



TABLE OF CONTENTS

Overview of Infrastructure	2
Deployment and Network Terminology	2
User Workloads	3
Factors Driving User Workloads	3
Typical User Workloads	4
Exceptions to Typical Workloads	5
Server Infrastructure	5
Server Architectures	5
Allocating Server Resources	6
Peak CPU Demands	6
Server Locations	6
System VM Requirements	7
Pano Maestro Requirements	7
Storage Infrastructure	7
Storage Capacity Requirements	7
IOPS Requirements	9
DVM Collections and Storage Choices	10
Network Infrastructure	12
Traffic Sources	12
Network Bandwidth	12
Network Latency	13
Network Ports	14
Peak Network Demands	15
Other Network Infrastructure	16
Client Infrastructure	17
Pano Zero Client Requirements	17
Pano Remote Host Requirements	17
More Information	17

Pano System Infrastructure Sizing Redbook

Planning Server, Storage and Network Infrastructure for Pano System Deployments

Virtualized Desktops hosted on centralized servers promise to radically reduce the costs and staffing demands of increasingly complex desktop computing. Pano Logic® is the first company to offer a complete, purpose-built solution for full native Windows® virtual desktops, combining a unique zero client endpoint with centralized management tools designed specifically for managing virtual desktops.

This redbook provides a detailed overview of how to plan and size the underlying infrastructure of servers, storage systems, and networking hardware needed to support Pano System virtual desktop deployments ranging from a handful of seats to many thousands.

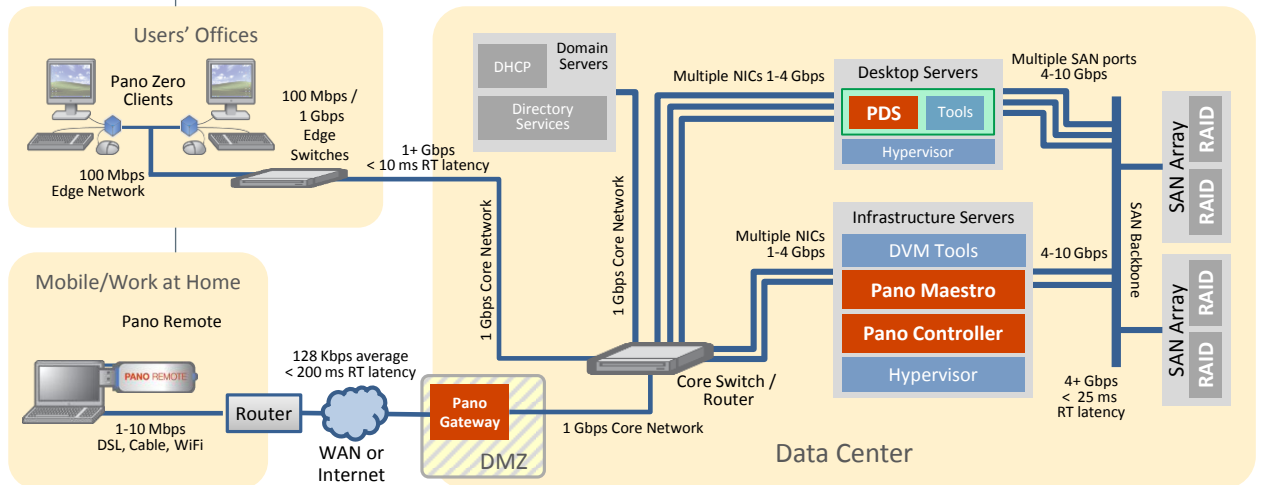
It covers best practices for estimating server hardware requirements based on desktop virtual machine (DVM) allocations, storage hardware IOPS and capacity requirements, and network bandwidth and latency demands.

Pano Logic redbooks assume readers have good familiarity with both their virtualization platform's operation and with the basic structure and operation of the Pano System. They are not intended to replace your platform vendor's documentation or the Pano Logic online help. This redbook is supplemented by the *Remote Deployments Redbook*, which provides more detailed information on network planning, sizing and troubleshooting.

Overview of Infrastructure

This redbook provides information on how to plan for Pano® deployments by assessing the requirements for your Pano System™. This section also provides information about configuring for scalability and redundancy and best practices for deploying a Pano System.

Figure 1:
Pano System deployments and their supporting infrastructure



DEPLOYMENT AND NETWORK TERMINOLOGY

To help understand the terminologies used in this document, below are definitions of several terms used to describe Pano System deployments. Figure 1 above illustrates most of these components and their interconnections.

High-level components in Pano deployment architectures include:

- **Pano Zero Clients** – software-free virtual desktop endpoints that connect users' peripherals (monitor, keyboard, mouse, audio and USB devices, etc.) to desktop virtual machines (DVMs) running on a desktop server.
- **Desktop Servers** – shared servers used to host DVMs and supporting hypervisors. May store DVM images to local direct attached storage or access them over a storage network.
- **Infrastructure Servers** – servers used to host the Pano Controller virtual machine, alternate connection brokers, virtualization platform components for DVM provisioning and management and storage optimization, as well as supporting hypervisors. In smaller remote sites or deployments, these components may be hosted on one or more desktop servers rather than on physically separate infrastructure servers.
- **Management Workstations** – Windows workstations or browser sessions used to connect to management front-ends, including Pano Controller and the virtualization platform software. These workstations are used by IT staff to manage and monitor Pano Zero Clients and DVMs. For Pano Controller and most platform management tools, only a web browser session is required. After initial setup, a Pano Zero Client can also be used as the management workstation.
- **Edge Network** – peripheral networks used to connect Pano Zero Clients and other devices, at user locations, back to the central, or core, network. Edge networks, often called local area networks (LANs), typically use 100 Mbps to 1 Gbps Ethernet connections and switches with higher speed backbones connecting multiple switches and the core network.

- **Core Network** – the central network connecting edge networks to the servers supporting a site or organization. Typically, network components in a data center's core network provide higher throughput and additional types of interfaces than edge network components. For example, a core network router or switch could provide connections ranging from 1 Gbps for servers and edge network switches up to 10 or 40 Gbps for campus-wide connections, as well as interfaces to wide area networks (WANs) or the Internet.
- **Storage Area Networks (SANs)** – dedicated connections from desktop and other servers to specialized storage controllers and drive arrays. Provide high-speed (4 to 10 Gbps), low latency access to shared storage without creating contention on core networks.
- **Domain Servers** – provide directory services, like Microsoft® Active Directory® from Windows Server®, for user and system authentication. These servers also typically provide network addressing resources, like the dynamic host configuration protocol (DHCP) service, which is important for Pano clients to be able to connect to the network.

Software components used in Pano deployments include:

- **Hypervisors** – allows multiple virtual machines (VMs) to run concurrently on a host server. Hypervisors are used to host both DVMs running on desktop servers and system VMs, like Pano Controller or platform DVM management tools, on infrastructure servers. Hypervisors supported by Panos are VMware® vSphere™ ESX and ESXi, Microsoft Hyper-V™, and Citrix® XenServer®.
- **Pano Controller™** – primary management console for Pano Zero Clients and DVMs. Has both Zero Client Controller and Virtual Desktop Broker roles. Can optionally integrate with alternate connection brokers like VMware View™ Connection Server or Citrix XenDesktop® Controller.
- **Pano Maestro™** – a centralized management front-end used to create, monitor, and manage groups of Pano Controllers.
- **Pano Direct Service™** – installed into the DVM's Windows® operating system, it provides connections to the Pano Zero Client and Pano Controller.
- **DVM Management and Provisioning Tools** – various platform-specific tools, often deployed as VMs on infrastructure servers, which are used for automated provisioning, management tasks or optimizing the storage used by DVMs.
- **Client Tools** – platform-specific software installed into the DVM's operating system to provide connections to the supporting hypervisor and platform DVM management tools.

Acronyms used in this redbook as well as by the virtualization platform software from Citrix, VMware, and Microsoft include:

- **DVM** – Desktop Virtual Machine – the virtualized Windows XP or 7 operating system that runs on a server and connects to the Pano Zero Client.
- **vCPU** – the virtualized CPU resources allocated to a specific virtual machine.
- **vRAM** – virtualized RAM allocated to a specific virtual machine.
- **IOPS** – a storage metric called input/output operations per second.

User Workloads

An important part of correctly sizing your infrastructure for Pano virtual desktops is correctly assessing the user and application workloads your desktops must support. These user workload estimates can help you estimate the server and storage resources needed for a deployment's specific mix of users.

FACTORS DRIVING USER WORKLOADS

Key factors to consider in estimating user workloads include:

- Number of applications run simultaneously

- The amount of CPU and RAM resources needed by the applications
- The complexity and rate of change of the application’s display
- Whether rich media like full-screen video will be played
- The resolution and number of monitors connected to the Pano Zero Client
- If local peripherals that cause added traffic back to the Pano Zero Client, like printers or storage devices, are being used
- If peripherals that need consistent low latency (e.g. isochronous USB devices like webcams) are attached to the Pano Zero Client
- The operating system being used (Windows XP or Windows 7) and whether it is a 32-bit or 64-bit version (for Windows 7 only)
- The load on the server from other infrastructure or platform software

For example, Pano Direct Service (PDS) needs a certain amount of CPU time in order to process and transfer the user’s desktop experience over the network to the Pano client. The amount of CPU time needed is directly proportional to the activity on the desktop’s display. Active displays (e.g. video playback, graphically rich web sites, rapidly scrolling through documents, resizing or dragging windows, etc.) cause higher vCPU consumption.

It is important to ensure that the vCPU usage by applications (1 above) does not consume most or all the vCPU allocation, because if it does, then PDS (2) won’t have enough resources available to deliver the experience.

TYPICAL USER WORKLOADS

Despite the possible complexity of estimating user workloads, it is possible to break them down into several typical user workloads. Table 1 provides some guidelines for typical user workloads and corresponding sizings for memory, vCPU and IOPS allocations. These in turn provide general guidelines on the number of DVMs per CPU core and the number of users per server for each workload.

Table 1: Typical User Workloads

Attribute	Light Workloads	Medium Workloads	Heavy Workloads
Application Usage	Task workers running only 1 or 2 applications, for example, a web browser or a billing application	Knowledge workers running multiple applications simultaneously, including Microsoft Office applications	Power users of financial modeling or scientific applications; users viewing full-screen or HD video and rich media
Memory Allocation per DVM	768 MB for Win XP 1 GB for Windows 7	1 GB for Windows XP 1.5 GB for Windows 7	2 GB for Windows XP 2+ GB for Windows 7
vCPUs per DVM	1 vCPU	1 – 2 vCPUs	2 vCPUs
IOPS needed per DVM	Approximately 30 IOPS per DVM	Approximately 40 IOPS per DVM	Approximately 50 or more IOPS per DVM
DVMs per Server CPU Core	Approximately 6 - 7 DVMs per core	Approximately 4 – 5 DVMs per core	Approximately 3 – 4 DVMs per core, but for best performance, allocate as many resources as necessary in the hypervisor.
Average Users per Server	40 users per server	30 users per server	25 or fewer users per server

In general, these user workload factors are the same whether the user is connecting via a Pano Zero Client or via a Pano Remote session although some impacts on performance such as multiple displays won't apply to Pano Remote sessions.

EXCEPTIONS TO TYPICAL WORKLOADS

Exceptions to these typical user workloads can be driven both by the peripherals connected to the Pano Zero Client and by the applications and utilities run within the DVM.

Multiple monitors generally result in the resource needs of both applications and PDS being higher as users with dual displays will tend to have more applications active, thereby consuming more CPU time. They may also generate more display updates as they jump between applications or drag applications between screens. Playing video on one screen while working in Excel, which generates substantial display updates, on another screen is supported, but definitely constitutes a heavy workload.

In addition, applications and peripherals that require low latency (e.g. isochronous USB) will also perform better when there are sufficient CPU resources available. For example, use of web conferencing generally requires two vCPU, while audio-only VoIP can benefit from two vCPUs. Even so, users running web conferencing and VOIP applications may be fine with one vCPU, depending on whether they are running other applications simultaneously.

You can use the Windows Task Manager inside the DVM (as you can on a regular PC) to look at the relative % of CPU time being consumed by specific application or system process, or by the Pano Direct Service process.

Server Infrastructure

This section describes the hardware requirements for servers used in virtual desktop deployments. In general, just four factors influence the density of users per server:

- Processor cores per server and the speed of the cores
- RAM per server, especially on desktop servers
- Available storage system IOPS
- Bandwidth from physical network (NIC) and SAN connections

SERVER ARCHITECTURES

Decisions on server architecture need to be driven by a balance between server CPU capabilities, RAM capacity and bandwidth and network connection bandwidth. For Pano deployments of several hundred or more seats, this suggests deploying a large number of 1U or 2U servers equipped with moderate CPUs, such as dual quad-core or six-core Intel® Nehalem or Westmere CPUs (such as E5620 running at 2.4 GHz or better), 48 to 96 GB of RAM and multiple NICs/ Ethernet ports (and optionally 4+ Gbps SAN) interfaces.

Additional processing capacity can be found by deploying multi-socketed servers with Intel 7500 series Xeon or AMD CPUs, potentially with hundreds of GB of RAM. However, using these larger servers rather than a greater number of smaller dual-socketed servers can both increase per-seat costs and shift the bottleneck from the CPU over to memory capacity, storage IOPS, or network bandwidth. Using fewer servers also reduces the level of availability provided by server redundancy. If rack space requirements are a driving factor, blade-servers can be used in place of 1U servers, although typically at a higher relative cost.

Using servers with earlier, pre-Nehalem processor architectures will reduce the number of seats that can be supported per server CPU core, as the latest microprocessor technology directly incorporates native support for virtualization. Although it might be tempting to employ old servers that might be available after a server consolidation or which might have some unused capacity for virtualization, this approach can result in an unsatisfactory user experience due to slow server performance.

Since any interruption in the servers' operations can at least temporarily suspend the access to virtual desktops, deploying servers with fault-tolerant hardware, such as

redundant power supplies, hot-swap drives and fans, and RAID controllers with battery backup for write-caching, may be a useful investment.

See Table 1 on page 4 for information on the typical capacity factors per server based on specific user workloads.

ALLOCATING SERVER RESOURCES

vCPUs, or virtual CPUs, represent a virtualized CPU core provided to the virtualized operating system by the virtualization platform's hypervisor. A single vCPU allocation to a virtual machine would correspond to permission to concurrently share one physical CPU core.

Because of the scheduling done by the hypervisor across all available CPU resources, it is possible to provision more vCPUs than available physical CPU cores. In addition, the frequent pauses in processing loads for desktop virtual machines make their resource needs less strenuous than that of server VMs that continually process many incoming requests. As a rule of thumb, each physical CPU core provides roughly four vCPUs that can be allocated to DVMs.

This means that a modern server with dual CPU sockets and two quad-core CPUs (with 8 physical cores) would have up to 32 vCPUs cores available at any time. When a desktop virtual machine with a specific vCPU allocation needs processing time, and sufficient physical cores are available, the hypervisor permits that processing to go forward. However, if all vCPUs are allocated, the DVM will be put into a wait state until resources are available. Because of this, and depending on the user workload patterns they support, you should keep the number of vCPUs allocated to concurrently active DVMs close to the actual vCPUs available to ensure predictable performance.

On some virtualization platforms, you can also set MHz reservations (minimums) and limits (maximums) on allocations of vCPU processing resources. By default, vCPUs are assigned the same speed in GHz as the physical CPU core. These reservations and limits can provide finer control over resource allocations, especially for VMs that have fairly consistent loads. However they generally aren't recommended for virtual desktop DVMs, as their loads can have substantial bursts of activity. For DVMs, this sort of vCPU restriction can introduce unnecessary wait states and either degrade the user experience or over-allocate resources.

PEAK CPU DEMANDS

With centralized virtual desktops, server resources are shared across all active user sessions. In most cases, you can safely assume that users are performing different activities at any point in time and that periods of high usage (peaks) and periods of low usage (valleys) will tend to cancel each other out. In some cases, such as in a computer lab or training center, all users may be asked to perform the same task at the same time (such as logging in – causing a “login storm”), thus creating a situation in which individual peaks will coincide. In such cases, you will need to allocate sufficient resources to cover needs during these periods of simultaneous peak usage.

SERVER LOCATIONS

You should review the data center location for the desktop and infrastructure servers. Make sure there is sufficient physical space and cooling for both the infrastructure servers and any other related infrastructure, such as SAN arrays or network infrastructure. Because of the bandwidth requirements for virtual desktops, both for traffic coming in from Panos on an edge network and for connections between servers and your SAN, multiple network drops may be required for each server.

Since both desktop and infrastructure servers can require considerable time to power on and restart their virtual machines, be sure to provide adequate backup or uninterruptable power for all of the servers to allow operations to continue through power interruptions.

Refer to your server's documentation for its power and operating conditions requirements.

SYSTEM VM REQUIREMENTS

The Pano Controller appliance VM has varying server vCPU and RAM requirements depending on the number of Pano clients and DVMs it is servicing and which roles are active in the appliance (ZCC role, VDB role or Full role). Tables 2 and 3 below list the reserved and available vCPU and vRAM resources needed on different platforms per Pano Controller VM.

Table 2: Pano Controller VM Resource Requirements, VMware and Citrix Platforms

Deployment Size (# of Seats)	Reserved vCPU	Available vCPU	# of vCPUs	Reserved vRAM (MB)
0-50	512 MHz	2.0 GHz	1	1024
50-250	1.0 GHz	4.0 GHz	2	2048
250-500	1.5 GHz	6.0 GHz	4	4096

Table 3: Pano Controller VM Resource Requirements, Hyper-V Platform

Deployment Size (# of Seats)	Reserved vCPU	Available vCPU	# of vCPUs	Reserved vRAM (MB)
1-50	1.0 GHz	4.0 GHz	1	2048
51-200	1.0 GHz	4.0 GHz	2	2048

PANO MAESTRO REQUIREMENTS

Pano Maestro has more limited requirements. The minimum requirements for the Pano Maestro appliance VM are 1 vCPU, 1 GB vRAM, and 14 GB virtual hard disk capacity.

Storage Infrastructure

This section discusses how to estimate the storage requirements for your Pano deployments.

STORAGE CAPACITY REQUIREMENTS

To estimate storage capacity needed, determine how much virtual disk space is required for the local operating system and installed applications for each DVM, allowing an extra 3 – 5 GB for Windows temporary and page files. The size of the virtual disk for DVMs should be kept to a minimum – typical figures range from 15 to 20 GB. Larger DVMs will consume more storage resources and make provisioning of new DVMs take longer. If user files are redirected to a file server and users don't install numerous applications after provisioning, the storage requirements for each DVM should be fairly static.

Multiply this size by the number of DVMs, and add 50 GB overhead for Pano Controller VMs and other platform software. In counting DVMs to estimate required storage capacity, be sure to allow for both active DVMs and inactive or template DVMs. If you're not sure how many template or inactive DVMs will be used, a general rule of thumb would be to allow for 1.5 DVMs per user. This should provide a rough estimate of the capacity needed for the deployment. However, the performance of most storage systems and hard drives drops quickly as they exceed 80% capacity utilization. To avoid these performance problems, add another 25% to the calculated capacity to reserve at least that much unused capacity.

Allowing for RAID Overhead

DVM images are typically stored on RAID arrays that also store redundant parity information to allow the array to continue running after the loss of a drive (with RAID 5 designed to survive the loss of one drive and RAID 6 designed to survive the concurrent loss of two drives). Because of this you also need to add sufficient extra storage capacity

overhead for that parity information. The amount of this overhead will vary depending on the number of drives in the array and the RAID level used. A good rule of thumb is to add another 15-25% to your storage capacity requirements to account for this RAID overhead. The storage summaries in the sample architectures are based on an estimate of 20% for storage capacity overhead from RAID.

Table 4: Summary of Storage Capacity Calculations

Factor / Calculation	Value	Description
Number of Seats or Users	# Users	For permanently assigned desktops, number of named users using Panos. For pooled desktops, number of concurrently active users.
Number of DVM Images	DVMs/User	Include active DVMs, template DVMs, inactive DVMs – typically 1.5 DVMs/user.
Average Size of DVMs	DVM Size	Size of virtual disk image for DVMs – typically 15 to 20 GB each assuming folder redirection used for user files.
% reduction from Deduplication	Varies	Potential % reduction in capacity needed due to deduplication or cloning of DVMs.
+ System VM Overhead	+ 50 GB/server	Add 50 GB overhead per server for Pano Controller and platform VMs.
= Net Capacity needed	Net Capacity = (# Users x DVMs/User x DVM Size x (1 – Deduplication %)) + Overhead	
Reserve free space %	25 %	Keep 20% free space in storage to avoid performance drops.
Allow for RAID overhead %	15 - 25 %	Subtract RAID overhead % for parity information storage.
= Raw Capacity needed	Raw Capacity = Net Capacity x (1- Free Space %) x (1 – RAID overhead %)	

DVM Cloning and Deduplication

Some optional platform tools, like VMware View Composer and Citrix Provisioning Services, can reduce the overall storage capacity needed for DVMs via de-duplication techniques like linked clones. This allows a static file (or block within a file), like an operating system component or application DLL, to be stored only once and then be streamed or provisioned to multiple DVMs as they are put into use. Depending on the type and composition of your DVMs, these tools can reduce DVM storage capacity requirements by up to 50% but at the cost of additional complexity and, in some cases, license fees.

Storage System Deduplication

If the storage subsystem hardware implements data deduplication, it will typically come at a cost in terms of latency and total IOPS supplied by the storage subsystem. Although deduplication can deliver storage capacity reductions of as much as 20 to 1, this sort of storage processing – either done in-line (or on-the-fly) as data is written or post-write after the data is saved – is better suited for large static data like backups rather than smaller, more dynamic data like virtual machine images.

In-line data deduplication is typically much slower than the storage processing needed by virtual machines and can introduce latencies that greatly reduce the user experience of DVMs. On the other hand, post-write deduplication can be performed by storage systems after the data is initially saved, providing more time to scan for and remove redundant data. While this reduces the real-time load on the storage system it still creates a lot of activity that can reduce the number of IOPS available for real-time processing of virtual

machine workloads.

IOPS REQUIREMENTS

In addition to storage capacity, a performance metric called input/output operations per second (IOPS) is the other critical factor in sizing storage systems to ensure adequate storage performance for DVMs. Inadequate IOPS will cause DVM processing to pause until storage resources becomes available, degrading the interactivity of the user interface. Available IOPS will depend on the IOPS capability of the drives servicing the DVMs, times the number of drives, less any overhead from RAID-array write-processing of parity information, as described below.

Hard drive vendors will often provide a raw IOPS number for each drive model. However, these figures may be overstated and by themselves they don't take into account IOPS used inside RAID arrays to write parity information. Smaller 2.5" drives typically also have higher IOPS than larger 3.5" drives, due to the higher areal density of the drive media and the shorter distance the drive head needs to move to retrieve or write information. Although disk technology continues to evolve and solid-state disks (SSDs) promise very high levels of IOPS in the future, as a rule of thumb estimate that enterprise-level 10-15K RPM SAS hard drives are the most common storage options for virtual desktop.

The total number of raw IOPS from all drives then needs to be adjusted for the overhead of RAID arrays. This overhead only applies to write operations – for read operations, RAID drives can perform much like the combined performance of the included drives. DVMs tend to use a much higher proportion of read IOPS than write IOPS, especially if certain optimizations, like disabling the Windows page file and indexing, are used to limit unnecessary writes.

Table 5 below summarizes how to calculate the overall available IOPS for a RAID array,

Table 5: Summary of Storage IOPS Calculations

Factor / Calculation	Value	Description
Number of Drives	# Drives	Add up the number of either direct attached or storage array drives.
Determine IOPS per Drive	IOPS/Drive	Calculate or lookup the number of IOPS supplied by the drives you are using – rotation speed (RPM), average disk access latency and seek times. Typical values for 2.5" enterprise-grade SAS drives are 140 raw IOPS per 10K RPM drive and 180 raw IOPS per 15K RPM drive.
= Raw IOPS	Raw IOPS = # Drives x IOPS/Drive	
RAID overhead factor	IOPS/Write	For RAID 5 arrays, this RAID overhead factor can require up to 4 IOPS for each write operation, while RAID 6 arrays can require up to 6 IOPS per write for parity information.
Estimated % Writes	Varies	Estimate the percentage of write operations – average value for Pano DVMs is around 25% provided the correct Windows optimizations are applied.
= Total RAID Overhead in IOPS	Total RAID Overhead = Raw IOPS x % Writes x RAID Overhead/Write	
= Available IOPS	Available IOPS = Raw IOPS – Total RAID Overhead	

Provided it delivers adequate bandwidth and low latencies (well under 25 ms round-trip for reads and writes), the technology (such as direct Fibre Channel links, iSCSI over teamed gigabit Ethernet, and so on) used to interconnect storage systems like SAN arrays and desktop servers won't have a significant impact on the number of IOPS delivered.

Estimating IOPS Requirements per Seat

Storage IOPS available to DVMs on the desktop hypervisor servers will tend to be the gating factor for many virtual desktop deployments. Using servers with more robust CPUs or more RAM can slightly lower the total number of servers required, but the number of available IOPS is critical. To estimate the number of IOPS needed, assume each concurrently active DVM needs about 30-50 IOPS. See Table 1 on page 4 for more information on typical user workloads and suggested allocations.

Keep in mind that IOPS are a transitory resource unlike storage capacity. While you have to allow for storage capacity for both active and inactive DVMs, IOPS only need to be allocated to active DVMs. For example, if your users have an average duty cycle of 80% (i.e. 20% of them are inactive on average) you can drop the inactive portion of your seats from any IOPS allocation totals, lowering the required number of drive spindles.

For DVMs, expect to see a high level of IOPS with a ratio of roughly 90% read and 10% write operations during boot up or powering on of a DVM and then a reversal of that ratio to a much, much lower level of IOPS but with only 20% read and 80% write once Windows is fully loaded. The first ratio only applies to the initial powering on of a DVM – if it is left running and users simply login in and out via the Pano client's dialogs or by using the Pano Button®, only the 2nd ratio will continue, often for many days until the DVM is restarted to apply updates or to correct a temporary problem. This ratio also assumes that network folder redirection is in place and that users aren't working from documents stored locally within their DVMs. Latency between the shared storage and the desktop and infrastructure servers should stay well below 25 ms for both read and write operations.

IOPS and Memory Page Files

One source of unnecessary IOPS consumption is from memory page files in VMs. Best practice is to minimize Windows and hypervisor page file access and leverage physical server RAM instead of using the disk system. To accomplish this, allocate sufficient physical server RAM per DVM to support the needs of the operating system and applications, while keeping paging to a minimum. Then set the Windows page file size inside the DVMs to 50% of the allocated RAM.

IOPS and system Virtual Machines

For Pano Controller, storage IOPS from its VM are minimal as almost everything gets cached into the infrastructure server's memory. For some platform-specific DVM provisioning and management tools, like vCenter, shared storage IOPS can be minimal if the administrator configured them to rely on a separate network-based database server. However, if the supporting database instance is on the same infrastructure server as the platform tool (not recommended for deployments over 200 seats), you may see a significant use of IOPS depending on whether the administrator has enabled non-default reporting settings and or is using the database for additional workload such as update managers, etc.

Shared storage IOPS needed for some other platform tools can be significant depending on whether you are using a SAN that works well them or not. As such it is difficult to give general IOPS estimation. Check with the virtualization platform vendor for their recommended best practices and be prepared to monitor storage traffic and make adjustments to the storage architecture if problems arise.

DVM COLLECTIONS AND STORAGE CHOICES

Pano Controller supports several different forms of DVM collections to simplify provisioning. Collections are used to configure and manage groups of similar DVMs. The types of collections used, along with the level of availability and fault-tolerance required, can determine what storage architecture choices will work best for your Pano deployment.

DVMs can be assigned based on user, device and third-party connection broker collections. For instance, if you want your users to be able to roam freely throughout the workplace and always have access to their DVM, use one of the user-based collections. Use device-based collections to allow a specific Pano client to always connect to a specific DVM.

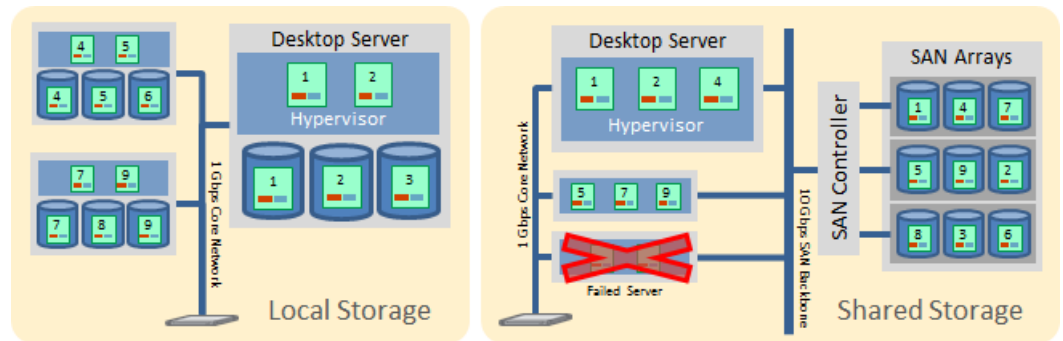
Users can be assigned multiple desktops. For example, a doctor in a clinic might have one permanently-assigned desktop available in their office and a different desktop assigned to an examination room, where each doctor needs to be able to access and update patient records using a certain application.

Local Storage vs. Shared Storage

The decision to use direct attached storage (local storage) vs. a storage area network (shared storage) is largely dependent on the availability requirements for users and the types of DVM collections used. Any DVM stored on local storage is at risk of being unavailable if the server is down due to failure or maintenance. Figure 2 shows how each server using local storage is only able to run the DVMs it stores.

If this risk is acceptable, or if there are alternate, equivalent DVMs on other servers, then local storage can be used to reduce the cost of storage. However, if DVMs are unique for each user and high availability is required, such as in cases where DVMs have been permanently assigned to users to replace their previous desktop PCs, then shared storage should be used, despite its higher cost. Shared storage, as illustrated in Figure 2, allows any desktop server to host any DVM and provides fault tolerance in the event of a server failure, but with the added cost and management complexity of the required storage networks, controllers and arrays.

Figure 2:
Local vs. Shared
Storage of DVMs



Pooled Desktop – Local Storage

For this type of DVM collection, local storage can be used to reduce storage costs because users don't need access to a specific DVM, but rather draw DVMs from a shared pool, based on a standard DVM template. Users are assigned the first available DVM from the pool at user login. An example might be standardized desktops used by task workers in a call center or workstations in a training classroom. After the user logs out, the DVM is returned to the pool. If a desktop server fails, a DVM from the same pool or template on another server can be provided to the user. Any sessions that are active on the failed server when the interruption occurs will be terminated. Because users are not guaranteed to be assigned the same DVM on subsequent logins, you may need to leverage folder redirection and/or third-party solutions to ensure persistent user personalization, user-installed applications and so on across different desktops in the pool.

Permanently Assigned – Local Storage

This type of DVM collection can use local storage to reduce storage cost; however, this entails some availability risk. In this type of collection, users are tied to specific DVMs. An example might be where DVMs are replacing a knowledge worker's previously dedicated PC. If a desktop server becomes unavailable, the sessions for the users with DVMs running on the server will be terminated and the desktops will be unavailable. These users will have to wait for the system administrator to assign alternate desktops or until the server storing their DVMs is back online.

Permanently Assigned – Shared Storage

This type of DVM collection requires shared storage, such as a SAN, if any level of availability after server failures is required. This maintains high availability for users tied to a specific DVM, such as DVMs replacing a knowledge worker's PC, but involves higher cost. With shared storage, the users' DVMs may be migrated to other desktop servers in the cluster either manually or automatically.

Network Infrastructure

Pano virtual desktops are typically deployed on a switched Ethernet network, which tends to avoid or minimize network bottlenecks. The edge network links where Pano clients are located should have 100 Mbps or better switched (rather than shared) Ethernet connections to the core network. Make sure desktop and infrastructure servers are connected via one or more 1 Gbps or better links to the core network. Intervening Ethernet switches or routers are fine, as long as there is a routable connection between Pano clients and DVMs.

To understand the impact of deploying Pano Zero Clients on your network, you need to evaluate the topology of your network, potential sources of network traffic and bottlenecks and the average and peak bandwidth needs generated by your Pano users' workloads.

TRAFFIC SOURCES

To estimate network traffic, evaluate each potential source of traffic. For example, if you consider a desktop server hosting DVMs, it will generally have up to four distinct sources of network traffic:

- Pano Direct Protocol traffic between Pano clients and servers
- Application and file sharing traffic
- Shared storage traffic, such as iSCSI or NFS traffic from the server to shared storage systems like a network attached storage (NAS) or SAN
- Management tool and other infrastructure traffic, such as DVM migration, fault tolerance, image backup, etc.

Deploying Pano virtual desktops will change your current network's traffic characteristics. For example, if you are currently running client/server applications (such as a database) that result in a lot of edge-to-core network traffic between distributed PC clients and application servers, replacing the PCs with Panos will redistribute the traffic. The client-server traffic will now flow between the centralized DVMs (on the desktop server) and application servers over the data center's core network. This will reduce application traffic travelling out to the edge network and may more than offset the Pano Direct Protocol traffic that now goes from the DVMs directly to the Pano Zero Clients.

NETWORK BANDWIDTH

When deploying Pano Zero Clients on networks that include links with constrained bandwidth like WANs, it's important to understand the end-to-end bandwidth utilization for your planned configuration and expected workloads. Bandwidth utilization depends on a number of factors, including the applications, workload patterns, user behavior and the characteristics of the underlying network links.

The amount of bandwidth used by a Pano Zero Client will adapt to network conditions. If network congestion is experienced, the Pano Direct Protocol will slow down to adapt to the lack of bandwidth. This will in turn result in a reduced user experience commensurate with the available bandwidth.

Actual bandwidth utilization of Pano Zero Client sessions is highly dependent on the type of application being used in the DVMs. Certain applications are very rich graphically and update the display very often – others are much simpler, and don't update large portions of the display often or use UI elements with solid colors that use less bandwidth than shaded or patterned areas. Because of these dependencies it is difficult to state an exact amount of bandwidth used by Pano Zero Client sessions. Pano Remote's bandwidth

usage can also vary significantly as its RDP protocol will try to use as much bandwidth as it needs depending on what the user is doing.

Estimating Bandwidth Requirements

To estimate and provision network bandwidth for Pano clients, you need to look at two different bandwidth metrics:

- Average bandwidth – bandwidth consumed over an extended period (such as 8 hours) and averaged. This number will include periods when Pano clients are actively updating the display and periods when they are idle or even disconnected. This figure can be used to estimate the number of sessions that can be provisioned on a given network link. Pano Zero Client users running typical office application workloads will use an average bandwidth of around 500 Kbps per session. Pano Remote sessions might consume as little as 128 Kbps on average.
- Peak bandwidth – maximum bandwidth used for short but large intermittent traffic bursts. Typically these result from activities like moving to a new slide in PowerPoint, minimizing a window, etc. Peak bandwidth bursts can use 5 to 10 Mbps of bandwidth and last between 10 and 20 ms. And while Pano Remote's average usage might be very low, some options, such as printing to local printers, can cause more sustained usage peaks of several Mbps.

For Pano Zero Client users with typical office applications workloads, a good rule of thumb would be to provision around 500 Kbps (average bandwidth) per user, which would support about 160 users on a 100 Mbps Ethernet link (after allowing for 20% loss of bandwidth due to Ethernet overhead). However, you also need to allow headroom for peak bandwidth bursts which, depending on your users' behavior patterns, may require that you drop the number of users provisioned per link by up to 30 – 50%. Using 1 Gbps or better links for connections carrying multiple Panos sessions can assure there will be adequate bandwidth even during peak demand periods. Ideally, once your deployment starts, you should measure both average and peak bandwidth usage for a representative sample of your users to adjust provisioning estimates.

Pano Remote Bandwidth Requirements

Pano Remote™ uses the Remote Desktop Protocol (RDP) rather than the Pano Direct Protocol® used by Pano Zero Clients. Pano Remote's bandwidth usage can vary significantly, as the RDP protocol will try to use as much bandwidth as it needs depending on what the user is doing. A typical average bandwidth usage might be only 128 Kbps per

NETWORK LATENCY

Latency is another key factor in providing an adequate Pano user experience. Latencies above 10 ms round-trip (from a Pano Zero Client to the desktop server and back) impact the performance of the Pano user interface, causing some display changes to lag on the Pano client. Increasing latency will result in a noticeable lag between typing a character and the character showing up on the screen, lag between clicking on a button and seeing the response, and so on. If latency gets too high – above 100 ms – due to server or network congestion, the session between the Pano Zero Client and the DVM will disconnect.

Pano Zero Client Network Latencies

Latencies above 10 milliseconds (ms) round-trip impact the performance of the Pano Zero Client user interface, causing some display changes to lag on the Pano client. Increasing latency will result in a noticeable lag between typing a character and it showing up on the screen, lag between clicking on a button and seeing the response, etc. If latency gets too high, above 100 ms, due to server or network congestion, the Pano session will disconnect.

A typical modern local area network designed for Pano deployments should see about 2 ms round-trip latency. Even on a larger campus with multiple layers of switches and routers, a well-designed network should have no more than 5 ms latency. To achieve this you should use 1 Gbps or better links between switch layers and to servers, cut-through

rather than store and forward switching/routing, and 10 Gbps links for campus backbones.

If you believe latency might be an issue, you can measure it by connecting a computer to the network port planned for a Pano and using the ping command to test the transmission time to the Desktop Server.

Pano Remote Network Latencies

Pano Remote, which uses the Remote Desktop Protocol (RDP) rather than the Pano Direct Protocol used by Pano Zero Clients, can tolerate much higher latencies and lower bandwidth on inter-network links although not with the higher interactivity and multimedia support of a Pano Zero Client.

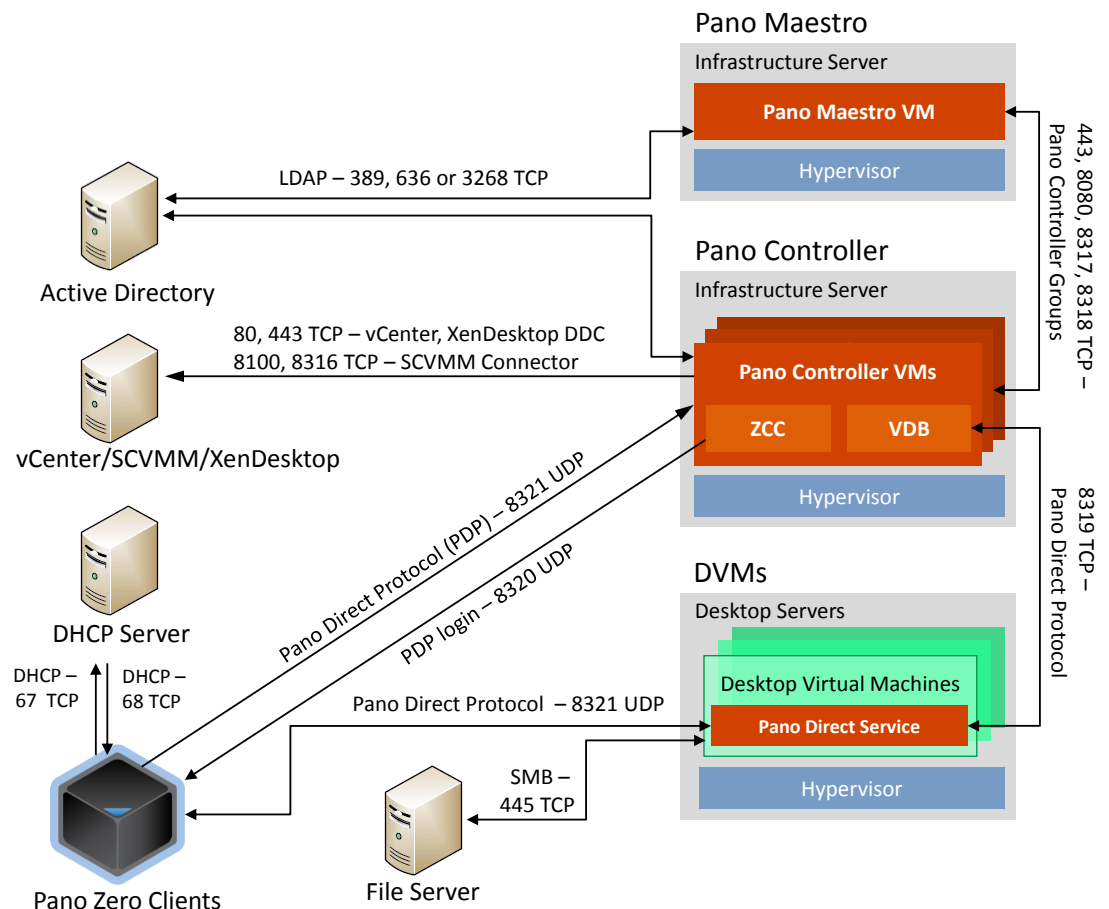
To allow Pano Remote to work over WANs like the public Internet, it is designed to tolerate round-trip latencies as high as 200 ms and still provide a very usable user experience.

NETWORK PORTS

Unlike simpler centralized Pano deployments on a single LAN, remote deployments require some consideration of the IP addressing and network access needs of Pano System and platform components. For example, you need to take into account the subnets, firewalls, and network address translation (NAT) between Pano Clients, Pano Controller and Pano DVMs. Ideally all Pano traffic should stay on the same virtual LAN (VLAN)/broadcast domain in order to use Pano Controller's Enable Local Broadcast method. This option will only work where your Pano Zero Client and your Pano Controller instance are within the same subnet.

Figure 3 shows the main network ports and protocols used by Pano System and platform components.

Figure 3:
Pano System
Ports and
Protocols



If your Pano clients are on a subnet that is external to that of your Pano Controller, you'll need to choose a different Discovery Configuration such as Remote Broadcast Networks, or Probe Address Range. Keep in mind that for these options to work you'll need to make sure that any intervening network switches or routers allow broadcast traffic across segmented subnets. In production environments, Pano Logic also recommends implementing a vendor class in your DHCP services. The benefit is quick and accurate discovery of your Pano clients.

In addition, multicast transmission needs to be enabled within the data center backbone Ethernet switches to enable Pano Controller Group communication. It is not uncommon for multicasting to be disabled by default.

Table 6 lists some of the TCP and UDP ports used by Pano System components and the virtualization platform software.

Table 6: Pano System Ports, Protocols and Services Information

Port	Transport	Protocol	Usage
67, 68	TCP	DHCP	Dynamic IP configuration/lease requests from Pano clients and DVMs, depending on where DHCP server(s) is located.
80	TCP	HTTP	Web server port for incoming http requests. All the requests to 80 are forwarded to 8080
445	TCP	SMB	File server traffic for user files depending on where file server is located
443	TCP	HTTPS	Encrypted (via SSL) connections to the web-based UI used to control Pano Controller and Pano Maestro
3268, 3269	TCP	LDAP	Active Directory, directory services requests from Pano Controller depending on where the Domain Controller is located
8317	TCP	PDP	Pano Maestro port for incoming secure requests from Pano Controller for license verification or FTS License server for inventory validation
8318	UDP	PDP	Connection brokering to the Pano Zero Clients; communication from Pano Controller to Pano Direct Service (PDS) in the DVMs
8319	TCP	PDP	Connection brokering to the Pano Zero Clients; communication from Pano Controller to Pano Direct Service (PDS) in the DVMs
8320	UDP	PDP	Communication from the Pano Zero Clients to Pano Controller to request login screens
8321	UDP	PDP	Bi-directional communication to/from Pano Zero Clients and Pano Controller
8080	TCP	HTTP	Pano Maestro port for incoming non-secure http requests (for example http://<IPAddress>)

PEAK NETWORK DEMANDS

Your network may experience periods of high, peak demand when a number of users log in simultaneously, such as when a branch office opens or when a class starts in a training center that uses Pano clients. These "login storms" take up network bandwidth, as well as CPU cycles and input/output operations per second (IOPS) (see "Storage IOPS Requirements" on page 8 for information about IOPS). This high, peak usage can cause delays in Pano user interface responsiveness. To mitigate this potential problem, you can

configure your system so the operating systems start up before being requested by a user, stagger periods of high activity or over-provision network and infrastructure resources to accommodate these periods of high, peak demand.

Anticipating Network Bottlenecks

Be sure to identify and characterize any potential bottlenecks between Pano clients and desktop or infrastructure servers. A typical bottleneck might be due to a large number of DVMs sharing a limited number of physical Ethernet ports on a desktop server.

Ideally you should provide sufficient physical networking capacity in the servers to handle both sustained average and peak bandwidth needs. Ensure that the anticipated network traffic can be handled by the physical Ethernet network ports installed on the servers. For both performance and redundancy, you should consider installing additional network interface cards (NICs) with multiple network ports. The number of physical network ports should scale with the workload of the server instead of relying on the hypervisor to multiplex numerous virtual network connections from DVMs and system VMs through only one or two physical network ports.

If you are deploying across a multi-facility campus-type environment, there may also be topology constraints, such as at routers or bridges that potentially result in bottlenecks between Pano clients and associated infrastructure or desktop servers.

Prioritize Network Traffic

Networking equipment, like switches and VPN appliances, may allow you to prioritize traffic on LAN and WAN links based on Quality of Service (QoS) tags. The Pano Direct Protocol doesn't directly implement QoS tags, but you may still be able to prioritize Pano traffic based on the ports used or the MAC addresses of the Pano clients and servers. In general, you should only prioritize Voice over IP (VoIP) traffic (and possibly IP video), over Pano traffic, with all other traffic given a lower priority.

In addition, multicast transmission needs to be enabled within the data center switches to enable Pano Controller group communication. It is not uncommon for multicasting to be disabled by default.

OTHER NETWORK INFRASTRUCTURE

You will also need to provide the following network infrastructure:

- **Directory Service:** a directory service (for example, Microsoft Active Directory, version 2003 or later) is required for most production deployments to provide user authentication and enable additional user-based functionality in Pano Controller. For a list of supported directory services, see "Supported Directory Services" in the online help. If you choose not to integrate a directory services server, all DVMs will only be able to use device-based DVM collections – for information on DVM collections, see "DVM Collections and Storage Choices" on page 10.
- **DHCP Server:** Make sure the DHCP network service is available on any network segment that hosts Pano clients - this is typically provided by either Windows Server or a router. The DHCP server should have a sufficient number of IP addresses for both the Pano clients and DVMs. Pano clients and DVMs can be on different network segments and IP subnets as long as they are routable.
- **File Server for User Files:** As a best practice, keep user files on a file server separate from the DVM image. This improves performance by not bogging down the desktop server's storage capacity and IOPS with user file traffic. It also improves availability by reducing the user's dependence on access to a specific DVM image containing their files.

Client Infrastructure

PANO ZERO CLIENT REQUIREMENTS

Finally, check on the locations where the Pano Zero Clients will be deployed. For each Pano Zero Client, the following will be required at the location where it is to be installed:

- An AC power outlet for the Pano Zero Client power supply
- A 100 Mbps or better Ethernet connection
- A DVI- or VGA-compatible monitor, with a recommended resolution of at least 1024x768 - for information on compatible monitor resolutions, see “Supported Monitor Resolutions” in the online help
- A USB keyboard and mouse – for information on compatible USB devices, see “Supported USB Devices” in the online help.

You can find more information on Pano Zero Client requirements and specifications in the *Pano System Data Sheet* on the Pano Logic web site (www.panologic.com).

PANO REMOTE HOST REQUIREMENTS

If using Pano Remote clients, a hosting Windows PC or laptop with a Windows (XP SP3, Vista SP2 or 7) operating system is required at each user's location. Pano Remote clients also need network connections back to either Pano Controller, if deploying internally, or over the Internet to your Pano Gateway server, if used externally.

More Information

For more detailed information on setting up and managing the Pano System, consult the online help available at help.panologic.com and the support knowledgebase in the Pano Logic Customer Center at support.panologic.com.

More information on the Pano System can be found in in the *Pano System Data Sheet* at www.panologic.com/datasheet/panosystem.

Information on deployment planning for remote locations, such as branch offices, distributed facilities, and mobile workers, can be found in the *Remote Deployments Redbook* at www.panologic.com/redbook/remote.

General information on Pano deployment planning, platform choices, scalability and redundancy options, best practices and sample architectures for 25-, 1,000-, and 10,000-seat deployments can be found in the *Deployment Architecture Overview Redbook* at www.panologic.com/redbook/overview.

For more information on the Pano System or to obtain a Starter Kit, please visit www.panologic.com or call us at 650-454-8940.

Pano Logic, Inc.
2000 Seaport Blvd, Suite 200
Redwood City, CA 94063

This document was issued in November 2011 and is specific to the features and capabilities of Pano System 5.0 [RB-IS-110211]

© Copyright 2011 Pano Logic, Inc.

Pano, Pano Logic, Pano Button, Pano Direct Protocol, and Pano Direct Technology are registered trademarks of Pano Logic, Inc.

Pano Device, Pano Gateway, Pano Controller, Pano Remote, Pano Direct Service and Pano System are trademarks of Pano Logic, Inc.