




# Desktop Virtualization Top 10 Mistakes Made





# Contents

- Contents .....2
- Overview .....3
- 10. Not Calculating Network Impact.....3
- 9. No Profile Strategy .....4
- 8. Lack of Application Virtualization Strategy.....5
- 7. Improper Resource Allocation .....7
- 6. Not Optimizing Antivirus .....8
- 5. Not Managing Boot Storms .....9
- 4. Ignoring Virtual Desktops Optimizations .....9
- 3. Not Enough Cache.....10
- 2. Default Controller Configuration.....11
- 1. Improper Storage Design .....13
- Honorable Mention .....16
- Revision History.....17



## Overview

With the excitement and promise of desktop virtualization, many organizations are trying to quickly implement a VDI-type solution to realize the expected benefits. Unfortunately, this exuberance has resulted in some implementation issues due to improper planning and design. This white paper focuses on the top 10 mistakes made, identified by Citrix Consulting, when implementing a desktop virtualization solution.

### 10. Not Calculating Network Impact

Regardless of the flavor of virtual desktop being implemented (hosted shared, hosted VM-based VDI, local streamed, etc), the network plays a critical role. The user experience degrades as the latency increases and the bandwidth decreases. Proper network planning must be based on the type of work users are performing and the overall network topology. Many organizations estimate 20 Kbps per user, as this was used for XenApp hosted applications. However, virtual desktops require more bandwidth because:

- In default configuration, virtual desktops provide a greater and richer experience than the default XenApp configuration.
- Virtual desktop users are idle less often than hosted application users. For example, if a user is sitting in front of a workstation and is accessing a hosted application on XenApp, the desktop and XenApp applications both appear as active. However, when the user is not working with the XenApp application, they are identified as idle by XenApp even though they might be active on their desktop. Due to this difference, the desktop will have less idle time than XenApp hosted applications.
- Desktops access and consume more multimedia and graphical mediums than most applications. A user's desktop session also includes internet activity, which often includes more graphics, video and sounds. When using a XenApp hosted application, these multimedia activities are only utilized if the application allows.

Estimating network impact is not a trivial matter because the ICA/HDX protocol tunes itself based on the amount of bandwidth available. The less bandwidth available means more compression is applied. Also, any estimate must include percentages for different user activities: typing, graphics, Internet, video (Flash, WMV, etc), and printing. With this information, the following can be used to create an **ESTIMATE**:

Parameter (medium workloads)	XenDesktop Bandwidth	XenDesktop with Branch Repeater
Office	43 kbps	31 kbps
Internet	85 kbps	38 kbps
Printing	553-593 kbps	155-180 kbps
Flash Video	174 kbps	128 kbps
Standard WMV Video	464 kbps	148 kbps
High Definition WMV Video	1812 kbps	206 kbps
Idle	Minimal	Minimal

By calculating the percentage of time a user is expected to be doing certain activities, a rough estimate can be determined for ICA bandwidth requirements. If multiple users are expected to be accessing the same type of content (videos, web pages, documents, etc), integrating the Branch Repeater into the architecture can drastically reduce the amount of bandwidth consumed. However, the amount of benefit is based on the level of repetition between users.

Additional details on the bandwidth estimates can be gathered by referring to the following Citrix white paper: [CTX124457 - Performance Assessment and Bandwidth Analysis for Delivering XenDesktop to Branch Offices](#).

## 9. No Profile Strategy

The user’s profile is one of the major ways the pooled virtual desktop becomes personalized. If users are going to accept a new desktop strategy, they must have the ability to personalize their desktop environment. The personalization must not negatively impact the performance of the environment. When organizations do not properly plan the profile strategy, one or more of the following will likely happen:

- Slow logon/logoff performance
- Inconsistent results
- Lost settings

These challenges will result in a negative perspective of the entire solution. As an example:

An organization had a profile strategy in place. Users started working in the new system. One day, a user had a profile corruption issue which resulted in the deletion of their entire profile. This meant the user had to recreate their entire personalized environment. After the profile was deleted, the user also quickly noticed all of their documents were deleted. Upon closer inspection, the user stored their documents in the “My Documents” folder. When the profile was deleted, the My Documents folder was also deleted.

Another example:



An organization was running a hosted VM-based VDI solution for a few months and decided the profile solution required modifications. Upon the updates, every user lost all of their personalization configurations.

To overcome these potential challenges, a profile strategy must be put into place that includes items like:

- **Folder Redirection:** Have portions of the profile stored on a network drive outside of the roaming profile. This allows the profile to load faster and protects these items from profile deletion.
- **Group Policies:** Utilize group policies to configure the user's virtual desktop profile. These policies should only be used when the user logs onto a virtual desktop.
- **Persistence:** Utilize a profile solution that allows for the extraction and storage of the personalized components of a user's environment outside of the profile.

Additional profile strategy recommendations can be found in the following the following article: [CTX124799 User Profiles for XenApp and XenDesktop](#)

## 8. Lack of Application Virtualization Strategy

Desktop virtualization can be successful without an application virtualization strategy, but only in certain situations, and typically only in smaller, less complex environments. One of the primary goals of desktop virtualization is to simplify the management of the desktop environment. One way of accomplishing this goal is to reduce the number of images. However, the primary factor that often dictates the need for additional desktop images are application sets.

Organizations typically deal with this challenge in one of three ways:

1. Installing every application into a standard desktop image
2. Creating multiple images with different application sets based on different user groups
3. Removing the applications from the desktop image and delivering via application virtualization

All three of these options are valid and work in different scenarios. However, applying one of these options to the wrong environment will result in major challenges as the following examples demonstrate.

One organization installed all of their business applications within a single desktop image. After the image was created and tested, it worked fairly well until certain applications required updates. Those updates sometimes caused issues with the other applications that did not appear until the image was fully deployed. Users also started to express confusion why they could see all of the applications but were unsure what they were supposed to do



with them. The user experience could have been improved by removing the non-standard applications from the desktop image and delivering via application virtualization.

In another example, an organization tried to over design a virtual desktop solution by doing the following:

A small organization consisting of 200 users implemented a virtual desktop solution. Following the complete virtualization guidelines, the organization virtualized all of the applications via hosting and streaming technologies. Although the solution functioned for the users and integrated seamlessly, trying to maintain the different components became a struggle. As the organization only had 4 different application sets, it would have been easier to implement 4 desktop images instead of a complete application virtualization solution.

A proper application virtualization strategy must determine

1. If the number of desktop images is too great to manage effectively. As the number of images increase, the environment becomes more difficult to maintain. By virtualizing the applications, the number of images can be reduced significantly. If, on the other hand, only a few images are required, the time and effort to support an application virtualization solution outweighs the benefit.
2. If traditional (non-virtualized) desktops are still required within the organization. If the applications are virtualized, the traditional desktop management is simplified as these devices can utilize the virtualized applications.
3. If hosted applications are required or if all applications can be streamed to the desktop. By removing the hosted application component, the application virtualization aspect of the environment is simplified as fewer resources are required.

In many implementations, application virtualization is a necessity. Integrating those applications into the virtual desktop must also be done correctly. As a general recommendation, applications should be integrated into a virtual desktop as follows:

	Base	Anomalous	Resource Intensive	Technically Challenging
Description	Core applications required by all users	Unique custom-built Uncertified Terminal Services support	Heavy system resource consumption	Large, complex with many moving parts and dependencies Frequent updates
Examples	Office, Acrobat		CAD/CAM, data processing	Epic, Cerner, SAP
Suggested Approach	Installed in desktop image	Streamed to desktop image	Streamed to desktop image	Installed on XenApp server

## 7. Improper Resource Allocation

Most users only consume a fraction of their total potential desktop computing power, which makes desktop virtualization extremely attractive. By sharing the resources between all users, the overall amount of required resources is reduced. However, there is a fine line between maximizing the number of users a single server can support and providing the user with a good virtual desktop computing experience.

Trying to push the hypervisor, any hypervisor, too hard results in a poor user experience.

Parameter	Hypervisor	Description
<b>CPU Allocation</b>	Citrix XenServer Microsoft Hyper-V VMware ESX	Users should start with a single vCPU and be granted a second if needed due to the following: <ul style="list-style-type: none"> <li>• Most virtual desktops should only be configured with a single vCPU. Most user-based applications are only single-threaded and will not benefit from a multiple CPU configuration. Many user applications do not require significant amounts of processing, which negates the need for more CPU power.</li> <li>• By allocating multiple vCPUs for each virtual desktop, extra resources are required to switch requests across the different vCPUs.</li> </ul>
<b>Command Tuning</b>	Citrix XenServer Microsoft Hyper-V VMware ESX	The XenDesktop controller sends low-level commands to the hypervisor layer to perform tasks on the virtual machines (start, stop, reboot, etc). If too many tasks are sent out simultaneously, the connection to the hypervisor layer can become sporadic. These tasks often have a large impact on the server resources, which impacts the users. It is advisable to throttle the number of commands sent, which was mentioned in 5. Not Managing Boot Storms section.
<b>Transparent Page Sharing</b>	VMware ESX	Transparent Page Sharing allows the ESX hypervisor to share portions of memory that are identical between virtual machines. This has the potential to improve the virtual desktop performance by having a positive impact on memory consumption, although no 3 <sup>rd</sup> party tests have been able to recognize any noticeable benefit. It should also be noted that transparent page sharing does take CPU cycles to compare memory blocks.
<b>Memory Ballooning</b>	VMware ESX	Memory ballooning shifts RAM dynamically from idle virtual machines to active workloads. Memory ballooning artificially induces memory pressure within idle virtual machines, forcing them to use their own paging areas and release memory for active virtual machines. In practical applications, this has shown to be an impediment to positive user experiences. Forcing virtual

		desktops to page to disk requires additional processes. If a large group of idle virtual desktops become active (after lunch, for example), those items must be retrieved from disk, which takes time. And if the server is hosting many active desktops, there is a strong possibility that the server will run out of RAM, which forces the hypervisor to page more memory out to disk. It is advisable to disable this feature.
--	--	--

## 6. Not Optimizing Antivirus

Antivirus solutions are important, even in a virtual desktop environment. If using hosted shared desktops or hosted VM-based VDI desktops, those virtual desktops are located within the data center with other critical systems. If a virus made it into the data center, the environment is at serious risk. However, simply adding an antivirus solution to the virtual desktop can have a major impact on the virtualization infrastructure, and even cause users to experience poor virtual desktop performance.

If the virtual desktops are streamed with Provisioning services, and those desktops start a full system scan at roughly the same time, those virtual desktops will eventually request the entire vDisk image. This not only overwhelms the network and Provisioning services, but also impacts the storage infrastructure as the write cache is utilized. Overcoming these issues is a fairly easy matter and is based on the following recommendations:

1. The vDisk image must be free from viruses. It is recommended to do a full system scan in private image (read/write) mode. This guarantees the image is clean.
2. When the vDisk image is in standard mode (read-only), the antivirus should be configured as follows:
  - a. Only scan create/modify activities of files
  - b. Scan on write events only
  - c. Scan local drives only
  - d. Exclusions
    - i. Pagefile
    - ii. Print Spooler directory
    - iii. Write cache file
    - iv. EdgeSight database
    - v. ICA client's bitmap cache directory





- e. Remove the antivirus configurations from the HKEY\_LOCAL\_MACHINE\Software\Microsoft\Windows\Current Version\Run registry key
3. Reconfigure antivirus so that the virus definitions file is stored on a persistent disk so antivirus doesn't have to download the entire definition file on each startup.

## 5. Not Managing Boot Storms

Most organizations have users arriving and logging into their desktops at roughly the same time. It is recommended to use the XenDesktop workstation group idle settings to prepare the environment for the user logon storm by booting desktops so many minutes before the users arrive. This makes the desktops immediately available for users and allows the system to recover from the massive boot storm. However, when the workstation group's defined boot time is reached, the controller might try to start thousands of virtual desktops simultaneously.

A virtual desktop startup has the single largest impact to any virtual desktop implementation. The XenDesktop controller must tell the hypervisor to start a new desktop and the hypervisor must allocate resources. Sending too many requests to the hypervisor can overwhelm the hypervisor's management layer (VMware ESX, Microsoft Hyper-V and Citrix XenServer). This should be mitigated by configuring the maximum number of simultaneous startups the controller can request. This is done by doing the following:

- On the XenDesktop Master Controller, edit the file: C:\Program Files\Citrix\VmManagement\CdsPoolMgr.exe.config
- Locate the MaximumTransitionRate entry and use a value of 20 (change based on actual environment parameters). The value entered forces the XenDesktop controller to limit the number of requests that are sent to the hypervisor's management layer.

This setting should be made to all XenDesktop controllers in the event the master fails and a backup takes over.

## 4. Ignoring Virtual Desktops Optimizations

Organizations often spend time to create a customized standard operating environment for their desktop operating systems. This often involves specific location settings, default application settings, and desktop descriptions. However, when delivering an operating system into a virtual desktop, many organizations do not go far enough to optimize the desktop for the virtualized environment. Whether the desktop is a hosted VM-based VDI desktop, a local streamed desktop or a hosted shared desktop, certain optimizations allow the hardware to focus on user-related tasks as opposed to extraneous system-related tasks. The following are examples of virtual desktop optimizations:

- **Disable Last Access Timestamp:** Each time a file is accessed within an operating system, a time stamp is updated to identify when that file was last accessed. Booting up an operating system accesses hundreds and thousands of files, all of which must be updated. Each action requires disk and CPU time that would be better used for user-related tasks. Also, if Provisioning services is used to deliver the desktop image, those changes are removed when the desktop is rebooted.
- **Disable Screen Saver:** Utilizing a graphical screen saver consumes precious memory and CPU cycles when the user is not even using the desktop. Those processes should be freed and used by other users. If screen savers are required for security purposes, then simply blanking the screen should be invoked as this does not impact the memory and CPU consumption.
- **Disable Unneeded Features:** Windows 7 contains many valuable components like Media Center, Windows DVD Maker, Tablet PC Components, and Games. These applications are memory, CPU and graphics intensive and are often not required in most organizations. If these components are made available to users, they will be used. It is advisable to remove unneeded services before deploying the first images.

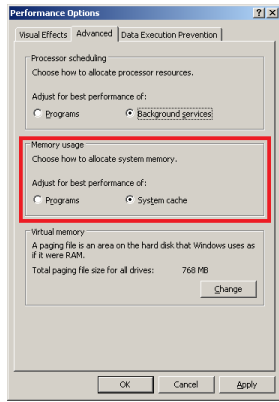
These are only a few recommendations, but it is obvious that optimizations have a major impact on the virtual desktop environment.

*Note:* The virtual desktop optimization best practices for Windows XP is located in [CTX124239](#)

### 3. Not Enough Cache

System cache is a powerful feature allowing a server to service requests extremely fast because instead of accessing disks, blocks of data are retrieved from RAM. Provisioning services relies on fast access to the blocks within the disk image (vDisk) to stream to the target devices. The faster the requests are serviced, the faster the target will receive. Allocating the largest possible size for the system cache should allow Provisioning services to store more of the vDisk into RAM as opposed to disk. However, many environments are not configured optimally. Simply adding RAM to a Provisioning services server is not enough; the system must be configured appropriately.

Parameter	Description
<b>Operating System</b>	<p>The operating system plays a large role in how large the system cache can become.</p> <ul style="list-style-type: none"> <li>• Windows Server 2003/2008 x32: 960 MB</li> <li>• Windows Server 2003/2008/2008 R2 x64: 1 TB</li> </ul> <p>Because the 64 bit operating system can have a larger system cache, a larger portion of the vDisk can be stored in RAM, which is recommended.</p> <p>Windows 2008 is recommended over 2003 because of the improvements in the memory manager subsystem, which has shown some improvements.</p>
<b>RAM</b>	<p>16-32GB of RAM</p> <p>The more RAM allocated for the server, the larger the system cache can become.</p> <p>The larger the cache means vDisks reads will be faster.</p>

<b>vDisk Storage</b>	<p>The vDisk can be stored on just about any type of storage (iSCSI, Fiber, local, NFS, CIFS, etc). However, there are a few instances where the storage selected will have an impact on how the Provisioning services server's operating system caches the vDisk blocks.</p> <ol style="list-style-type: none"> <li>1. Network Drive: If the Provisioning services server sees the vDisk drive as a network drive via a UNC path, the server will not cache the file.</li> <li>2. CIFS Share: If the storage infrastructure is a network CIFS share, Provisioning services will not cache the vDisk in memory.</li> </ol>
<b>Optimizations</b>	<p>In Windows Server 2003, large system cache must be enabled by configuring the server's performance options, which is shown in the figure to the right.</p> <p>In Windows Server 2008, this setting is not required due to the enhancements in the memory allocation system. Windows 2008 utilizes a dynamic kernel memory assignment that reallocates portions of memory on-the-fly, while previous versions had these values hard set during startup. As Windows 2008 requires more system cache, the operating system will dynamically allocate.</p> 

## 2. Default Controller Configuration

At a high level, the XenDesktop controller is responsible for:

- Authenticating users against Active Directory
- Enumerating available resources
- Creating registrations for newly started virtual desktops
- Maintaining an active heartbeat with online virtual desktops

When XenDesktop is installed, many architects fail to optimize the controller and simply install with default configurations (farm master, XML broker, Web Interface). This configuration is perfectly reasonable. XenDesktop can function with a single controller. However, during boot and logon storms, where hundreds or thousands of users connect to the environment in a short amount of time, the controller can become a bottleneck.

The controller bottleneck can result in long logon times or even denied service. By separating controller functionality across multiple servers, the overall XenDesktop farm can support more virtual desktops and respond faster.

First, all XenDesktop implementations should have redundant controllers to provide fault tolerance. For environments that are expected to host more than 1,000 hosted VM-based VDI desktops, it is often recommended to separate the controller functionality across a minimum of five servers all of which can be virtualized. The role designation across the five servers would be as follows:

Parameter	Description	Values
<b>Master Controller</b>	<p>The master controller is in charge of the overall XenDesktop farm operations. It also maintains the correct number of idle desktops within the workstation groups by communicating with the virtualization layer (XenServer, Hyper-V or ESX).</p> <p>If the Master fails, a secondary, dedicate server can be used. If one is not defined, one of the XML brokers will be used.</p>	<p>Specify the XML controllers as backup controllers by modifying the registry:</p> <p>Key:</p> <ul style="list-style-type: none"> <li>• HKLM\Software\Citrix\IMA\RUNTIME\UseRegistrySetting</li> <li>• Type: Dword</li> <li>• Value: 1</li> </ul> <p>Key:</p> <ul style="list-style-type: none"> <li>• HKLM\Software\Citrix\IMA\RUNTIME\MasterRanking</li> <li>• Type: Dword</li> <li>• Value: 1</li> </ul> <p>To disallow the Master controller from accepting virtual desktop registrations, modify the following registry key:</p> <p>Key:</p> <ul style="list-style-type: none"> <li>• HKLM\Software\Citrix\DesktopServer\MaxWorkers</li> <li>• Type: Dword</li> <li>• Value: 0</li> </ul>
<b>XML Controller (x2)</b>	<p>The XML controllers are responsible for registering virtual desktops, authentication and enumerating users.</p>	<p>Configure Web Interface to use the XML Controllers as the farm servers. This will only use these two servers for authentication and enumeration activities.</p> <p>Specify the XML controllers as backup controllers by modifying the registry:</p> <p>Key:</p> <ul style="list-style-type: none"> <li>• HKLM\Software\Citrix\IMA\RUNTIME\MasterRanking</li> <li>• Type: Dword</li> <li>• Value: 2</li> </ul> <p>Key:</p> <ul style="list-style-type: none"> <li>• HKLM\Software\Citrix\IMA\RUNTIME\UseRegistrySetting</li> <li>• Type: Dword</li> <li>• Value: 1</li> </ul>
<b>Web Interface (x2)</b>	<p>The Web Interface servers provide the end user interface for authentication and desktop access.</p>	<p>Configure Web Interface on redundant servers to offload the processes from the Master and XML Controllers.</p>

By separating out the controller roles into three roles across five or more virtual servers, a XenDesktop farm is able to host 5,000+ hosted VM-based VDI desktops. Of course the maximum number for any implementation is based on numerous factors like logon rates, acceptable logon time, and hardware.

# 1. Improper Storage Design

When it comes to storage, most people think only about size... How much space do I need to allocate per user? With desktop virtualization, storage goes beyond simple size calculations. Virtual desktops rely on the storage infrastructure to load different parts of the operating system and user environment. Each one of the requests impacts the storage infrastructure. Without a properly designed storage subsystem, a user's virtual desktop will slow down to the point of becoming unusable because the storage becomes a bottleneck.

In order to properly design the storage infrastructure, the architect must first be able to calculate the expected Input/Output Operations per Second (IOPS) requirements. The IOPS calculations must take the following into account:

Parameter	Description	Values
<b>Disk Speed</b>	The speed that the disk spins has a direct impact on how fast the disk can read the correct sector.	15,000 RPM: 150 Random IOPS 10,000 RPM: 110 Random IOPS 5,400 RPM: 50 Random IOPS
<b>Read/Write</b>	IOPS are broken down into reads and writes. Certain processes are read intensive while others require more writes. The ratio between reads and writes impacts the overall IOPS.	It has been shown that most desktop implementations result in the follow read/write ratios: <ul style="list-style-type: none"> <li>• Reads: 20%</li> <li>• Writes: 80%</li> </ul>
<b>RAID Level</b>	The RAID configuration impacts how many actual write IOPS are available due to the different types of redundancy in place. The write penalty reduces the overall IOPS for each disk spindle.	RAID 0: No RAID Penalty RAID 1: Penalty of 2 RAID 10: Penalty of 2 RAID 5 (4 disks): Penalty of 4 RAID 5 (5 disks): Penalty of 5
<b>Desktop Lifecycle</b>	Each desktop goes through six phases, with each phase incurring different hits on the storage subsystem.	Boot: 26 IOPS Logon: 14 IOPS Work: 8 IOPS Idle: 4 IOPS Logoff: 12 IOPS Offline: 0 IOPS

Taking the six different parameters into account, an architect can calculate the IOPS requirements on a server-by-server basis and for the entire desktop virtualization architecture. The formula is as follows:

$$Total\ Raw\ IOPS = Disk\ Speed\ IOPS \times \#\ Of\ Disks$$

$$Functional\ IOPS = \left( \frac{(Total\ Raw\ IOPS \times Write\ \%)}{RAID\ Penalty} \right) + (Total\ Raw\ IOPS \times Read\ \%)$$

For example, if there are eight 72GB 15K SCSI3 drives in a RAID 10 configuration, the storage would have 720 functional IOPS.



$$\text{Total Raw IOPS} = 150 \times 8 = 1200$$

$$\text{Functional IOPS} = \left( \frac{(1200 \times .8)}{2} \right) + (1200 \times .2) = 720$$

With the functional IOPS calculated for the disk subsystem, the number of desktops that can be supported is based on the phase of the desktop lifecycle. Identifying how many desktops can be simultaneously logged in with this configuration is as follows:

$$\text{Support Desktops} = \frac{\text{Functional IOPS}}{\text{Desktop Lifecycle IOPS}}$$

$$\text{Supported Desktops} = \frac{720}{14} \approx 51$$

These calculations help identify what is possible when all desktops are doing the same activity. However, this is not likely to be the case. In fact, on each hypervisor, a portion of the desktops will be in one of the five phases. To determine the storage requirements, one must calculate the peak IOPS load during a logon storm or boot storm.

In addition to IOPS requirements,

Parameter	Recommendation	Values
Write Cache RAID	RAID 1 or RAID 10	The write cache is write intensive, which eliminates RAID 5 as an option.
Write Cache IOPS	Plan for 14 IOPS	Using 14 IOPS as the scalability calculation or IOPS allows for intense logon storms, which cannot be controlled. Although boot storms have a larger IOPS impact, the storm can be managed with the XenDesktop configuration.
vDisk RAID	RAID 5	The vDisk will only experience reads. By optimizing the RAID configuration for reads (RAID 5), the speed of vDisk block delivery is increased.
Disk Type	Fixed	Fixed disks overcome the performance challenges of dynamic disks, such as: <ul style="list-style-type: none"><li>• Fragmentation as disks continue to expand</li><li>• Simultaneous expansion of hundreds of disks during virtual desktop startup</li><li>• Misalignment of the storage, which results in each I/O operation within the disk requiring two I/O operations on the storage infrastructure as the blocks within the dynamic disk cross block boundaries on the storage infrastructure.</li></ul>





## Honorable Mention

In addition to the Top 10 mistakes made, the following list includes a few other items that, while important, did not make it into the top.

1. **NIC Teaming:** Provisioning services streams the desktop image to the virtual desktop. Provisioning services NICs should be teamed for throughput/aggregation and not just for failover/redundancy. More than two NICs can be teamed for aggregation based on network throughput requirements.
2. **NIC Optimization:** although Provisioning services can run with the default NIC configurations, the environment can run faster with a few optimizations: Disable Large Send Offload
3. **Common Image:** Reducing the number of images helps simplify management and updates as fewer image updates are required. However, using a single image across multiple physical end point platforms can become difficult to maintain. Specific hardware drivers can potentially conflict and installing multiple device drivers results in image bloat. It is often better to create different images for different hardware (not applicable if the end point is virtualized).
4. **VDI for Wrong Reason:** Organizations should do virtual desktops because there is a business reason to provide users with a Windows XP/Windows 7 desktop interface. Without a business reason, the virtual desktop solution will be seen as extravagant and costing too much money for no value.





## Revision History

Revision	Change Description	Updated By	Date
1.0	Document created	Daniel Feller – Lead Architect Doug Demskis – Sr. Architect Tarkan Koçoglu– Sr. Architect Nicholas Rintalan – Architect	July 1, 2010
1.1	Updated registry values for Farm servers	Daniel Feller – Lead Architect	August 16, 2010

### 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.

©2010 Citrix Systems, Inc. All rights reserved. Citrix®, Access Gateway™, Branch Repeater™, Citrix Repeater™, HDX™, XenServer™, XenApp™, XenDesktop™ and Citrix Delivery Center™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.