Sizing Guide

August 2009

Sales Tool

This document is intended to provide general guidelines for sizing and configuring the infrastructure supporting your Pano System deployment to help ensure a successful desktop virtualization experience. The right configuration is a balance between CPUs, storage, and memory, more of one may allow for less of another. Our strongest recommendation is to thoroughly evaluate the application environment/demands and test the final configuration.

VMs per Core: 5 – 8

CPU load of a virtual machine highly depends on the workload

- Typical VM load utilizes up to 12-20% of physical CPU, 5-8 VMs / core
- For typical users, 5 VMs per core is recommended
- 8 or more VMs may be deployed per core for lighter workloads

VMs per Disk (15K): 10 – 15

Disk performance is key to VM responsiveness

- Both sufficient IOPS and low latency are required for optimal performance
- High performance disks (15K SAS) can support 10 – 15 VMs per non-parity disk
- Lower performance disks may be used with fewer or less demanding VMs
- Either direct attached storage (DAS), or iSCSI/FC SANs may be used
- Large on-controller caches contribute to performance
- Test SAN accordingly, actual results vary depending on application load and use case

VM physical memory: 410MB – 716MB

Adequate VM memory ensures a positive individual VM experience

- 80% of 512MB or 70% of 1024MB allocated to each virtual machine in physical memory
- Depending on user type, 512K–1GB should be allocated
- Over constrained VM memory can result in paging and over-burden the disk sub-system
- Total memory on host can be over-subscribed – the hypervisor will make intelligent paging decisions for VM's
- Typically, 70-80% of VM memory will be in physical memory, the rest In ESX swap.

15 - 30 Users, ESXi

1 server:

1 quad-core CPU

12 – 16 GB physical memory

3-4 SAS 15K drives in RAID5

Simplest server configuration, vCenter optional.

80 - 120 Users, ESX, Pooled

2 servers:

2 quad-core CPUs

32GB physical memory

6 SAS 15K drives in RAID5

Simple configuration utilizing local storage.

200 - 300 Users, ESX, Pooled

5 servers:

2 quad-core CPUs

32GB physical memory

6 SAS 15K drives in RAID5

Simple configuration utilizing local storage.

80 - 120 Users, ESX, Shared Storage

2 servers:

2 quad-core CPU

32GB physical memory

2 drives in RAID1

1 SAN (iSCSI or FC)

12 SAS 15K drives

RAID5 Configuration

Redundant controllers

Two well configured servers supporting 80-120 users, each utilizing SANs for shared storage. SAN is configured with 12 SAS drives, assuming 1 parity and 1 spare.

200 - 300 Users, ESX, Shared Storage

5 servers:

2 quad-core CPU

32GB physical memory

2 drives in RAID1

2 SAN (iSCSI or FC)

12 SAS 15K drives

RAID5 Configuration

Redundant controllers

Five well configured servers supporting 200-300 users, each utilizing two SANs for shared storage. SAN is configured with 12 SAS drives, assuming 1 parity and 1 spare. One larger 24-drive SAN in 2 RAID groups may also be used.

For more information visit www.panologic.com or call 650-454-8940/877-677-PANO

Pano Logic, Inc. 1350 Willow Rd. Suite 202 Menlo Park, CA 94025

Sizing Guide

August 2009

Understanding Workloads

Light Users

- Task or knowledge workers running only 1 or 2 applications; i.e. web browser or a billing application
- Allocate 512MB of memory per virtual machine and as many as 8 virtual machines per core may be adequate

Typical Users

- Knowledge workers running multiple applications simultaneously, including Office applications
- Allocate up to 1GB of memory per virtual machine and 5 virtual machines per core.

Heavy Users

- Workers using scientific applications, high end graphics or software development
- Usage will typically spike load on CPU, disk, or result in demanding graphics.
- For best performance, allocate as many resources as necessary; create reservations in hypervisor and/or provide dual virtual CPUs per VM, even if more VMs than cores on host.

Do I need a SAN?

The decision to use direct attached storage (local) vs. SAN (shared) storage is largely dependant on the availability requirements and the type of virtual machine collections being used. Any virtual machine stored on local storage is at risk of not being available if the host is unavailable due to a failure or maintenance. If this risk is acceptable, or there are alternate equivalent virtual machines on other hosts, then local storage can be used to reduce costs. If virtual machines are unique and availability is a concern, such as in the case of VMs that have been permanently assigned to users, then shared storage should be used.

Collection Types:

Pooled Desktop – Local Storage

Local storage is used to reduce storage costs.

Users are assigned the 1st available VM out of the pool at login time. After log out, the VM is returned to the pool.

If the server goes down, a VM from the same pool on another server can be served instead.

Any active sessions when the interrupt occurs will be terminated.

Permanently Assigned – Local Storage

Local storage is used to reduce storage costs. Some availability risk.

Users are tied to a specific VM.

If a host becomes unavailable, the corresponding users will loose their sessions.

The users will have to wait for the administrator to assign alternate desktops or wait until the server is back online.

Permanently Assigned – Shared Storage

Leverages a SAN or NAS for shared storage. Maintains high availability for users with permanently assigned desktops.

Users are tied to a specific VM.

If a host becomes unavailable, the corresponding VMs will be migrated to other servers in the cluster either manually or automatically.

VM Best Practices

Allocate plenty of memory and the hypervisor will swap out any unused VM memory

The hypervisor makes intelligent choices because it sees memory utilization across the entire host

Try to keep VMs from paging.

Be aware of disk intensive applications – see notes under Heavy Users, above

Storage Best Practices

Low latency is a priority - keep access latency low by over provisioning your SAN interconnect and storage infrastructure below load maximums.

Keep RAID groups between 6 - 12 drives, do not exceed 15 VMs per LUN

At or near load, SAN access latency will degrade. Disk sub-system performance is the single largest factor in VM success.

Network Best Practices

Switched LAN environments are best for deploying virtual desktops and delivering the *full* desktop experience

Make every effort to reduce packet loss, latency, and jitter. If slow links are causing bottlenecks, use QoS.

Desktops are an interactive service – only prioritize VoIP traffic above desktops. The protocol will adapt to match your network conditions.

For more information visit www.panologic.com or call 650-454-8940/877-677-PANO
Pano Logic, Inc. 1350 Willow Rd. Suite 202 Menlo Park, CA 94025