XenApp 6.x Planning Guide: Virtualization Best Practices

Table of Contents

CITRIX®

Overview

Desktop virtualization comprises of many different types of virtual desktops. One option is to use a Hosted Shared Desktop model, which consists of a published desktop running on a Citrix XenApp server.

One of the goals when creating a design for Hosted Shared Desktops is to try and maximize scalability while still providing an adequate user experience. Hosted Shared Desktops provide an advantage over other desktop virtualization techniques by only requiring the use of a single operating system, which significantly reduces user resource requirements and helps improve scalability numbers.

However, in order to get the most users, making correct design decisions as to the resource allocation is important. Creating too many virtual machines or too few might negatively impact scalability.

This planning guide provides best practices for virtualizing Citrix XenApp. Even though these best practices are based on the Hosted Shared Desktop model, they are still relevant in a non-desktop model where users only connect to published applications without the desktop interface.

Scalability

For most organizations, the overall goal is to obtain the greatest user density from a physical hardware configuration without impacting user experience. As with the Hosted VDI model, the number of users that can be hosted on a single server depends upon user load, CPU and memory. As XenApp 6.5 is 64-bit only, it is assumed that the processor subsystem is the primary bottleneck. If 32-bit operating systems are used, it is more likely that memory will be the primary bottleneck.

Based on production deployments, numerous case studies and taking into account the guidelines mentioned within this document, the following table provides some general guidance on how to estimate user density based on the number of physical cores available. The move from dual to quad sockets is not linear and this has been accounted for by a 15% drop in user density.

Number of Sockets	Users per Physical Core		
	Light	Normal	Heavy
Dual	18	12	6
Quad	15	10	5

The following assumptions were made during the creation of these estimates:

- XenApp vCPU: The optimal number of vCPUs assigned to each virtual XenApp server will vary according to the characteristics of the users and applications supported. However, optimal density is typically obtained when 4 vCPUs are assigned to each virtual XenApp server.

- Processor Speed: The speed of the processors has a direct impact on the number of users that can be supported per processor. The estimates provided are based on a processor speed of 2.7 GHz.

- Workloads: Light, normal and heavy workloads are not mixed within a single virtual XenApp server or physical virtualization hosts. It is assumed that all workloads include antivirus and standard monitoring tools.

- Hypervisor Overhead: The overhead from supporting the hypervisor has been accommodated by reducing the estimated number of users per core rather than specifically reserving virtual CPUs.

- XenApp Optimization: The recommended XenApp optimization recommendations have been applied. For more information, please refer to the Citrix Knowledgebase Article CTX131577 – XenApp 6.x Optimization Guide.

The host density estimates in the following table were obtained by multiplying the number of physical cores available by the estimated user density per core values.

Sockets	Cores	Total Physical Cores	VM Count	User Density per Host		
				Light	Normal	Heavy
2	6	12	6	216	144	72
2	8	16	8	288	192	96
2	10	20	10	360	240	120
4	6	24	12	360	240	120
4	8	32	16	480	320	160
4	10	40	20	600	400	200

As a general rule, RAM requirements should be calculated by multiplying the number of light users by 341MB, medium users by 512MB and heavy users by 1024MB. Therefore, each virtual machine hosted on a dual socket host should typically be assigned 12GB of RAM and each virtual machine hosted on a quad socket host should be assigned 10GB of RAM.

Important: *Although these estimates provide a good starting point it is still important that scalability testing be performed to account for areas of variance, including – processor speed, processor architecture, application set, usage patterns and number of idle users.*

Guidelines

In general, when making the decision to virtualize XenApp servers, the main goal is to support the most number of users on a piece of hardware. However, other guidelines can also come into play that might not follow the approach that gives one the best scalability. Business structure, costs and operational items can also play an important role in selecting the most appropriate virtualization strategy for XenApp.

Operations

More virtual servers often mean more management activities. This can be a deciding factor when determining if more or less virtual machines are needed. For example, if ten servers are to host six virtual servers instead of four, that yields 20 more servers to manage, maintain, troubleshoot and support.

However, if single image management solutions, like Citrix Provisioning Services, are used, the operational aspects of managing more virtual server instances becomes irrelevant as all of the virtual servers are based on a single master image.

Application Requirements

Applications influence how a server should be virtualized based on the application characteristics and user activity.

- Resource Requirements: Certain applications have a bigger impact on CPU while other on RAM. The application characteristics plays an important role when determining how many vCPUs to allocate or how much RAM to assign, assuming these requirements don't go beyond the maximum limitations of the operating system.

- User Activity: In many implementations, organizations silo or create separate XenApp servers to host different applications. Silos have a direct impact on resources allocated to a virtual server. If the server contains a set of applications that are rarely utilized, it would make sense to create smaller virtual machines as opposed to a virtual server instance expected to host many more users.

The virtual machine specifications must align with the requirements of the application set and expected usage.

Flexibility

Many organizations virtualizing XenApp servers opt for live migration functionality, where a virtual machine can move from one host to another without impacting active user sessions. This is especially useful when scheduled maintenance is needed on a host or an organization wishes to reduce power consumption by shutting down hosts after hours. Simply shutting down the virtual machine will most likely impact the few remaining XenApp sessions still active, which is why some organizations utilize live migration capabilities. However, the size of a virtual machine can potentially limit flexibility.

When all things are equal, smaller virtual machines provide greater levels of flexibility, especially from a virtual machine placement perspective. For example, it is often easier to move 4 virtual servers (4 vCPU and 8GB of RAM each) than it is to move a single, large virtual machine (16 vCPUs with 32 GB of RAM). In order to successfully migrate a virtual server, the new host must have

enough available resources to support the virtual machine. This is much easier when the virtual machine resource allocations are smaller.

Licensing Costs

Microsoft licensing is very explicit in stating how many licenses are required for physical/virtual servers

- Windows Server 2003/2008 Standard: Each license provides one physical server and one virtual server, although standard is limited to four physical processors.
- Windows Server 2003/2008 Enterprise: Each license provides one physical server and four virtual servers, although Enterprise is limited to eight physical processors.
- Windows Server 2003/2008 Datacenter: A Datacenter license is required for every processor on the physical server, and allows for any number of virtual servers.

The costs associated with licensing the Windows Server operating system can play a role in determining how many virtual machines are created. As an example, would it be better, from a licensing standpoint, to create 4 virtual machines with 6 vCPU or 6 virtual machines with 4 vCPU assuming scalability is similar? The four virtual machine model would result in fewer required Microsoft licenses, which will help keep the costs lower.

In addition, using a feature like XenMotion, Live Migration or vMotion can also impact licensing. The Standard and Enterprise licenses for virtual machines are tied to the physical server. If the virtual machine on one Enterprise server is migrated to another physical server, the new physical server must have enough licenses to support another VM as well as its current virtual machine load.

Note: *Calculating licensing costs and requirements can be done by utilizing the Windows Server Virtualization Calculators (http://www.microsoft.com/windowsserver2008/en/us/hyperv-calculators.aspx).*

Best Practices

The following recommendations are almost always advisable when virtualizing XenApp servers:

Decision	Justification	Hypervisor
Overcommit CPU: No	It is advisable not to allocate more vCPU than there are logical cores within the given hardware. Experience has shown that greater levels of scalability are achieved by not overcommitting CPU.	Hyper-V XenServer vSphere
Utilize Hyper-threading: Yes	Newer processors have the ability to do hyper-threading, where each core is two logical cores. Utilizing hyper-threading in a XenApp environment has been shown to improve user density. However, if it is recommended to NOT enable CPU pinning for VMs where hyper-threading is enabled.	Hyper-V XenServer vSphere

Disable ASLR: No	As many organizations try to protect their XenApp servers from viruses, malware and other OS exploits, it is advisable to keep Address Space Layout Randomization enabled, which is the default setting. The functionality is included with Windows 2008, Windows 2008 R2, Windows Vista and Windows 7.	Hyper-V XenServer vSphere
Enable Transparent Page Sharing: Depends on OS	Enabling or disabling Transparent Page Sharing has not been shown to either help or hurt performance on newer systems (Windows 2008, Windows 2008 R2, Windows Vista and Windows 7). However, older systems (Windows 2003 and Windows XP) have benefited, mostly because the page sizes are smaller (4K), thus making it easier to share pages of memory.	vSphere
Optimize for XenApp: N/A	On systems utilizing pre-Nehalem processors, the XenServer setting "Optimize for XenApp" provided increased scalability. Since the release of the Nehalem processors, much of the functionality has been placed on the hardware so this particular XenServer setting can be ignored.	XenServer
Memory Allocation: Fixed	As users are dynamically load balanced across XenApp servers, memory usage between different virtual machines should be similar, helping negate the need for dynamic memory allocation techniques. Also, if VM migration strategies are used, this could cause memory overcommit resulting in aggressive paging and poor performance across all XenApp virtual machines. It is advisable to set fixed values for memory reservations for XenApp virtual machines.	XenServer Hyper-V vSphere
Host Swapping: No	In most environments, all XenApp servers are actively hosting users at the same time. Swapping out memory from one XenApp host will degrade performance for all virtual machines as the memory keeps getting transferred to/from disk.	vSphere
Dynamic Power Savings: No	To save power, the BIOS settings allow certain processors to slow down when full power is not required. However, hypervisors might not be able to adjust the BIOS settings to allow for full power when required. This results in resources not being fully utilized.	XenServer Hyper-V vSphere

Conclusion

A Hosted Shared desktop offers a feature rich and cost effective solution for light and normal user workloads. However, the performance requirements of heavy users means that they are often better suited to a Hosted VDI desktop so that administrators can granularly assign memory and processor resources on a per-user basis.

Although there are many different factors to consider, the use of dual socket virtualization hosts and virtual XenApp servers with 4 x vCPUs typically offers the best balance between cost, user density and flexibility.

Product Versions

Product	Version
XenApp	6.5, 6

Revision History

Revision	Change Description	Updated By	Date
1.0	Document Created	Daniel Feller – Lead Architect Dimitrios Samorgiannidis – Principal Consultant	May 12, 2011
1.1	Virtualization table updated	Daniel Feller – Lead Architect	June 13, 2011
1.2	Overcommit justification updated (physical cores changes to logical cores)	Daniel Feller – Lead Architect	June 21, 2011
1.3	Added Dynamic Power Savings section	Daniel Feller – Lead Architect Dimitrios Samorgiannidis – Principal Consultant	August 2, 2011
1.4	Updated hyperthreading practice to include note on CPU pinning	Daniel Feller – Lead Architect	January 19, 2012
1.5	Updated server sizing estimates based on updated processor specifications and latest information available	Andy Baker - Architect	August 24, 2012

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.