XenDesktop Planning Guide

Integration with XenServer

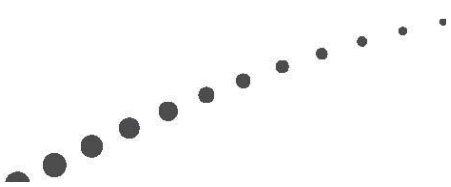

Contents

Introduction 3

Guidelines 3

Pool Master Role 6

Hardware Specification 7

Networking 10

Storage 14

Introduction

This document provides design guidance for Citrix XenDesktop 5 deployments that leverage Citrix XenServer 5.6. It should not be considered as a replacement for other Citrix XenServer or XenDesktop design guidance, but rather an addendum that will assist in design decisions specifically related to using Citrix XenServer as the hypervisor. For further planning guides, please refer to the XenDesktop and XenServer Design Handbooks.

Guidelines

Resource Pool Sizing

A resource pool consists of multiple XenServer hosts bound together to form a single entity for the purpose of management and high-availability. Hosts, within a pool, share common storage and networking components so that virtual machines can be started on, or migrated between (XenMotion), any host within the pool where sufficient resources exist. Although not enforced, Citrix supports up to 16 hosts in a single resource pool. Exceeding this number may result in degraded performance, both for virtual machines and management activities such as maintaining the idle pool count within a desktop group.

The number of resource pools required to support an implementation of XenDesktop depends on various factors, including:

- Number of XenServer Hosts – The scale of the proposed XenDesktop solution may require more than 16 XenServer hosts to support the anticipated number of virtual servers and desktops. As this figure exceeds the recommended number of supported hosts, consider separating out the virtual servers into a dedicated resource pool(s) away from the virtual desktops. This will improve scalability and help prevent resource intensive desktops from negatively affecting the performance of virtual servers.

- Performance – Most businesses will have desktops groups that require guaranteed levels of performance. To address this requirement in a small environment, consider the implementation of dedicated XenServer hosts within an existing pool. For larger environments, it is sometimes necessary to create dedicated pools to meet the service level agreements associated with these desktops.

- High Availability – There may be a requirement for desktop groups to offer varying levels of redundancy. For example, a desktop group used by financial traders could require N+100% redundancy whilst a desktop group accessed by human resources may only require N+10%. In order to accommodate increased levels of redundancy, sufficient capacity must exist within the pool to handle the required number of host failures. In such situations, it may be pertinent to isolate desktop groups into their own resource pools based on the level of redundancy required.

- Application Set – It may be beneficial to dedicate resource pools for large desktop groups as they share a common, predictable resource footprint and application behavior. Alternatively, grouping desktop groups together based on differing resource footprints could help to improve desktop density per host. For example, splitting processor intensive desktops across several pools will help to distribute the impact from processor saturation.

- Physical Network – Some environments may have complex network configurations which require multiple resource pools to be deployed. For example, some desktop groups may need to be isolated onto specific subnets for reasons of security, whilst others may have requirements for network connectivity which can only be satisfied within specific data centers or network locations.

- Virtual Network – Depending on the environment, it may be overly complex to trunk every VLAN to every resource pool. As such, it may be necessary to define resource pools based on the VLANs to which they are connected.

- Processor Architecture – Although XenServer 5.6 simplifies resource pool expansion by allowing disparate processors, there are specific requirements which, if not satisfied, will require the creation of additional resource pools:

    - The server joining the resource pool must share the same CPU vendor (i.e. AMD, Intel) as the servers already in the pool, though the specific type (family, model and stepping numbers) need not be identical.

    - The CPUs of the server joining the pool must support either Intel FlexMigration or AMD Enhanced Migration. Almost all Intel and AMD processors now include this functionality.

    - The features of the older CPUs must be a subset of the features available in the new CPUs.

    - The server joining the pool must be running the same version of XenServer software, with the same hotfixes installed as servers already in the pool.

    - An Enterprise or Platinum XenServer license is required.

- Administrative Boundaries – After virtualization, organizations may need to maintain the separation of administrative duties at a departmental, regional or countrywide basis. Although XenServer role based access, available in the Enterprise and Platinum editions, will be able to satisfy the majority of use cases, there may still be situations where dedicated resource pools are required. For more information on XenServer role based access, please refer to CTX126441 and CTX126442.

- Security Restrictions – Some XenServer resource pools may provide a platform for desktops which are sensitive in nature or which must be isolated in order to meet regulatory compliance. A separate pool may be mandatory under these circumstances to ensure the required segregation and administrative separation.

- Dedicated Hardware – Certain desktop groups may require more expensive hardware than others, including faster network cards, Host Bus Adapters (HBAs), extra memory and additional redundant components. If so, consider separating out these hosts into separate resource pools so that they can be dedicated to the relevant desktop groups.

In many cases, virtual workloads will fall into multiple categories, for example high priority users will require a high level of both performance and high-availability. By capturing all of the requirements, it should be possible to understand the most appropriate resource pool design for the environment and to allocate resources accordingly. In some situations, it may become necessary to configure multiple resource pools within a single desktop group.

Decision	Options	Recommendation
Number of resource pools	1+	Two, one for servers and one or more for virtual desktops.
XenServer hosts per pool	1-16	Exceeding 16 hosts per resource pool is not supported by Citrix.
Resource Pools per Desktop Group	1+	For simplicity, try to maintain one resource pool per desktop group.
Additional resilience per pool	N+	At least, N+1 so that the pool can continue to operate in the event of a single server failure.

Pool Master Role

Each resource pool has one host allocated to the role of pool master that is responsible for handling all pool-wide management activity as well as exposing an administrative interface to the XenDesktop Hosting Management Module, Command Line Interface, xsconsole and XenCenter. All management calls submitted using the XenServer API (xapi) to the resource pool are directed to the pool master where they are proxied to the relevant host within the pool. The XenDesktop Hosting Management Module uses this interface to regulate the idle pool count for each desktop group it manages. As such, it is typically advantageous to host a reduced number of virtual machines on the pool master due to the additional overhead associated with this role. This can be achieved through the use of Workload Balancing, available in the Enterprise and Platinum editions of XenServer. By default, Workload Balancing adjusts the critical metric thresholds applied to the pool master to reduce the load applied, however these metrics can be customized as appropriate:

1. In the Resources pane of XenCenter, select XenCenter > your-resource-pool.
2. In the Properties pane, click the WLB tab.
3. In the WLB tab, click Configure WLB.
4. In the left pane, select Critical Thresholds.
5. In Critical Thresholds page, accept or enter a new value in the Critical Thresholds boxes.

If the pool master host becomes unavailable, the XenDesktop Hosting Management Module will be unable to control the power state of the virtual machines in the affected pool. A typical side effect would be that desktop groups could run out of available powered-on desktops and new users may be unable to connect. To avoid this scenario, each resource pool should be configured for high availability so that a new pool master can be automatically elected – CTX119717. Every pool member is capable of redirecting pool management requests to the pool master, via http redirects, so also ensure that each desktop group is configured to communicate with multiple XenServer hosts within each pool by adding one or more 'HA Servers' for each 'Host'.

Decision	Options	Recommendation
Dedicated Pool Master	Dedicated Reduced Load Full Load	Where sufficient capacity exists, dedicate the Pool Master role.
Resource Pool HA	Enabled Disabled	Enabled in situations where a suitable iSCSI or FC storage repository (356MB or greater) exists for the heartbeat.
Host servers specified per desktop group	1+	A minimum of two host servers should be specified within each desktop group. The first server specified should be the pool master.

Hardware Specification

The hardware selected for the XenServer hosts has a direct impact on the performance, scalability and resilience of the XenDesktop solution. As such, it is critical that the following key areas are considered during the hardware selection process:

- Compatibility – All server hardware should be listed on the current XenServer Hardware Compatibility List (HCL). XenServer 5.6 supports up to 64 logical processors, 256 GB of RAM and 16 NICs per host.

- Scale Up/Out – The decision to scale up or out should be based on the amount of space available, cooling/power capacity and maintenance/hardware or hosting costs. Depending on requirements, it may be advantageous to scale up, rather than out, so that a smaller number of more powerful XenServer hosts are deployed.

- Processor – Citrix XenServer requires 64-bit Intel VT or AMD-V processors to support Windows workloads. All servers in a common pool must be configured with processors capable of executing a common instruction set. One core should be reserved for the use of the control domain.

  As a rough rule of thumb, the following calculation can be used to determine how many virtual desktops can be supported per XenServer host, where processor is the primary bottleneck:

$$Virtual\ Desktops\ per\ Server = Virtual\ Desktops\ per\ Core * (Server\ Cores - 1)$$

User Group	Virtual Desktops Per Core
Light	8-10
Normal	6-8
Power	4-6
Heavy	2-4

Table 1: Virtual Desktops per Core

Note: *Each implementation will have different resource requirements. Scalability testing should be completed to provide more accurate figures.*

Note*: For additional information on virtual machine based resource allocation, please refer to CTX127277.*

Note*: XenServer hosts using the Nehalem or Westmere CPU architecture may become unresponsive, lose network/serial console connectivity and local console access when the CPU attempts to enter a power saving state during idle periods. This functionality can be disabled by turning off C-States in the host's BIOS menu. For additional information on this issue, please refer to CTX127395.*

- Disk – At a minimum, XenServer requires 16GB of locally attached storage, however, 60GB is recommended. The install process creates two 4 GB partitions for the control domain, whilst the remaining space is available for virtual machines.

- Memory – Sufficient RAM should be specified for the host control domain as well as the anticipated number of virtual machines. The memory configured for the control domain should be increased when running more than 50 virtual machines per XenServer host - CTX126531. For information on virtual machine based resource allocation, please refer to CTX127277.

  Dynamic Memory Control (DMC) can be used to increase desktop density per XenServer host by automatically adjusting the memory of running virtual machines between specified minimum and maximum memory values. This allows virtual desktops to borrow/lend additional memory when required. However, desktop users typically utilize more memory as the day progresses. As such, it is important to calculate memory usage per desktop group as a maximum and not an average. DMC should be used with caution as a shortfall in memory will result in slow performance and a negative user perception.

- Component Redundancy – The hardware selected for the XenServer hosts should have sufficient redundancy to meet the requirements of the proposed XenDesktop solution, for example disk, power, cooling, storage connections and network. Depending on requirements, it may be necessary to offer different levels of redundancy per resource pool.

  The use of blade servers introduces specific redundancy requirements not associated with traditional server deployments. For scenarios where a high level of redundancy is required, consider the following recommendations:

  - Network Switches – Within each enclosure, each blade should be logically attached to diverse blade enclosure switches. The blade enclosure should provide redundant network uplinks to the backbone network.

  - Redundant FC-SAN Switches – Within each enclosure, each blade should be logically attached to diverse SAN switches. Each blade enclosure should provide redundant connectivity to the SAN fabric(s).

  - Redundant Chassis Power – Each chassis should be configured with redundant power supplies.

  - Redundant Chassis Cooling - Each chassis should be configured with redundant cooling fans.

  - Redundant Admin Connectivity – Where out of band management interfaces are provided, these too should ideally be redundant.

- Network Interface Cards (NIC) – There are a number of factors which must be considered when determining the number of NICs required in each XenServer host, please refer to the networking section of this document for further information:

    - At a maximum, XenServer 5.6 supports 16 network cards and 8 bonds.

    - The implementation of NIC bonding allows for increased resiliency by allowing up to two network cards to function as a single entity. In the event of one NIC failing, XenServer automatically routes traffic to the second NIC.

    - Routing virtual machine traffic via a Source Level Balancing (SLB) bond will help to improve the throughput and redundancy of the XenDesktop environment.

      Note: *Load balancing associates traffic from each virtual interface to one of two NICs in the bond. Load balancing does not allow a single virtual interface to utilize both NICs in a bond simultaneously.*

    - If sufficient infrastructure exists, performance may be improved by separating out different types of network traffic across multiple physical NICs, for example management, virtual machine, provisioning, backup and storage (iSCSI and NFS) traffic can all be isolated from each other.

  Sharing network cards between virtual desktops can lead to a network bottleneck. The use of fast network adapters/switches (1 Gbps or greater) will help prevent the network from becoming a bottleneck.

  If virtualizing Provisioning Services, the selection of a Single Route I/O (SR-IOV) capable network card will offer enhanced performance and scalability. Please refer to the Network section of this document for additional information.

- Storage Controllers – High performance solutions include Fiber Channel SAN and hardware iSCSI, while lower throughput can be achieved using standard network adapters and configuring software iSCSI or NFS.

Decision	Options	Recommendation
Scaling	Up Out	Scaling up typically offers the optimal return on investment.
Processor	Intel VT AMD-V	Intel VT or AMD-V processors are required to support windows workloads. Disable C-State functionality for Nehalem or Westmere CPUs.
Disk	16GB+	60GB of resilient local storage
Dynamic Memory Control	Enabled Disabled	Disabled, RAM should typically be allocated based on peak utilization.
Network Cards	1 to 16	6 NICs – 2 x NIC bond for management and virtual machines, 2 x NIC bond for provisioning and 2 x NIC bond for storage and backup.

Networking

If unfamiliar with the concepts in this section, please refer to the XenServer 5.6 Administrators Guide – CTX124887 and the Introduction to XenServer Networking document - CTX127885.

When integrating XenServer with XenDesktop it is important to consider the following key networking topics:

- Compatibility – Network adapters should be selected from the XenServer Hardware Compatibility List (HCL).

- NIC Configuration – The host's networking resources are shared by the virtual desktops it supports. If insufficient bandwidth exists, users will experience a degraded level of performance. As such, Citrix recommend the use of fast network cards and switches (1Gbps or greater) to help address this concern.

  If sufficient infrastructure exists, performance may be improved by separating out different types of network traffic across multiple physical NICs, for example management, virtual machine, storage, provisioning and backup traffic can all be isolated from each other. For details on how to setup a multi-homed virtual desktop, please refer to CTX120955.

- Resource Pools – Networking is a pool-level feature in XenServer. Networking changes made on the pool master via XenCenter are automatically synchronized across all hosts within a pool to facilitate XenMotion, High Availability and Workload Balancing. For this reason, it is absolutely critical that the physical cabling and NIC configuration match across all hosts in the pool.

- NIC Bonding – The implementation of NIC bonding allows for increased resiliency by allowing up to two network cards to function as a single entity. In the event of one NIC failing, XenServer automatically routes traffic to the second NIC. Ideally, NICs within each bond should be diversely routed so that a single switch failure does not bring down the bond.

  Source Level Balancing (SLB) allows virtual machine traffic to be load balanced across two physical NICs. Failover support is provided for all other traffic types, including iSCSI and NFS storage traffic. Routing virtual machine traffic via one or more bonds with load balancing enabled will help to improve the throughput and redundancy of your XenDesktop environment. For additional information on bonding, please refer to CTX124421 and CTX127885.

  Note: *Load balancing associates traffic from each virtual interface to one of two NICs in the bond. Load balancing does not allow a single virtual interface to utilize both NICs in a bond simultaneously.*

- Provisioning Services – The significant level of network I/O generated by Provisioning Services can be a hindrance to virtualization. The introduction of Single Route I/O (SR-IOV) capable network cards addresses this concern by allowing virtual machines to directly exploit the hardware without any mediation by the hypervisor resulting in improved performance and scalability. For additional information on SR-IOV, please refer to CTX126624.

  Note: *SR-IOV is not compatible with XenMotion, HA or Workload Balancing as these actions cause any assigned virtual functions to be lost. These limitations are not significant in a Provisioning Services deployment as Provisioning Services implements its own HA and load balancing mechanisms.*

  Provisioning Services is a network intensive application that could negatively impact the performance of other network traffic. For example, Provisioning Services typically streams 166MB of data across the network during the boot process of a single Windows 7 x86 desktop. Depending on the impact Provisioning Services will have on the underlying network, it may be necessary to implement a high-capacity network or to create a separate physical network dedicated to Provisioning Services traffic.

  For details on how to create a separate Provisioning Services network with XenDesktop, please refer to CTX120955. For details on the average data transfers during boot-up by operating system, please refer to CTX125744.

- Addressing – IP addresses need to be assigned to the XenServer management interfaces and individual virtual machines. As such, the design must consider the IP addressing requirements for these components. If DHCP is used to provide the IP configuration for the management interfaces on the XenServer hosts, ensure that reservations are created for the appropriate MAC addresses.

- Resiliency – Many servers today are supplied with NICs offering two or more ports. As such, it is important that any bonds created consist of connections from two separate physical NICs so that a single card failure does not bring down the bond. Figure 1 demonstrates how a host with three dual port NICs can be configured to provide network bonds for management, virtual machine and provisioning traffic.

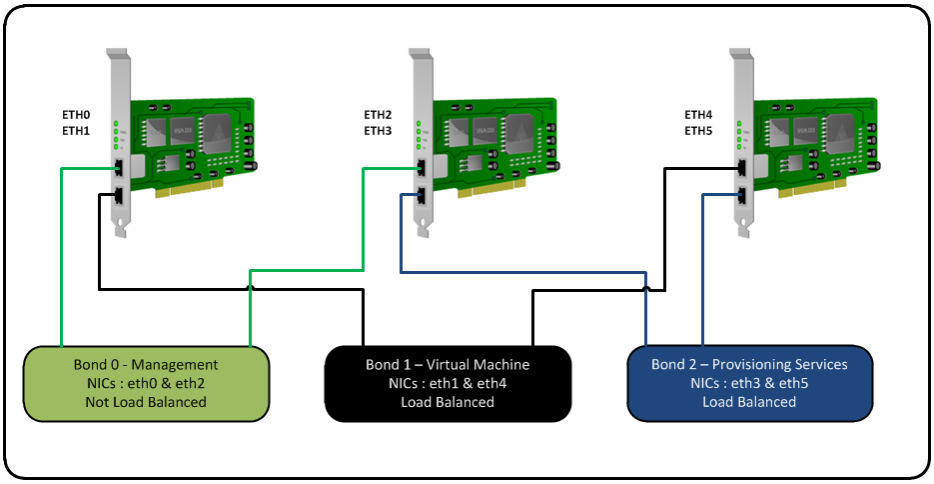

Figure 1: Six port XenServer Bonding Configuration for XenDesktop

Redundancy should also encompass the external network. Bonded NICs should be diversely connected to external switches to help reduce the risk from a single switch failure.

- Quality of Service (QoS) – Some desktop groups may use an unfair share of network resources. In such situations, it may be useful to limit the maximum transfer rate available to these desktops. Quality of Service may be implemented using the XenCenter console or command line interface.

- VLANs – Many network environments take advantage of VLAN technologies to reduce broadcast traffic and enable complex virtual network configurations which would otherwise not be possible. Requirements may exist for separate VLANs per resource pool or even desktop group. Citrix XenServer supports the configuration and use of 802.1Q tagged VLANs. Each XenServer host in a pool will need to have its physical NICs connected to specific VLAN trunk ports to allow for the correct routing of VLAN tagged traffic. For additional information on VLANs, please refer to CTX123489.

  Note: *Configuring the XenServer management interface on a VLAN network is not supported.*

- Security – High security environments may require firewalls to be placed between various Citrix components within the data center. When integrating XenServer with XenDesktop it is essential to ensure that either port 80 or 443 is open (depending on encryption requirements) between the XenDesktop Controllers and the XenServer hosts to facilitate machine state queries and power management operations. For more information on the ports required by XenDesktop and other Citrix technologies, please refer to CTX101810.

Decision	Options	Recommendation
NIC Hardware	HCL	Selecting components from the XenServer Hardware Compatibility List ensures they are certified and supported.
NICs per host	1 to 16	6 x NICs – • 2 x NIC bond for management/virtual machine • 2 x NIC bond for provisioning • 2 x NIC bond for storage and backup
NIC Configuration	Consistent across pool	Same across all hosts within the pool. When using multiple port NICs, bonds should be created using ports from disparate physical NICs for additional redundancy.
Bonding – HA	Enabled Disabled	Enable bonding to improve redundancy provided sufficient NICs, switches and ports are available. Bonds should be diversely routed to help reduce the risk from a single switch failure. Use SLB bonding to improve the throughput of virtual machine traffic, when required.
Quality of Service	Enabled Disabled	Implement when network capacity is limited.
LAN / VLANs	2+	For reasons of security and performance, the XenDesktop infrastructure should be hosted on a separate LAN/VLAN to the virtual desktops.
SR-IOV NICs	Enabled Disabled	If virtualizing Provisioning Services, select SR-IOV capable NICs for improved performance and scalability.

Storage

For an introduction to storage, please refer to CTX118397. Storage has a major impact on the performance, scalability and availability of the XenDesktop implementation. As such, storage requirements must be considered for each key component:

1. XenServer Hosts – At a minimum, XenServer requires 16GB of locally attached storage, however, 60GB is recommended. The install process creates two 4GB partitions for the control domain whilst the remaining space is available for virtual machines. In addition, enabling XenServer HA requires a pool wide shared storage repository (iSCSI or Fibre Channel LUN of 356MB or greater) for the heartbeat and metadata volumes.

2. Virtual Machines – Those virtual Machines which do not utilize Provisioning Services or Machine Creation Services will require storage to host their virtual disk image(s).

3. Provisioning Services – Sufficient storage must exist for the Provisioning Services store to support the estimated number of vDisks required, including backups and future updates.

   Each target provisioned from a shared vDisk must have sufficient storage available to host its write-cache, which can either be hosted on the target itself (RAM/local storage/shared storage) or a Provisioning Services Server (local storage/shared storage). In most situations, consider using either local or shared storage on the target device due to the following concerns:

   a. Device-RAM – This can be an expensive use of memory and targets will fail if memory is exhausted.

   b. Provisioning Services-Storage – Adds additional latency as requests to/from the cache must cross the network.

   Note: *XenMotion, HA and Workload Balancing will not be available when the write-cache is hosted on the target's local storage.*

4. Application Streaming – The Application Hub is a file share or web server used to host streaming profiles. For details on sample profile sizes, please refer to CTX114838.

5. Database Storage – Storage must be allocated for each Citrix XenDesktop database based on current and future requirements:

   - XenDesktop farm
   - XenApp Power and Capacity Management
   - Smart Auditor and Command Center
   - XenApp farm
   - Provisioning Services farm
   - XenApp Configuration Logging
   - EdgeSight for NetScaler
   - EdgeSight for XenApp / Endpoints
   - Workload Balancing

6. User profiles / Data – Central storage is required to share profiles and data between multiple desktops / XenApp Servers.

7. Machine Creation Services – Storage must be allocated for the Master, ID and Difference Disks. XenServer supports Machine Creation Services with the following storage solutions:

    a. Local Disks – Only supported with Virtual Hard Disk on Logical Volume Manager. Virtual Machines created on local disks will not support XenMotion, High Availability or Workload Balancing.

    b. NFS – Recommended storage solution for Machine Creation Services as NFS requires less lock/unlock operations than block level storage, thus increasing storage I/O.

    c. Block Storage – High number of lock/unlock operations when many hypervisors communicate with a single LUN.

The XenDesktop storage design should focus on the following key areas so that the most appropriate storage solution is selected:

- Local/Shared Storage – Virtual deployments typically utilize shared storage in preference to local storage. This can be an expensive strategy, especially for those SMB customers without an existing enterprise storage solution to leverage. As an alternative, it may be possible to achieve the required level of performance and redundancy through the use of local storage. Although using local storage may require additional disks and array controllers to be purchased per server the overall cost is likely to be less than that of an enterprise storage solution.

  A disadvantage of local storage is a reduced level of scalability due to the hardware limit on the number of disks supported, particularly for blade systems. As the number of virtual desktops per host increases, additional disks may be required to accommodate the number of IOPS generated. Another limitation is that shared storage is required to support XenMotion, Workload Balancing and High Availability. Although these features are less critical when supporting virtual desktops, they are still very important for server workloads. As such, it may be pertinent to implement a hybrid environment where virtual servers utilize shared storage and virtual desktops use local storage.

- Tiered Storage – A one size fits all storage solution is unlikely to meet the requirements of most virtual desktop implementations. The use of storage tiers provides an effective mechanism for offering a range of different storage options differentiated by performance, scalability, redundancy and cost. In this way, different virtual workloads with similar storage requirements can be grouped together and a similar cost model applied.

- Performance – The performance requirements for each desktop group must be considered during the design of the storage repositories:

- Storage Architectures – If Network Attached Storage (NAS) is to be used, consider isolating this traffic onto a separate physical network to help prevent congestion.

- Storage Controllers – High performance solutions include Fiber Channel SAN and hardware iSCSI, while lower throughput can be achieved using standard network adapters and configuring software iSCSI or NFS. The performance of Fiber Channel and iSCSI can be improved by implementing multiple storage adapters and configuring them for multipathing, where supported.

  Note: *Bonding only provides load balancing for virtual machine traffic. Failover support is provided for all other traffic types, including iSCSI and NFS storage traffic.*

- Jumbo Frames – Configuring XenServer networks to use jumbo frames can improve performance for storage traffic. Jumbo frames are Ethernet frames containing more than 1500 bytes of payload. Jumbo frames are typically used to achieve higher throughput, reducing the load on system bus memory, and reducing the CPU overhead. The command line interface is used to enable jumbo frames, for example:

  xe pif-param-set uuid=<vif_uuid> other-config:mtu=<value>

  Note: *All NICs and switches between the XenServer hosts and the storage solution must support jumbo frames.*

- SSD/HDD – As the cost of Solid State Drives (SSD) falls, more and more organizations are benefiting from the performance of these drives. However, SSD is still significantly more expensive than the traditional Hard Disk Drive option (HDD).

- Disk I/O Interface – XenServer provides built-in support for local IDE, SATA, SCSI and SAS interfaces whilst remote support is provided for iSCSI, NFS, SAS and Fibre Channel.

- RAID – As virtual desktops are typically write-intensive (80% write / 20% read), RAID-1 or RAID-10 should be considered for the Provisioning Services write-cache in preference to RAID-5 due to the associated write penalties:

RAID Level	Write Penalty
RAID-0	1
RAID-1 & RAID-10	2
RAID-5 (3 data & 1 parity)	4
RAID-5 (4 data & 1 parity | 3 data & 2 parity)	5
RAID-5 (5 data & 1 parity | 4 data & 2 parity)	6

Table 2: IOPS Write Penalty for RAID Configurations

The disk activity for the Provisioning Services vDisk Store will primarily consist of reads, provided that it's not used for private vDisks or server side caching. In this scenario, RAID-5 offers an acceptable solution at a reduced cost to RAID-10.

- IOPS – The number of IOPS generated will vary based on application set, user behavior, time of day and operating system used. Scalability testing should be performed to determine the IOPS required during boot-up, logon, working and log-off phases. Guessing the number of IOPS required is likely to lead to performance issues or an over-priced storage solution.

  The boot process typically generates the highest level of disk I/O activity. As such, virtual desktops should be started in batches prior to the beginning of the business day to help reduce the load on the storage subsystem. In addition, disabling the automatic restart of virtual desktops following logoff will also help to reduce storage load.

  The Windows XP setup and diskpart utilities create misaligned boot partitions resulting in additional disk I/O activity. The diskpart utility included with Windows Server 2003 SP1 and Windows 7 addresses this issue. However, misaligned offsets may still occur depending on the size of the partition. The recommended approach is to always manually create the partition offset with diskpart on all virtual hard disks prior to formatting. For more information on calculating the correct offset, please refer to the following article from Microsoft – KB929491.

The use of dynamic virtual hard disks should be avoided in production environments. Dynamic virtual hard disks include a single byte at the end of the file which causes misalignment with the disk subsystem. Dynamic virtual hard disks also have an expansion algorithm (in 2MB blocks) that generates significant overhead on the storage device when the drive expands. As the dynamic drive expands the allocation table is updated and the drive's header and footer sections are rewritten for each of the file extension operations. For more information on the performance of virtual hard disks, please refer to the following Microsoft White Paper – Virtual Hard Disk Performance.

The number of IOPS generated by the virtual desktops can be further reduced through operating system optimizations. For more information, please refer to the Windows XP – CTX124239 and Windows 7 – CTX127050 optimization guides for XenDesktop.

Current and future requirements for multiple staggered shifts must be considered, as there is likely to be a significant impact on performance due to the increased logon and logoff activity.

The number of virtual disk images assigned to each repository will have a direct impact on performance. A balance must be found between performance, cost and management.

- Redundancy – Storage repositories must be designed to meet the redundancy requirements of the components which they support. This includes RAID levels, storage adapters and the back end storage configuration.

  The best practice for shared storage is to configure two NICs or HBAs in a bonded or multipath setup.

- Provisioning Services – The vDisk store should be hosted on a block-level storage device so that Provisioning Services can benefit from the Windows System Cache. For more information on memory considerations for Provisioning Services, please refer to CTX125126.

- Quality of Service (QoS) – XenServer supports priority based QoS for virtual disks. The use of this functionality should be considered when multiple desktop groups, with different levels of priority, share a common storage repository.

- Thin Provisioning – Administrator can use thin provisioning to present more storage space to the virtual machines than is actually available on the storage repository. Thin provisioning is typically used in development and test environments where virtual machines are frequently rebuilt.

Warning: *If using thin provisioning in production environments, take appropriate measures to ensure that sufficient storage exists. If available disk space is exhausted, virtual machines will fail to write to disk, and in some cases may fail to read from disk, possibly rendering the virtual machine unusable.*

- Data De-Duplication – Storage requirements may be reduced through the use of data de-duplication, whereby duplicate data is replaced with pointers to a single copy of the original item. Enabling de-duplication on the storage repository hosting the write-cache is unlikely to offer significant benefit due to the short lifetime of this data. De-duplication typically takes place outside business hours due to the performance impact associated with this activity.

Decision	Options	Recommendation
Storage Location	Local Shared Mixed	Local storage if existing enterprise storage solution is not available and no requirement for XenMotion/HA/Workload Balancing.
Jumbo Frames	Enabled Disabled	Enabled on dedicated storage NICs.
RAID for Write-Cache and Machine Creation Services	All RAID levels	RAID-1 or RAID-10 as virtual desktops are typically write intensive.
RAID for vDisk Store	All RAID levels	RAID-5, unless private vDisks or server side write-cache is required.
Virtual Hard Disks	Fixed Dynamic	Fixed virtual hard disks should be used in production due to the alignment and expansion issues with dynamic virtual hard disks.
vDisk Store Storage	Block Level Network	Block level storage so that Provisioning Services benefits from the Windows System Cache.
MCS Storage	Local Storage NFS Block Storage Storage Link	NFS storage is recommended as it generates fewer lock/unlock operations than block storage.
Quality of Service	Enabled Disabled	Enabled in situations where multiple desktop groups, with different levels of priority, share a common storage repository.
Thin Provisioning	Enabled Disabled	Disabled in a production environments due to the risk of running out of storage.
De-Duplication	Enabled Disabled	Enabled, although unlikely to offer significant savings for write-cache storage.

Product Versions

Product	Version
XenDesktop	5.0
XenServer	5.6

Revision History

Revision	Change Description	Updated By	Date
0.1	Document Created	Andy Baker	December 3, 2010
0.2	Incorporate feedback from Daniel Feller, Sarah Vallieres and Peter David	Andy Baker	February 7, 2011
1.0	Document Release	Andy Baker	February 22, 2011

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. Its Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2008 was $1.6 billion.