Virtual Desktop Infrastructure Sizing

May 2012

Solution Brief

This document provides general guidelines for sizing and configuring the supporting server and storage infrastructure to help ensure a successful Pano System deployment. The right configuration is a balance between CPUs, storage, memory and interconnect bandwidth – and more of one may allow for less of another. Our strongest recommendation is to thoroughly evaluate expected workloads for your Desktop Virtual Machines (DVMs) and expect to make some adjustments and fine tuning as your Pano virtual desktop deployment progresses.

DVMs per Core: 3 – 7

CPU load of a DVM is highly dependent on the workload

- Typical DVM load utilizes up to 15-25% of physical CPU, equal to 4-6 DVMs per core
- For medium workloads, allocating 4-5 DVMs per core is recommended
- 6-7 DVMs may be allocated per core for lighter workloads (see next page)

DVMs per Disk (15K): 4 – 6

Disk performance is key to DVM responsiveness

- Both sufficient IOPS (30-50 per DVM) and low latency (average read and write latency under 20 ms) required for optimal performance
- Enterprise-level 15K SAS RAID supports 4-6 DVMs/drive excluding parity drives
- Lower performance disks may be used with fewer or less demanding DVMs
- Either direct attached storage (DAS), or iSCSI/FC SANs may be used
- Large on-controller caches contribute greatly to storage performance
- Test SAN accordingly, actual results vary depending on DVM workloads

DVM physical memory: 768 MB – 2 GB

Adequate DVM memory ensures a positive individual user experience

- Depending on workload and DVM OS, 768 MB – 2GB should be allocated
- Under-allocation of DVM memory can result in Windows paging and over-burden the disk subsystem
- Although not desirable, the total physical memory on server can be over-allocated – the hypervisor will make intelligent paging decisions for DVM's
- Typically, 70-80% of DVM memory will be in physical memory, the rest will be in the hypervisor swap.

15 – 30 Users, ESXi

1 server:

1 quad-core CPU

32 – 48 GB physical memory

5 – 6 SAS 10K drives in RAID

Simplest server configuration, on free ESXi, VMware vCenter Server is optional but still recommended.

50 – 80 Users, Pooled

2 servers:

2 quad or six-core CPUs

48 GB physical memory

7 – 8 SAS 15K drives in RAID

Simple configuration utilizing local storage, pooled desktops.

125 – 200 Users, Pooled

5 servers:

2 quad or six-core CPUs

48 – 72 GB physical memory

7 – 8 SAS 15K drives in RAID

Simple configuration utilizing local storage, pooled desktops.

60 – 90 Users, Shared Storage

2 servers:

2 quad/six-core CPU

48 – 72 GB memory

2 drives in RAID1

1 SAN (iSCSI or FC)

18 SAS 15K drives

RAID Configuration

Redundant controllers

Two well configured servers supporting 60 – 90 users total, each utilizing SANs for shared storage. The SAN is configured with up to 18 SAS drives, including parity and spare drives.

150 – 225 Users, Shared Storage

5 servers:

2 quad/six-core CPU

48 – 72 GB memory

2 drives in RAID1

2 SAN (iSCSI or FC)

18 SAS 15K drives each

RAID Configuration

Redundant controllers

Five well configured servers supporting 150 – 225 users total, each utilizing two SANs for shared storage. The two SANs are configured with up to 18 SAS drives, including parity and spare drives. One larger 36 drive SAN divided into multiple RAID groups may also be used.

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

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

Virtual Desktop Infrastructure Sizing

May 2012

pano LOGIC

Solution Brief

Understanding Virtual Desktop Workloads

Light Workloads

- Task or knowledge workers running only 1 or 2 applications; i.e. a web browser or a billing application.
- Memory allocation of 768 MB (XP) to 1 GB (Windows 7) per DVM .
- As many as 6 – 7 or more active DVMs per core can be allocated.
- Storage system needs to provide roughly 30 IOPS per active user.

Medium Workloads

- Knowledge workers running multiple applications simultaneously, including Microsoft Office applications.
- Memory allocation of 1 GB (XP) to 1.25 GB (Windows 7) per DVM.
- About 4 – 5 active DVMs per core.
- Provide around 40 IOPS per active user.

Heavy Workloads

- Power users using scientific applications, high end graphics or software development.
- Memory allocation of 2 GB or more per DVM (both XP and Windows 7).
- Only 3 – 4 active DVMs per core.
- As much as 50 IOPS per active user.
- For best performance, reserve as many resources as needed – create reservations in hypervisor and/or provide dual virtual CPUs per DVM, even if more DVMs than cores on server.

Do I need a SAN?

The decision to use direct attached storage (local) vs. SAN (shared) storage is largely dependent on the availability requirements and the type of DVM collections being used. Any DVM stored on local storage is at risk of not being available if the server is unavailable due to a failure or maintenance. If this risk is acceptable, or there are alternate equivalent DVMs on the other servers, then local storage can be used to reduce costs. If DVMs are unique for each user and availability is a requirement, such as in the case of DVMs that have been permanently assigned to users, then shared storage should be used despite higher costs.

Collection Types:

Pooled Desktop – Local Storage

Local storage is used to reduce storage costs.

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

If the server goes down, a DVM 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 DVM.

If a server becomes unavailable, the corresponding users will lose 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 but at a higher cost.

Users are tied to a specific DVM.

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

DVM Best Practices

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

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

Try to keep DVMs 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 32 DVMs per LUN

At or near load, SAN access latency will degrade. Disk sub-system performance is the single largest factor in DVM 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 (under 10 ms round-trip), 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. 2000 Seaport Blvd, Suite 200 Redwood City, CA 94063