pixels for pictures, animations, and videos. This limit is reached only when the network is congested and the PCoIP session needs to conserve bandwidth. By decreasing the default to 40 percent of lossless, animation and video frame rates are improved by up to 10 percent under constrained conditions.

Hardware-Accelerated 3D Graphics

In response to user demand for an ever-richer set of applications to be supported in the virtual environment, VMware improved 3D graphics support in View 5.x, with additional improvements added in Horizon 6. The following 3D capabilities expand both the target user base and potential use cases that IT organizations can deliver with virtual desktops.

- **Soft 3D** – Introduced with View 5.0, support for software-accelerated 3D graphics is provided by a Soft 3D graphics driver without physical GPUs installed in the VMware ESXi™ (hypervisor) virtual machine host.
- **vSGA** (Virtual Software Graphics Acceleration) – Introduced with View 5.2, multiple virtual machines can leverage physical GPUs installed locally in the ESXi virtual machine hosts to provide hardware-accelerated 3D graphics to multiple virtual desktops.
- **vDGA** (Virtual Dedicated Graphics Acceleration) – Introduced with View 5.3, a single virtual machine is mapped to one physical GPU installed in the ESXi virtual machine host to provide high-end, hardware-accelerated workstation graphics where a dedicated GPU is needed.

The 3D graphics acceleration is built into the VMware vSphere® platform. Horizon with View fully leverages vSphere and delivers a robust set of 3D offerings to end users. You can administer support for 3D desktops from the View Administrator console. You can enable it on a per pool or per virtual machine basis using the vSphere client. More information and a detailed performance study about 3D graphics acceleration are available in the VMware community forum paper, [VMware Horizon 6 and Hardware Accelerated 3D Graphics – Performance and Best Practices](#).

In contrast to a physical workstation that has sole use of its GPU, GPUs in a virtualized environment are a shared resource. Therefore, it is important to ensure that each virtual machine does not waste the GPU resource. For instance, if View is configured to remote at a lower frame rate (the default is 30 FPS), it usually does not make sense for a 3D application to render hundreds of frames per second (FPS). For these situations, you can configure a View registry setting to limit the maximum application frame rate either in the template virtual machine or on a per-virtual-machine basis. The value is typically set to the maximum frame rate used by PCoIP. Setting the following registry entry for a 3D workload can significantly improve the performance and consolidation ratios achievable:

```
HKLM\SOFTWARE\VMware, Inc.\VMware SVGA DevTap\MaxAppFrameRate
```

VMware Virtual SAN

VMware Virtual SAN is a software-defined storage tier that simplifies and streamlines storage provisioning and management. Horizon 6 supports Virtual SAN 5.5, providing a low-cost storage solution for virtual desktops and RDSH sessions and applications. Virtual SAN allows IT to manage resources and allocate storage on an as-needed basis through storage-policy-based management—the administrator creates storage policies and applies the appropriate policy at deployment or when requirements change.

Virtual SAN offers the following advantages:

- Supports any type of desktop, stateless or persistent
- Reduces the complexity of scaling for virtual desktop infrastructure (VDI) deployments
- Scales to the maximum vSphere cluster size (32 nodes)
- Delivers performance that is almost equal to an all-flash storage system at a fraction of the cost




Figure 2: Virtual SAN and Horizon with View Architecture

Performance Results

To test the performance of the new and enhanced features, we used VMware View Planner 3.5, a workload generator. View Planner simulates typical end-user operations, such as typing in Microsoft Word, playing a PowerPoint slideshow, reading Outlook emails, viewing PDF files, browsing Web pages, and watching a video. View Planner also mimics user behavior by allowing for “think time” during activities. For more information, see the [View Planner Installation and User Guide](#).



Figure 3: Applications Used in View Planner Workload

To simulate a heavy user, all applications were selected, a fast moving video with many screen changes was played, and a think time of 2 seconds was used. For a medium user, all the applications were selected, but a slower video with fewer changes and a think time of 5 seconds were used.

View Planner was run for multiple iterations, with each iteration completing all user operations for the specific group. Each iteration has three phases: ramp-up, steady state, and ramp-down. During each iteration, View Planner reports the latencies for each operation performed within each virtual machine.

View Planner divides tests into groups. Group A represents interactive operations, and Group B includes CPU- and I/O-sensitive operations. Quality of service (QoS) is determined for Group A user operations as 1 second, and Group B user operations as 6 seconds.

RDSH Sizing

With sizing, the goal is to consolidate as many sessions as possible on a particular infrastructure without sacrificing quality. To assess sizing, three aspects of performance were examined:

- How many users or sessions per physical core can a desktop or application session support?
- How many vCPUs are used for an RDSH virtual machine?
- How many RDSH virtual machines are needed?

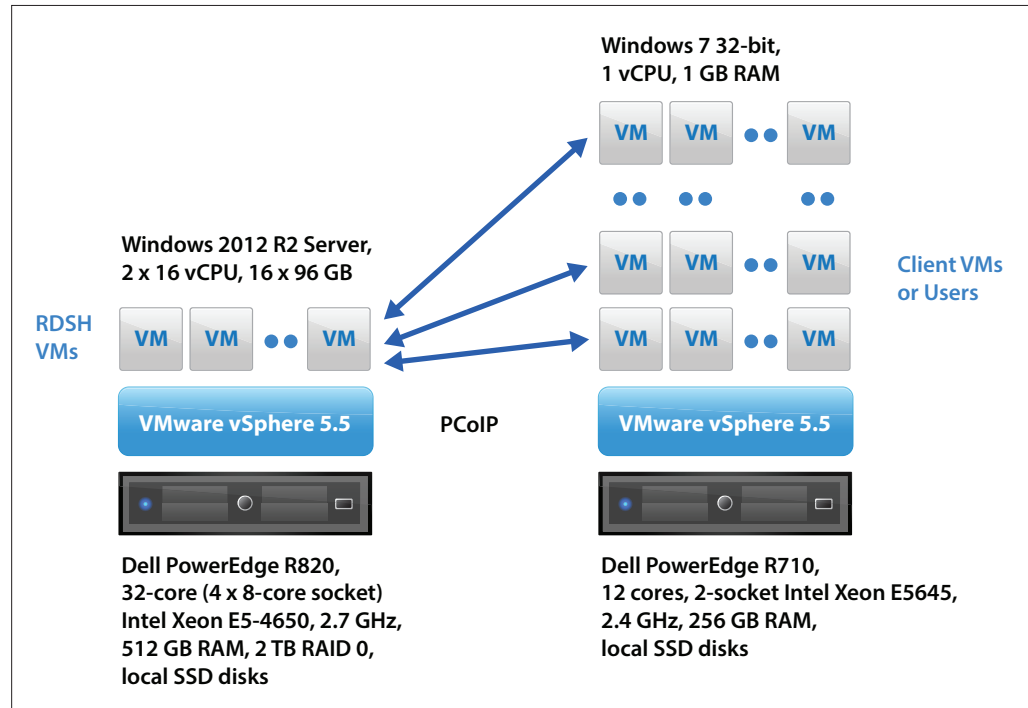


Figure 4: System Under Test for RDSH Server Virtual Machine Sizing

Calculating the Number of vCPUs and RDSH Virtual Machines

One of the main considerations in RDSH sizing is to decide whether to use a few large RDSH server virtual machines (VMs)—12 to 24 vCPUs—or many small RDSH server VMs (4 to 8 vCPUs). To arrive at the optimal number of RDSH vCPUs and instances, we experimented with several different configurations.

The experimental results show that the number of vCPUs for the RDSH server VM should fit within one CPU NUMA node; that is, the number of vCPUs is less than or equal to the number of cores in the CPU socket.

The total number of vCPUs across all RDSH VMs should be twice the total number of physical cores in the system across all CPU sockets. Because most modern processors have hyperthreading enabled, and each physical core has two hyperthreads, you can conceive of this 2:1 over-commitment ratio as a 1:1 ratio by taking hyperthreaded cores into account.

With these two practices in place, essentially two RDSH VMs are being placed in one NUMA node. For more information about RDSH sizing, see the [VMware Horizon 6 RDSH Performance and Best Practices white paper](#).

The system under test included a vSphere server hosting several RDSH virtual machines, each running Windows 2012 R2 Server and each configured with 2–16 vCPUs and 16–96 GB of memory. The RDSH virtual machines communicated with the client virtual machines using PCoIP. Each client virtual machine ran either a remote desktop or a remote application. Each client virtual machine ran 32-bit Windows 7 with 1 GB RAM.

Calculating the Number of Users per Core

Figure 5 shows the optimal number of users (sessions) per physical core that can be served on the RDSH VMs in terms of response time in seconds. Thorough testing by VMware has shown that users can tolerate delay in a remote desktop or remote application action for 6 seconds. The chart shows nine users per core is the point that intersects with a 6-second response time.




Figure 5: View Planner Group B Response Time with Increasing Users per Core

Calculating the Number of Application Sessions per Core

Similar to the remote desktop sessions, all applications were run as remote seamless applications. The number of sessions was increased from 40 to 68 on an 8-vCPU RDSH virtual machine. Figure 6 shows that the density is between 8 and 8.5 users per core, which is slightly lower than the number of desktop sessions, which was about 9 users per core.




Figure 6: View Planner Group B Response Time with Increasing Application Sessions per Core

Comparing Display Protocol Performance

An important aspect of using remote applications is to find the ideal user experience at different network conditions. Horizon 6 with View uses the PCoIP remote viewing network protocol. VMware compared the performance of PCoIP to Microsoft RDP 8 and Citrix ICA. The study looked at two primary factors that contribute to the speed of an application when viewed remotely: CPU usage and bandwidth usage.

The test consisted of 60 remote application sessions running on an 8-vCPU Windows 2012 R2 RDSH virtual machine. All View Planner 3.5 applications were run, except video, a minimum of three iterations with a 5-second think time. The resolution was 1152 x 864, and the color depth was 32-bit. The ICA compression level was changed to low to match the perceptually lossless quality of PCoIP.

CPU Usage

Figure 7 shows the percentage of the guest CPU used when performing the various workload tasks. PCoIP used on average 71.6 percent, RDP used 68 percent, and ICA 71.2 percent. All the protocols used about the same amount of compute power on the guest operating system.

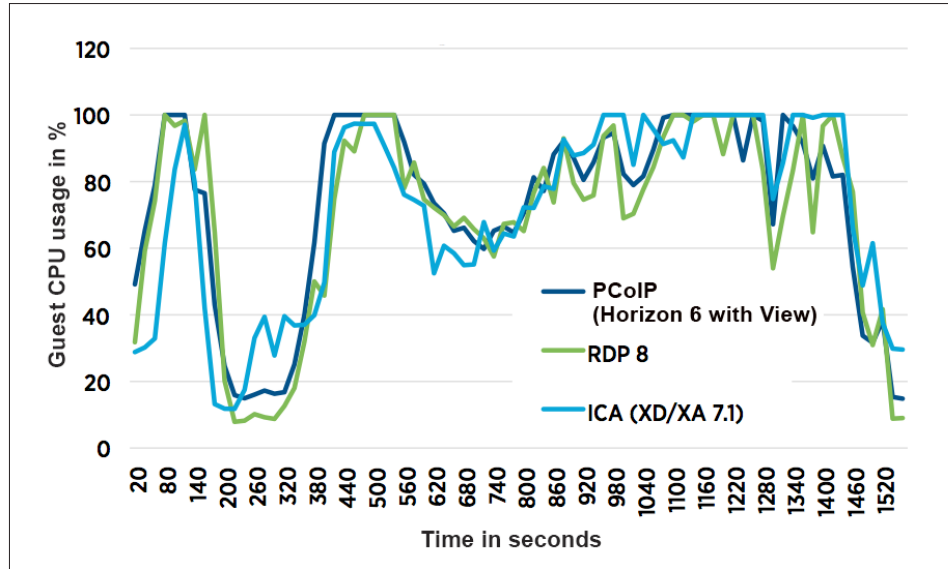


Figure 7: CPU Usage in the RDSH Virtual Machine for Three Competing Remote Display Protocols

Bandwidth Usage

Figure 8 shows the bandwidth usage of the 60 sessions for each display protocol. The average kilobits per second (Kbps) are calculated for the workload runs. PCoIP averaged 44.7 Kbps per session, RDP used 50.7 Kbps, and ICA 48.4 Kbps. PCoIP bandwidth performance is about 10 percent better than RDP and ICA.




Figure 8: Bandwidth Usage for View Planner Medium Workload per Remote Display Protocols

Performance Results of PCoIP Default Changes

To test the effects of changing the PCoIP defaults, we compared the performance of Horizon 6 with View to View 5.3.

Figure 9 shows the setup used for this evaluation. An HP ProLiant BL 460c G6 blade server with a 2.53 GHz Intel Xeon processor and 96 GB physical memory hosted the desktop VMs. The desktop VM data and operating system disk resided on an all-flash array. The VM ran 32-bit Windows 7 with one vCPU and 1 GB of virtual memory.

To run clients at scale, VMs were used on a separate host. An HP ProLiant BL 460c G6 blade server with a 2.13 GHz Intel Xeon E7-2830 processor and 256 GB physical memory was used. Windows 7 32-bit clients connected to the desktop VMs with PCoIP. The clients were configured with one vCPU and 1 GB of virtual memory.

The View Planner workload representing a heavy user was run on all 40 desktop VMs. All applications were tested, and a think time of 2 seconds was used.

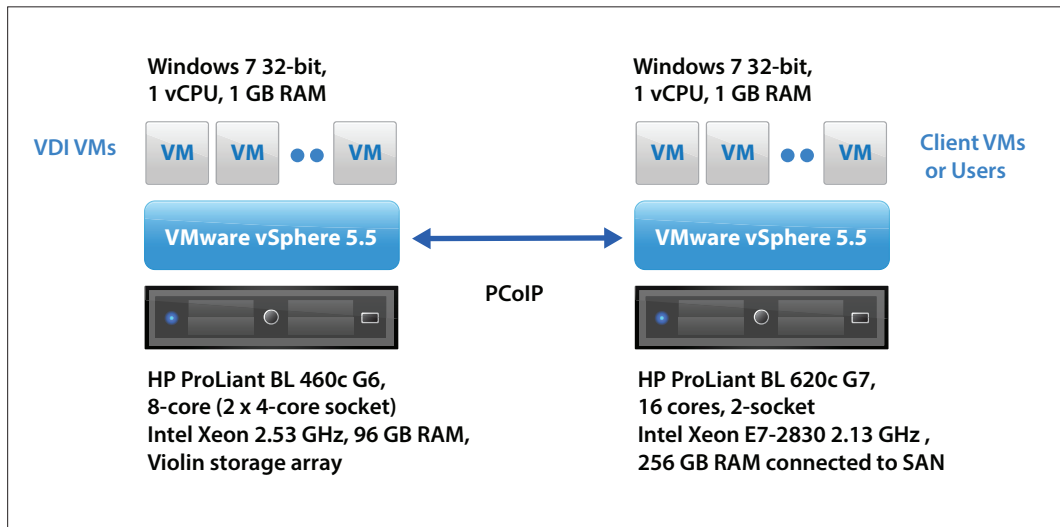


Figure 9: Setup to Test PCoIP Default Setting Changes

Figure 10 shows the response time for Groups A and B, comparing Horizon 6 with View and View 5.3. Horizon 6 with View shows about a 10 percent lower response time in both interactive and I/O-sensitive operations compared to View 5.3. Bandwidth usage was about 30 percent lower in Horizon 6 with View due to changes in PCoIP default settings.




Figure 10: View Planner Response Times Comparing Horizon 6 with View and View 5.3

Hardware-Accelerated 3D Graphics Performance

To measure the scalability of a VDI solution that uses vSGA to support 3D graphics, two different workloads that stress the vSGA solution in different ways were tested. Scalability is defined in terms of the consolidation ratio and the corresponding response time or frame rate during the runs.

The following workloads represent typical customer use scenarios:

- **Light 3D workload** – This workload consists of common desktop applications and usage, including Microsoft Office 2010, Adobe Acrobat, 720p video, browsing static content with Internet Explorer, and running Google Earth in the Chrome browser. All applications are launched at the beginning of the run and remain open for the duration of the run. Throughout the duration of the test, the workload performs a variety of operations using these applications. The ordering of the operations differs from desktop to desktop to mimic real-world workloads. The desktop VMs run Windows 7 at a resolution of 1600 x 1200 pixels and have Aero enabled.

- **CAD workload 2** – A Solid Edge CAD viewer is run in isolation for the duration of the test. During the test, a 3:1 reducer model was used, as illustrated in [Figure 11](#).




Figure 11: 3-to-1 Reducer Model Used in Performance Tests with the Solid Edge Viewer

A think time of 5 seconds was used. All tests used the default settings. The setup is shown in [Figure 12](#). The workloads were run on a single Dell R720 server with different VM consolidation ratios. The number of VMs that can be supported per GPU can be dictated by either the GPU's compute resources being exhausted or the GPU's available memory being exhausted.

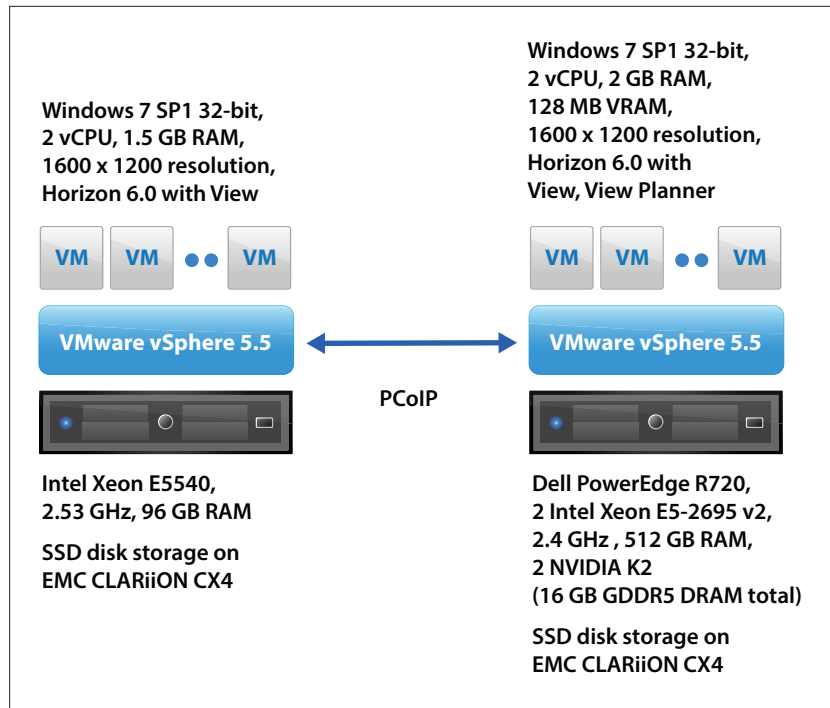


Figure 12: Setup for Measuring the Performance of the vSGA Stack and Horizon 6 with View

Light 3D Workload

For the light 3D workload, the number of VMs was gradually increased until the View Planner response threshold was exceeded. The results are shown in [Figure 13](#). The corresponding CPU utilization, as measured using *esxtop*, is shown by the line graph.

Based on this data, it is clear that the vSGA stack can support 120 users while each user is executing the light 3D workload. Running on higher-performance processors typically delivers even higher consolidation ratios.

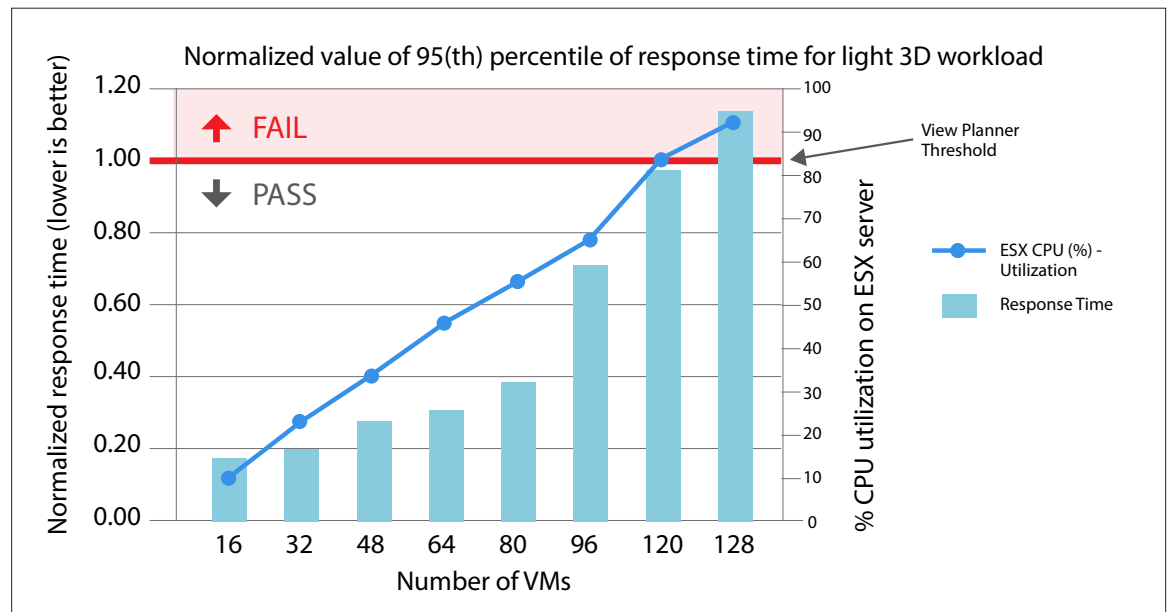


Figure 13: Response Time as the Number of VMs Is Increased

The results shown in [Figure 13](#) were obtained using desktop VMs configured with 128 MB video RAM. Because the test bed had two NVIDIA GRID K2 GPUs with 8 GB DRAM each, the GPUs can support only about 128 desktop VMs. For this light 3D workload, the maximum consolidation ratio achieved on the dual-socket server under test was 128. The test was stopped when the VM responsiveness exceeded the upper limit allowed by the View Planner responsiveness threshold.

CAD Workload

In this workload, a Solid Edge viewer runs a single model: a 3-to-1 reducer, as shown in [Figure 11](#). As with the light 3D workload, the simulated user’s interaction with the model mimics a real user’s pattern. [Figure 14](#) illustrates the scalability of the vSGA as the load on the server with two K2 cards is increased. As the number of VMs is increased from 1 to 72, the aggregate remotely delivered FPS increased by 35 times for K2, and 30 times for K1. Peak GPU usage was about 95 percent.

[Figure 16](#) presents the same results as [Figure 14](#), but shows the remotely delivered FPS data on a per-VM basis. This view of the data highlights that the performance of the individual VMs moderately decreases as the number of VMs is scaled.




Figure 14: Scalability of the vSGA Solution as the Server Load Is Increased with NVIDIA K2 Cards




Figure 15: Scalability of the vSGA Solution as the Server Load Is Increased with NVIDIA K1 Cards

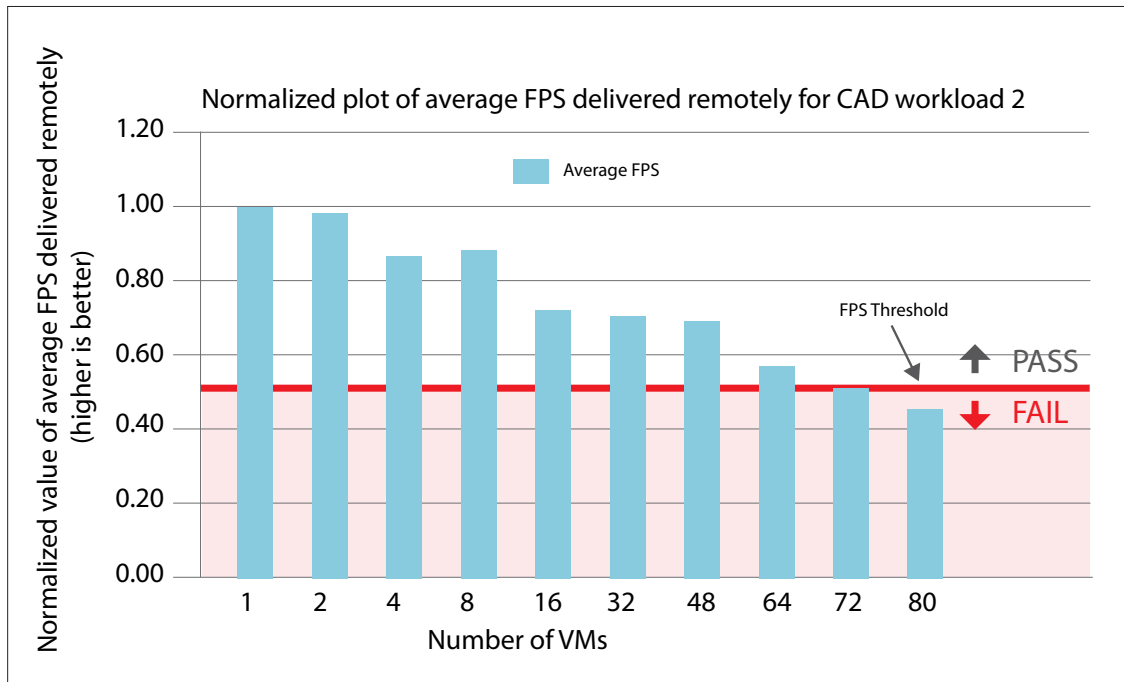


Figure 16: Remotely Delivered FPS Data on a per-VM Basis with K2 Graphics Cards




Figure 17: Remotely Delivered FPS Data on a per-VM Basis with K1 Graphics Cards

VDI Characterization on Virtual SAN

In this setup to test VDI characterization on a three-node Virtual SAN cluster, a host running the desktop VMs has 16 Intel Xeon E5-2690 cores running about 2.9 GHz. The host has 256 GB physical RAM, which is more than sufficient to run one hundred 1-GB Windows 7 VMs. For Virtual SAN, each host has two disk groups. Each disk group has one PCI-e 200-GB solid-state drive (SSD) and six 300-GB 15 K RPM SAS disks. We used View Planner to evaluate VDI and Virtual SAN performance.

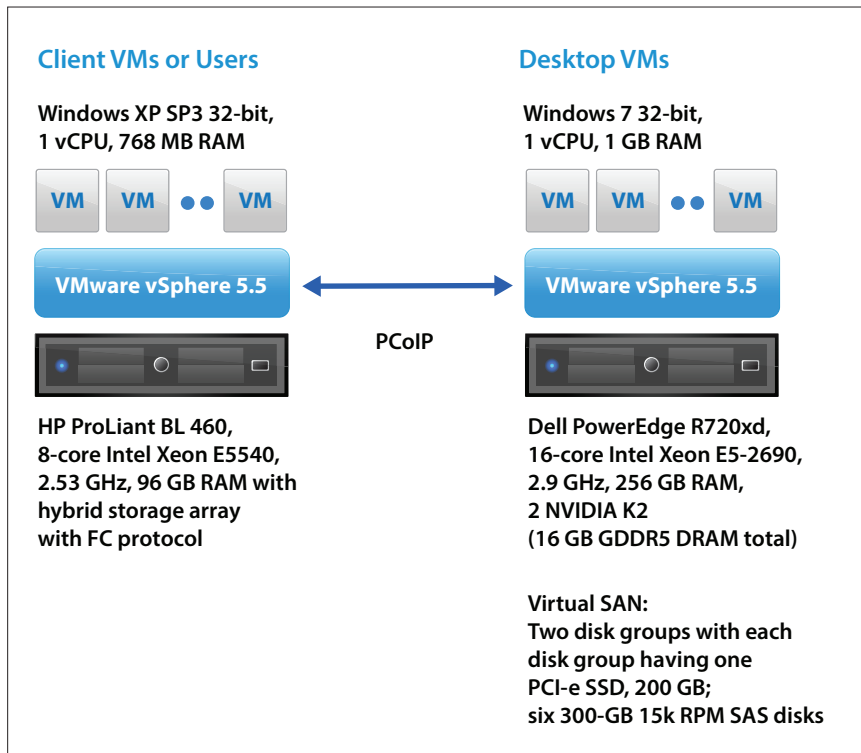


Figure 18: Setup for Virtual SAN (Host Configuration for a 3-Node Cluster)

To highlight the PCoIP default changes, we compared the View Planner score on Virtual SAN 5.5 (vSphere 5.5 U1) on View 5.3.2, and Horizon 6 with View and Virtual SAN public beta (vSphere 5.5). On a three-node cluster, VDI mark—the maximum number of desktop VMs that can run passing QoS criteria—is obtained for both Virtual SAN 5.5 and Virtual SAN public beta. The results are shown in [Figure 19](#).

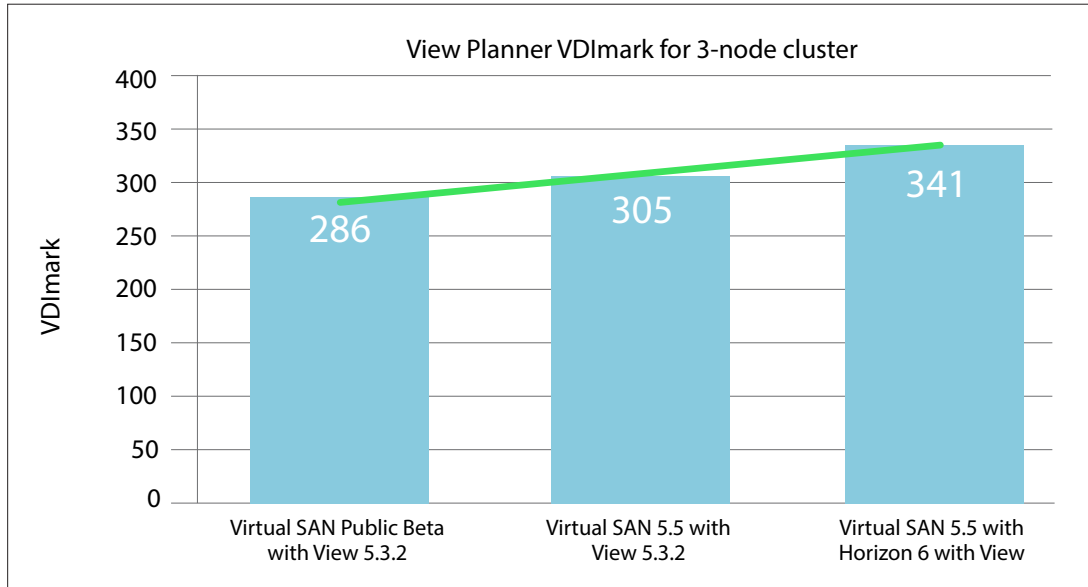


Figure 19: VDI Performance on Virtual SAN

The results show that with Virtual SAN 5.5, we can scale up to 305 VMs on a three-node cluster, which is about 5 percent more consolidation than Virtual SAN public beta. Compared to prior releases, Virtual SAN 5.5 has about 10 percent more consolidation with Horizon 6 with View. This improvement is primarily due to PCoIP default changes, which provide better host consolidation and lower bandwidth usage.

Horizon 6 Best Practices

Based on the performance results, we recommend the following best practices.

RDSH Virtual Machine Sizing

Using multiple RDSH server VMs instead of one large VM provides the best performance. Make the vCPU for the RDSH VM less than or equal to the number of cores in the socket so that the VM fits within the same NUMA node. For the number of instances, a 2:1 CPU over-commitment works better, for example, 64 vCPUs on 32 cores, or 64 hyperthreaded cores, in our experiments.

If the number of cores in the socket is less, and the requirement is to use a large vCPU RDSH VM that does not fit in the same NUMA node, try the **preferHT=TRUE** option in the VMX file of the RDSH server VM. This option forces the vCPUs to the fewest number of physical sockets by giving preference to hyperthreaded cores for scheduling purposes, increasing the probability of the VM staying in the same socket.

You can also set the CPU affinity so that each vCPU is mapped to one set of hyperthreaded cores. Because a 1:1 overcommit with hyperthreaded cores is recommended, setting the affinity in the VMX file can further improve performance. In the following VMX configuration example, an 8-vCPU VM is mapped to cores 0-7. Similarly, the other seven instances of the 8-vCPU VMs could be mapped to cores 8-63:

```
sched.cpu.affinity = "0,1,2,3,4,5,6,7"
```

RDSH Session Sizing

Based on nine desktop sessions per core for typical office workers (medium user) on a 2.7 GHz processor, about 300 MHz per user is needed. For seamless applications, the amount is about 350 MHz. Therefore, setting the CPU MHz between 300 and 500 MHz per session provides a good user experience.

A typical session for a medium user requires the following:

- **CPU** – 300–500 MHz per session
- **Memory** – 400–500 MB per session for nine applications, based on a typical working set of applications (your needs might be different)
- **Disk space** – 200–300 MB per user in the operating system disk for profiles, temporary files, and so on
- **Network** – 50 Kbps per session on average, but also plan for peak bandwidth

RDSH Server Virtual Machine Optimization

To optimize the RDSH server VM image, you can make several registry and group policy changes using the [Horizon 6 RDSH Optimization Tool](#) available on the VMware Community Web site. To view the registry changes, you can run process monitor (procmon) and set the filter on WriteRegistry to see which optimizations are being applied.

Guest Best Practices for Bandwidth and Storage

It is important to optimize the primary VM before the linked clones or full clones are created. Many optimizations can be applied to the guest VM. These optimizations save precious resources, such as bandwidth and storage.

Table 1 shows the settings applied to achieve the best bandwidth and minimize the need for storage resources. More guest optimizations are discussed in the [Optimization Guide for Windows 7 and Windows 8 Virtual Desktops in Horizon with View](#).

PARAMETER	CONFIGURATION
vCPU	One for Windows XP and Windows 7 and Windows 8; two for multimedia-intensive apps
Memory	768–1024 MB for Windows XP 1 GB for 32-bit Windows 7 and Windows 8 2 GB for 64-bit Windows 7 and Windows 8 1.5–2 GB for Windows XP, Windows 7, and 32-bit Windows 8 3 GB for 64-bit Windows 7 and Windows 8 for memory-intensive apps
Network adapter	vmxnet3, flexible
Storage adapter	PVSCSI or LSI Logic SAS
VMware Tools™	Latest installed (also make sure that the balloon driver is functioning properly)
Visual settings	Adjust to best performance (can provide an additional 10–20 percent bandwidth savings in WAN environments) Disable animations for Windows maximize and minimize operations Use default cursor for busy and working cursor
Disable services*	Windows Update, Super-fetch, Windows Index
Group policy settings	Disable hibernation Disable system restore Set screensaver to None
Other settings	Turn off clear type Disable fading effects Disable auto-play and external drive caching for quick release (recommended so that these services do not try to pull USB drive information over the WAN) Disable last access timestamps*
* Especially important to curtail redo log growth with linked clones.	

Table 1: Guest Best Practices to Save Bandwidth and Storage Resources

PCoIP Settings

After optimizing the primary VM, you can adjust PCoIP settings to realize the best user experience. Table 2 lists group policy object (GPO) settings to improve the user experience for WAN environments. But even the default settings save lots of bandwidth. More PCoIP settings are discussed in the [VMware View 5 with PCoIP: Network Optimization Guide](#).

SETTING	DEFAULT	RECOMMENDATION	DESCRIPTION
Session Audio BW limit	500 Kbps	50-100 Kbps	Reduces bandwidth usage of audio with usable quality.
Maximum frame rate	30	Change to 10-15 based on network settings	Provides decent video playback and fast graphics operations in typical WAN conditions.
Maximum session bandwidth	-	Set per network conditions and link bandwidth	Helps PCoIP better estimate bandwidth, and maximizes the user's experience of sharing the same link.
Client-side cache size	250 MB	Set per client-side memory available	Lets you configure the client-side image cache size.

Table 2: PCoIP GPO Settings and Best Practices

3D Graphics Best Practices

Horizon 6 with View and PCoIP dynamically adapt to the available CPU and bandwidth resources to present the optimal user experience. Even when multiple VMs are sharing a single physical GPU, vSphere ensures that the resource is fairly shared between them. As a result, minimal configuration is required to deliver peak performance.

Virtual SAN Best Practices

Get better performance from your Virtual SAN with these best practices.

- Use more, smaller SAS disks for capacity tier to get more IOPS.
- Separate the SSD device access path from the capacity tier path. Use a dedicated SAS HBA for disk-based flash or use a PCI-E or memory slot SSD.
- Minimize the number of components as much as possible per VM.
- Enable View Storage Acceleration.

Conclusion

This paper presented performance data and best practices for VMware Horizon 6 and highlighted some of the new and enhanced features. It describes some of the sizing practices for RDSH virtual machines to get the best performance for RDSH desktop sessions and applications. According to the findings, the vCPU for the RDSH should be configured so that the VM fits within one NUMA node, and the number of RDSH server instances should allow all hyperthreaded cores to be fully utilized. For RDSH desktop and application sessions, 300 MHz per session provides a good user experience for a medium workload.

In Horizon 6 with View, the PCoIP defaults have changed, resulting in a 10 percent improvement in the user experience and 30 percent lower bandwidth usage. Similar results are seen when the desktops are deployed on Virtual SAN 5.5.

For the vSGA feature, the results illustrated the ability of VMware hardware-backed 3D support to scale efficiently, demonstrating the benefits of GPU virtualization and the strength of the VMware 3D strategy.

References

[VMware View Planner Installation and User Guide](#)

[VMware View Planner download](#)

[VMware View Planner: Measuring True Virtual Desktop Experience at Scale](#)

[Comprehensive User Experience Monitoring](#)

[Optimization Guide for Windows 7 and Windows 8 Virtual Desktops in Horizon with View](#)

[Remote Desktop Protocol 8.0 update for Windows 7 and Windows Server 2008 R2](#)

[VMware Horizon View 5.2 and Hardware Accelerated 3D Graphics: Performance Study](#)

[VMware View 5 with PCoIP: Network Optimization Guide](#)

[VDI Benchmarking Using View Planner on VMware Virtual SAN \(VSAN\), part 1](#)

[VDI Benchmarking Using View Planner on VMware Virtual SAN \(VSAN\), part 2](#)

[VDI Benchmarking Using View Planner on VMware Virtual SAN \(VSAN\), part 3](#)

[VDI Performance Benchmarking on VMware Virtual SAN 5.5](#)

[Simulating different VDI users with View Planner 3.0](#)

[VMware Horizon with View](#)

[VMware Virtual SAN](#)

[VMware Horizon 6 Reference Architecture](#)

[VMware Horizon 6 RDSH Performance and Best Practices Whitepaper](#)

[VMware Horizon 6 and Hardware Accelerated 3D Graphics - Performance and Best Practices](#)

Authors and Contributors

Banit Agrawal, Senior Performance Engineer, VMware, is an expert in VDI, View, remote display protocols, VMware View Planner, performance benchmarking, and performance troubleshooting and has filed several patents addressing these areas.

Hari Sivaraman, Senior Member of the Technical Staff, VMware, works on 3D rendering performance and CUDA support for ESX.

Xing Fu, Member of the Technical Staff, VMware Performance Engineering group, works on the performance of virtualization solutions.

Pradeep Chikku, Senior Storage Engineer, VMware, has experience in file systems, storage, ESX kernel storage stack, storage devices, device drivers, and Fibre Channel storage.

Nachiket Karmarkar, Senior Performance Engineer, VMware, has expertise in the area of VDI performance, performance automation, and View Planner.

Rishi Bidarkar, Director, VMware, leads the VDI Performance and View Planner teams. He has filed several patents in the area of VDI performance and display benchmarking.

The authors would like to extend their sincere thanks to Nancy Beckus, technical writer and editor in the Technical Marketing team at VMware, and Julie Brodeur for their comments, feedback, and improvements in the quality of this white paper. The authors also thank the View Planner team and Warren Ponder for their comments and feedback.

To comment on this paper, contact the VMware End-User-Computing Technical Marketing team at euc_tech_content_feedback@vmware.com.

