



CITRIX[®]

XenDesktop Handbook

An architect's guide to desktop virtualization



Overview Quick Access Links

Introduction.....	3
Methodology.....	4
Project Plan.....	5

Project Accelerator

This is the struggle to achieve success within your user community



These are the deciding factors of desktop virtualization as well as other technologies. [Click Here](#) to visit [Citrix Project Accelerator](#) an interactive online tool creating customized sizing and design recommendations based on the methodology, best practices and expert advice identified within this handbook.

Introduction

In traditional business environments, workers suffer from productivity loss in many ways, including downtime during PC refreshes, patches and updates, or simply when they are away from the office. Application and desktop virtualization centralizes apps and desktops in the datacenter, rather than on local devices. This allows IT to deliver apps and desktops to users on demand, to any device, anywhere.

Take the following response from a desktop virtualization user:

Experience from the field

Take the following response from a desktop virtualization user: As a remote employee for [company], I struggled every time I needed to access the company's intranet, which forced me to VPN into the network. I also kept data on my local device because trying to access it over my broadband connection was too slow. Some coworkers did the same and lost data due to a virus, thankfully I was luckier.

Depending on my mood (and the weather), changing devices and locations was a challenge as I had to have my applications and data copied to many different endpoints. I know this was unsecured, but I didn't care because I was more concerned with flexibility.

Since moving to a virtual desktop, I'm able to use any device. I'm able to work from any location. And best of all, I don't have to worry about copying my data and applications onto all of my personal devices. I paid for these devices; I don't want work to clutter up my personal space.

Unfortunately, organizations sometimes struggle to achieve this level of success. Why does one organization succeed while another organization struggles?

If we compare the factors between success and failure between

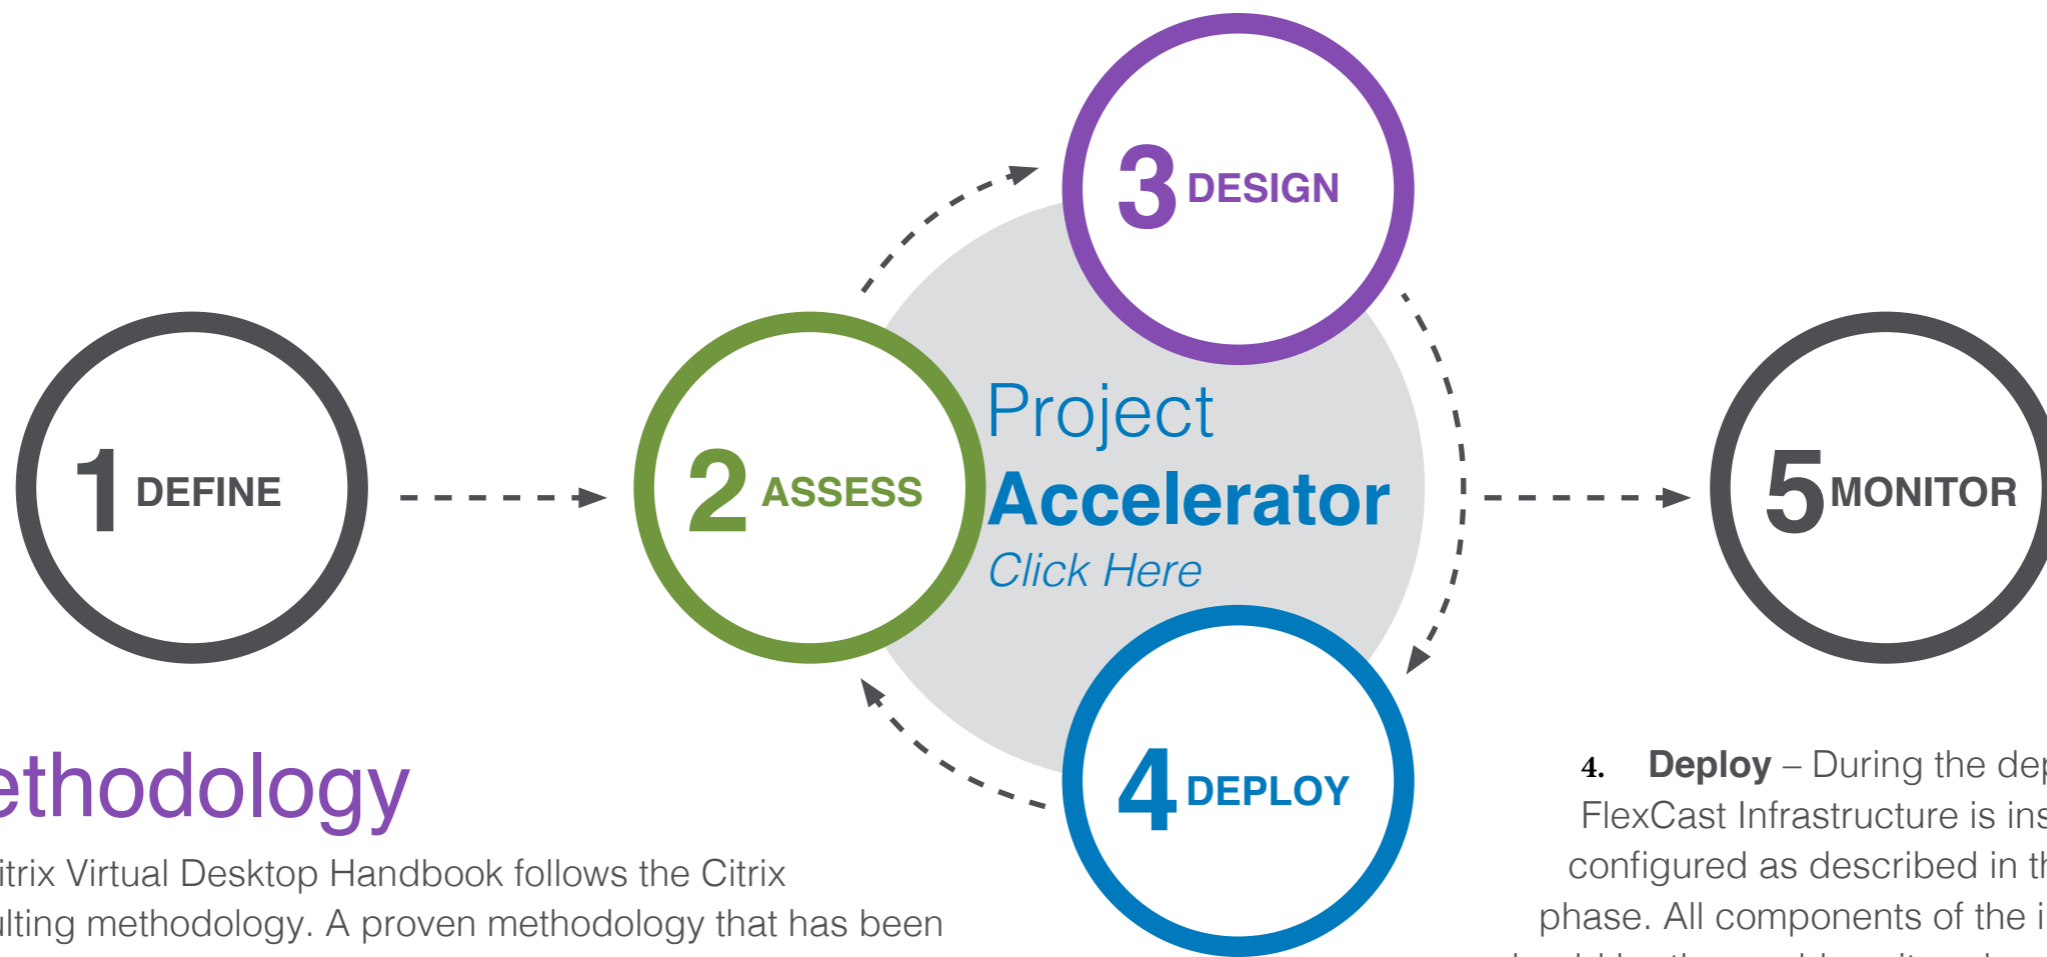
desktop virtualization and other technology related projects, we see that there is little difference:

- **Lack of justification** – Without a solid business reason, desktop virtualization is simply a new way to deliver a desktop. A business justification gives the project team a goal to strive towards.
- **Lack of a methodology** – Many people who try and struggle to deploy a desktop virtualization solution do so because they jump right in without understanding or implementing the appropriate prerequisites. A structured methodology provides the path for the project.
- **Lack of experience** – For many who embark on a desktop virtualization project, there is a lack of experience, which creates a lack of confidence in the design. Architects begin to second-guess themselves and the project stalls.

Our hope is that this handbook can alleviate the anxiety associated with desktop virtualization by showing how challenges can be resolved in a manner that is technically sound, but also feasible and effective for organizations facing deadlines and other organizational challenges.

Citrix Consulting has successfully employed the methodology, experience and best practices shared within this handbook across thousands of desktop virtualization projects.

The [Citrix Virtual Desktop 5.x](#) and [Virtual Desktop 7.x](#) handbooks are not the only resource to guide you through the desktop virtualization journey. Citrix also provides [Project Accelerator](#); an interactive online tool creating customized sizing and design recommendations based on the methodology, best practices and expert advice identified within this handbook.



Methodology

The Citrix Virtual Desktop Handbook follows the Citrix Consulting methodology. A proven methodology that has been successfully employed across thousands of desktop virtualization projects. Each phase includes guidance on the important questions to ask, what tools to use and tips to help you succeed. The Citrix Consulting methodology consists of five phases:

1. **Define** – Builds the business case for desktop virtualization by creating a high-level project roadmap, prioritizing activities and estimating storage and hardware requirements.
2. **Assess** – Key business drivers are rated so that work effort can be prioritized accordingly. In addition, the current environment is reviewed for potential problems and to identify use cases for the project. This information will be used to set the direction of the Citrix deployment, upgrade, or expansion.
3. **Design** – Define architecture required to satisfy key business drivers and success criteria identified during the assess phase. Topics such as environment scalability, redundancy and high availability are addressed.

4. **Deploy** – During the deploy phase, FlexCast Infrastructure is installed and configured as described in the design phase. All components of the infrastructure should be thoroughly unit and regression tested before users are provided with access to the environment.

5. **Monitor** – Define architectural and operational processes required to maintain the production environment.

The Citrix Consulting methodology follows an iterative Assess > Design > Deploy process for each major initiative of your project. In doing so, your organization is left with tangible improvements to the environment at the end of each engagement. For example, high priority user groups can progress through the assess, design and deploy phases earlier than other user groups

Note: The Virtual Desktop Handbook provides content on the Assess and Design phases of the Citrix Consulting methodology. Additional phases will be released soon

Assess Quick Access Links

Assess Overview	7
Step 1: Define the Organization	7
Step 2: Assess the Environment	9
Data Capture Strategy	9
Capabilities Assessment	11
Step 3: Define the User Groups	12
User Segmentation	12
Assign FlexCast Models	14
Step 4: Define the Applications	17
Application Rationalization	17
Link Apps to Users	18
Step 5: Project Management	19
Roadmap	19
Build the Right Team	20
Business and Technical Roles	21

Assess Overview

Creating a desktop virtualization solution begins with a proper assessment. Architects that fail to properly assess the current environment find that they require the assess information later on, forcing them to backtrack, which can potentially stall and put the project at risk.

By gathering all of the information from the outset, the architect will gain an appreciation for the current environment and be able to work from the beginning on properly aligning business and user requirements with the overall solution.

The assess phase is a five-step, simple to follow process:

- 1 - Define Organization**
- 2 - Assess Environment**
- 3 - Define User Groups**
- 4 - Define Applications**
- 5 - Plan Project**

Step 1: Define the Organization

The first step in your virtual desktop project should be to understand and prioritize the strategic imperatives of the organization. This enables the project management team to define success criteria and allows the design team to create a tailored and optimized architecture.

Requirements can be captured during meetings or by distributing questionnaires. Meetings are more time consuming, but allow for follow-up questions to be asked and help to simplify the prioritization process. It is important that this exercise be completed jointly by both business managers and IT decision makers since both groups will have significantly different viewpoints. Take the

following examples of what certain organizations faced, which drove their selection of desktop virtualization.

Experience from the Field

Finance – A large financial institution had a base of operations in the city designated as the host city for an upcoming G8 summit. As these types of meetings historically include riots, protests and other issues that can disrupt business and the safety of their employees, the financial organization needed an alternative allowing their users to work from the safety of their homes.

Agriculture – Due to thin margins, an agriculture organization wanted to save money by extending the life of desktop PCs while still being able to run the latest applications.

Healthcare – A large healthcare organization was in need of a solution to simplify application updates as the main application required updates on a weekly basis. Due to the distributed nature of the endpoint devices, the organization was in need of a better application delivery solution.

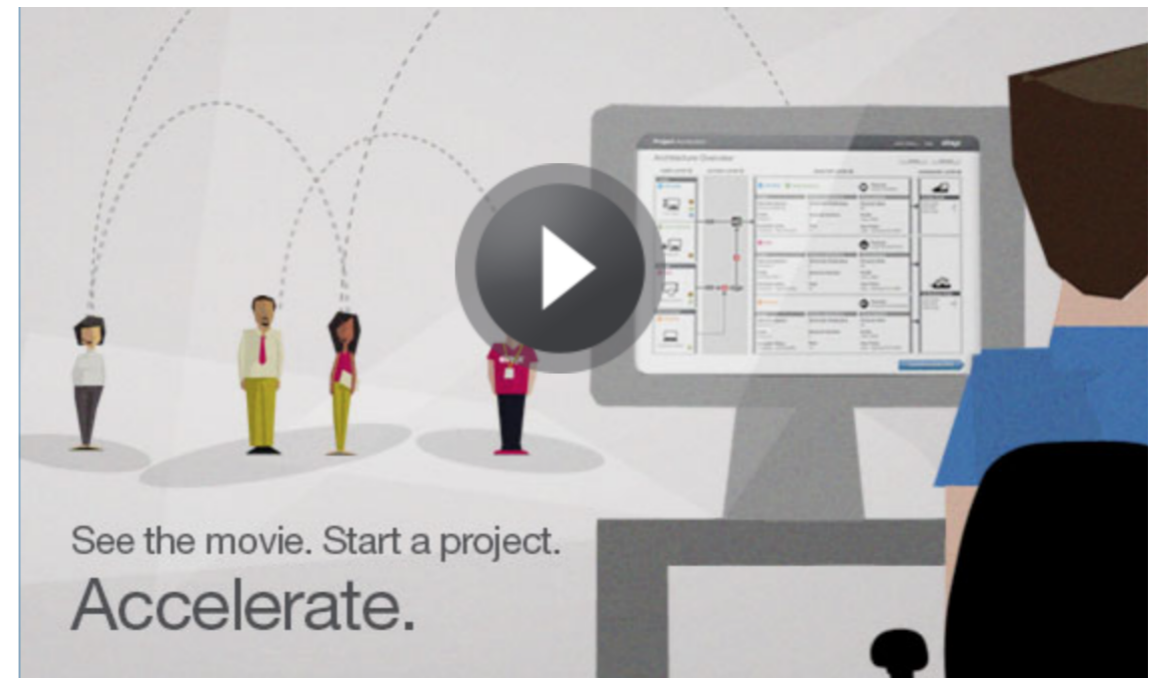
These are just a few examples, but they demonstrate how organizations think about their priorities. Most organizations do not focus on technology, they focus on the needs of the user and of the organization. These needs can be met with technical solutions but it is imperative the team understands the “Why” of the project.

In addition to the three real-world examples, the table on the next page identifies a few other priorities often stated from many organizations

Requester	Requirement
Business Managers	Better IT agility: Flexible desktop solution that is capable of accommodating periods of change such as rapid growth or downsizing. For example, enabling the business to setup project offices or temporary points of sale very rapidly without long delays or IT notification periods.
	Bring your own device: Empower employees to choose their own devices to improve productivity, collaboration and mobility.
	Collaboration: With an increase in both globalization and mobility, team members are often dispersed across multiple physical locations. Powerful collaboration capabilities are required to ensure high levels of productivity, efficiency and quality.
IT Decision Makers	Work from anywhere: The business needs to support home workers in order to attract and retain top talent, and / or traveling employees.
	Better desktop management: Simplify the management of desktop infrastructure. IT is not as proactive as they would like and spend too much time “fighting fires”.
	Increase security: Data theft or the loss of devices containing sensitive data is a big risk and preventive measures are a top priority.
	Extend desktop hardware lifecycle: Replacing workstations every three to five years in order to keep up with the requirements of the operating system or the applications has been very costly.
	Reducing cost: Cut costs associated with supporting and maintaining traditional desktops.

The prioritization process should be completed in collaboration with the project team, business managers and IT managers so that all views are considered.

Early on, organizations often need to estimate the hardware required to support the new solution. [Citrix Project Accelerator](#) provides a fast and accurate way to estimate hardware and storage requirements:



Step 2: Assess the Environment

The data capture process collects key information on users, devices, applications and infrastructure that will be required during subsequent phases of the project.

Data Capture Strategy

There are three main techniques that can be used to collect data:

- **Manual** - For small organizations, it is possible to capture user and application data by visiting each desktop in person or by connecting remotely. Performance counters can be used to acquire raw user data while the Add / Remove Programs applet in Windows XP or the Programs and Features applet in Windows 7 and Windows 8 can be used to provide a list of installed applications as well as information on how often they are accessed.

Most medium and large-sized organizations use an enterprise software deployment (ESD) tool such as Microsoft System Center Configuration Manager (SCCM). Since the ESD solution is typically the sole method of deploying application packages to workstations, the project team can query the ESD software to identify which applications are assigned to which desktops. However, most ESD solutions do not provide details on application performance requirements, application usage metrics or user installed applications. In addition, significant effort may be required to extrapolate, analyze and document application and desktop information from the ESD tool.

The problem with the manual approach is that it is difficult to gain a good understanding of each user and application's performance requirements over time. In addition, the manual approach is very time intensive per desktop making it

inappropriate for medium and large-sized organizations. It can also be difficult to query application requirements of mobile users who may spend extended periods of time away from the office.

- **Survey** - A questionnaire can be created and distributed to each business manager to identify user and application requirements within their departments. Questionnaires are far less time consuming for the project team than the manual approach because the focus is upon the department managers to identify the applications necessary for their employees to complete their jobs. However, it is unlikely that every manager within the organization will complete the questionnaire and a completion ratio of 70% should be considered normal. The results from a questionnaire are typically less accurate than the manual or automated approach. Although less time consuming than the manual approach, considerable effort is still required to extrapolate, analyze and document the results from the returned questionnaires.
- **Automated** - There are a variety of automated inventory tools available that allow for a wide range of information to be collected, including applications installed, access frequency, and performance requirements. This data is uploaded to a central database so that reports can be generated across the entire user and desktop estate. To help reduce the cost and complexity of desktop virtualization projects, Citrix has partnered with Lakeside Software to provide [Citrix Project Accelerator](#) users with a free 60-day license of Lakeside FastTrack. FastTrack is a sophisticated inventory tool that has been developed based on Citrix methodologies, terminology and best practices. An example [LakeSide FastTrack report](#) is available from the LakeSide website. An automated inventory tool is a great solution for medium and large-sized organizations; however the centralized infrastructure and agent deployment effort is unlikely to be appropriate for very small organizations due to the time required when compared to the manual method. In

addition, some organizations may not allow inventory agents to be installed across the desktop estate or may not have rights to install the agent on user-owned desktops.

The advantages and disadvantages of each approach are summarized in the following table:

Comparison of Data Capture Strategies

Approach	Data for all Characteristics	Results Returned	Admin Time Per Desktop	User Involvement
Manual	No	~80%	Long	Likely
Survey	Yes	~70%	Medium	Yes
Automated	No	~100%	Short	No

Although the automated method is accurate, fast, and does not inconvenience employees, there are a number of business characteristics that automated tools cannot identify. For example, what is the criticality of the user, will the user need to work offline and should access to the application be restricted due to security or licensing restrictions? Therefore, the recommended approach is to use the automated method to identify technical characteristics and a questionnaire to identify the business characteristics.

The [User Segmentation](#) and [Link Apps to Users](#) sections provide detailed information on the user and application characteristics that you will need to collect.

Key results from the user data gathering exercise should be documented in the [assess spreadsheet](#).

Citrix Consulting Tips for Success

1. **Representative users** – If you don't have enough time, resources, or licenses to inventory every desktop in your organization, make sure that you pick a representative subset of users. For example, deploying an agent on every desktop in the HR department but missing out the sales and finance departments will impact your results. Take time to ensure that you

select a representative subset of users from every department and role within the organization. And remember, not all users within a single department will have the same requirements.

2. **Check the data** – When using an automated inventory tool, regularly check that the data has been successfully uploaded to the central database. There have been cases reported where insufficient space has been allocated to the central database resulting in several weeks of agent data being lost causing the project to be delayed.
3. **Monitoring period** – It is extremely important that the automated inventory tool monitors the desktops over a sufficient period of time. At the very least, monitor the desktops for a minimum of one month; ideally monitor usage over periods of peak activity such as quarter end so that you have a good idea of average and peak application performance requirements. There may be situations where users only need access to a specific application during quarter end and unless you are running the automated inventory tool at this time you will not be aware of its usage.
4. **Remember the plugins** – Plugins are available for various applications, including Microsoft Internet Explorer, Microsoft Word and Microsoft Outlook. For example Citrix ShareFile and e-mail archiving solutions are frequently implemented as plugins. To prevent critical functionality being omitted, plugins should be treated as applications during the application assessment.
5. **Application dependencies** – It is imperative that you understand all of the interactions between your applications. For example, there may be an application which requires Microsoft Word and Excel be installed on the same system so that reports can be viewed. When it comes to the design phase of the project you will need to make sure that these applications are grouped together appropriately.
6. **Application consolidation** – It may be tempting to skip through the application consolidation phase of the assessment, but time

spent reducing the number of applications significantly reduces complexity and time spent on the assessment and design phases of the project.

- 7. Application owners** – To ensure that you don't incorrectly categorize applications or remove an important application from your inventory, work closely with the various application owners during the Rationalization, Characteristics and Compatibility Steps.
- 8. Final check** – Once the consolidated list of applications has been finalized, complete with characteristics and compatibility information it should be sent to the application owners for review to ensure that it is correct and no critical applications or information has been omitted.

Capabilities Assessment

The information captured during the capabilities assessment will be used to achieve the following objectives:

- 1. Identify risks** – Like traditional desktops, desktop virtualization is dependent on a wide range of supporting technologies, including storage, networking, directory services and applications. In many cases, issues reported with virtual desktops are a symptom rather than a cause. For example, slow performance and disconnected sessions are more often caused by periods of high-latency than a desktop specific issue. Therefore, a desktop virtualization project is an excellent opportunity to review an organization's existing infrastructure to ensure that it provides a good solid foundation upon which to build a virtual desktop solution. Any risks identified should be reported to the project manager so that they can be addressed appropriately.
- 2. Create roadmap** – The capabilities assessment provides the project team with a detailed understanding of the existing environment so that they can estimate implementation time for

each user group and prioritize the implementation order.

- 3. Virtual desktop design** – Provide the project team with detailed information on the current state of the environment so that they can successfully integrate the new virtual desktop solution. The capabilities assessment should also determine whether existing infrastructure components can be leveraged or whether new infrastructure needs to be purchased, for example shared storage and virtualization technologies.

Key results from the capabilities gathering exercise should be documented in the [Citrix Capabilities Assessment](#) template. A list of questions has been provided per infrastructure technology to highlight key information that needs to be collected. These questions are based on the experiences of Citrix Consulting across numerous desktop virtualization projects and should be answered by the appropriate technology architect for each section. Not all sections need to be completed, for example if the organization does not already have a Citrix environment this section can be left blank. The length of the capabilities assessment will vary based on the size and complexity of the environment but typically takes about three days to complete.

Citrix Consulting Tips for Success

- 1. Discussions** – Meet with the architects rather than sending them a questionnaire so that additional questions can be raised, diagrams drawn and detailed explanations provided.
- 2. Schedule meetings** – It is advisable to schedule meetings with the architects well in advance to ensure availability. Provide each architect with a copy of the questions that you plan to ask them so that they can prepare appropriately. Also, when scheduling the meetings request any relevant documentation for background reading as this may prompt additional questions and discussions.

3. **Documentation** – Use the capabilities assessment template to capture your discussions with the architects. The document can then be circulated amongst the project team to ensure that everybody has the same information.
4. **Future initiatives** – Ask the architects whether they are planning any future initiatives, for example upgrading to a new product version, or adjusting the architecture. This will ensure that the project team is aware of all upcoming changes.
5. **Identify risks** – It is important that all risks are identified as early as possible so that they can be tracked and appropriate action taken. Risks should be graded in severity and if possible remediation plans should be created stating an estimated timeframe and cost.

Step 3: Define the User Groups

Once the data capture is complete, you're ready to start dividing up the users into different groups based on a common set of requirements. This allows a FlexCast model to be assigned to each user group without compromising on performance or functionality.

User Segmentation

Users are often classified as task workers, branch office workers, mobile workers and the like. Unfortunately, this classification is too broad to offer meaningful segmentation because users can simultaneously be described as task workers, mobile workers, and branch office workers. Instead, group users together that have the same requirement for:

- **Primary datacenter (B)** – Each user will have a primary datacenter assigned that will be used to host their virtual

desktop, data, and application servers. Identify the datacenter that the user should be assigned to rather than the datacenter they are currently using. Users will be grouped based on their primary datacenter so that a unique design can be created for each one.

- **Personalization (B)** – Personalization requirements are used to help determine the appropriate FlexCast model for each user group. For example, if a user group requires complete personalization, a Hosted VDI desktop with Personal vDisk will be recommended as the optimal solution. There are three classifications available:

Personalization Characteristics

Personalization	Requirement
None	User cannot modify any user or application settings, for example - kiosk.
Basic	User can modify user-level settings of desktops and applications.
Complete	User can make any change, including installing applications.

- **Security (B)** – Security requirements are used to help determine the appropriate FlexCast model and policies for each user group. For example, if a user group requires high security, a Hosted-Shared Desktop, Pooled Desktop or On-Demand Apps FlexCast model will be recommended as the optimal solution. There are four classifications available:

Security Characteristics

Security Level	Data Stays in Datacenter or is Encrypted	Prohibit User Installs	Allow Multi-User Operating Systems	MAC / IP Address Auditing
Low	No	No	Yes	No
Medium	Yes	Yes	Yes	No
High	Yes	Yes	No	No
High + Audit	Yes	Yes	No	Yes

- **Mobility (B)** – Mobility requirements are used to help determine the appropriate FlexCast model for each user group. For example, if a user group sometimes connects remotely, the Streamed VHD FlexCast model will not be selected because it requires a high-speed local network. There are four classifications available:

Mobility Characteristics

Mobility	Requirement
Local	Always connects to an internal, high-speed and secured network.
Roaming Locally	Connects from different locations on an internal, high-speed, secured network.
Remote	Sometimes connects from external variable-speed, unsecure networks.
Mobile	Often needs access when the network is intermittent or unavailable.

- **Desktop Loss Criticality (B)** – Criticality will be used to determine the level of high availability, load balancing and fault tolerance measures required. There are three classifications available:

Desktop Loss Criticality Characteristics

Criticality	Requirement
Low	No major risk to products, projects or revenue.
Medium	Potential risk to products, projects or revenue.
High	Severe risk to products, projects or revenue.

- **Workload (T)** – Collecting user performance requirements will allow the desktop virtualization infrastructure to be accurately sized and an appropriate FlexCast model to be selected.

Workload Characteristics

User Type	Characteristics
Light	1-2 office productivity apps or kiosk.
Medium	2-10 office productivity apps with light multimedia use.
Heavy	Intense multimedia, data processing or application development.

Note: Performance thresholds are not identified based on processor, memory or disk utilization because these characteristics will change dramatically following the application rationalization and desktop optimization process. In addition, it is likely that the user's management tools and operating system will change during the migration process. Instead, workload is gauged based on the number and type of applications the user runs.

(T) Technical Characteristic

(B) Business Characteristic

Experience from the Field

Utility company – A large utility company collected data on every user in their organization. During the user segmentation process it was realized that the organization's existing role definitions were sufficiently well defined that all the users within a role shared the same requirements. This allowed a significant amount of time to be saved by reviewing a select number of users per group.

Government – A government organization discovered that there was significant deviation between user requirements within each role, particularly around security and criticality. As such, each user needed to be carefully reviewed to ensure that they were grouped appropriately.

The fastest and easiest way to identify your user groups within the user assessment workbook is to filter the results based on these key requirements. Once you have identified the users within one group, transfer the relevant information to the user segmentation worksheet within the [assess spreadsheet](#).

Assign FlexCast Models

As with physical desktops, it is not possible to meet every user requirement with a single virtual desktop type. Different types of users need different types of desktops. Some users may require simplicity and standardization, while others may require high levels of performance and personalization. Implementing a single desktop virtualization model across an entire organization will inevitably lead to user frustration and reduced productivity.

Citrix FlexCast offers a complete set of application and desktop virtualization technologies that have been combined into a single integrated solution. Because each FlexCast model has different advantages and disadvantages, it is important that the right model is chosen for each user group within the organization.

There are five FlexCast models available, the advantages and disadvantages of each model are described below:

- **Hosted shared** – With the hosted shared FlexCast model, multiple user desktops are hosted on a single server-based operating system and provisioned using Machine Creation Services or Provisioning Services. The hosted shared desktop model provides a low-cost, high-density solution, however applications must be compatible with a multi-user server based operating system. In addition, because multiple users are sharing a single operating system, users are restricted from performing actions that negatively affect other users, for example installing applications, changing system settings and restarting the operating system. There is also the potential that a single user could consume an unfair share of resources, which may negatively affect other users. The hosted shared FlexCast model is provided by Citrix XenDesktop in combination with Microsoft Remote Desktop Services (RDS).
- **Hosted VDI** – The hosted VDI FlexCast model provides each user with a desktop operating system. Hosted VDI desktops are less scalable than hosted shared desktops because each

user requires their own operating system. However, hosted VDI desktops remove the requirement that applications must be multi-user aware and support server based operating systems. In addition, the hosted VDI model provides administrators with a granular level of control over the number of virtual processors and memory assigned to each desktop. The hosted VDI model is provided by Citrix XenDesktop, and offers the following sub categories:

Random / Non-Persistent – Desktops are based on a single master image and provisioned using Machine Creation Services or Provisioning Services. Users are dynamically connected to one of the desktops in the pool each time they logon. Changes to the desktop image are lost upon reboot.

Static / Non-Persistent – Desktops are based on a single master image and provisioned using Machine Creation Services or Provisioning Services. Users are allocated a virtual desktop on first access. Once assigned, users will always be connected to the same virtual desktop. Changes to the desktop image are lost upon reboot.

Static Persistent – Desktops are based on a single master image and provisioned using Machine Creation Services or Provisioning Services. Users are allocated a virtual desktop on first access. Once assigned, users will always be connected to the same virtual desktop. Changes to the desktop are stored in a personal vDisk and retained between reboots. Desktops with a personal vDisk cannot be shared between multiple users; each user requires their own desktop. If high availability is required, the personal vDisk must be stored on shared storage.

- **Remote PC** – Physical desktops that have already been deployed. These desktops must be managed manually or with 3rd party desktop management tools.
- **Streamed VHD** – Desktops are based on a single master image

and provisioned using Provisioning Services. The streamed VHD FlexCast model allows Windows XP, 7 and 8 desktops to be run locally on the user's desktop computer. Streamed VHD is a great solution for high-end workstations because it allows them to leverage local processing power. Streamed VHD requires a LAN connection to be in place between the desktop and the provisioning servers and changes to the desktops are lost upon reboot.

- **Local VM** – Windows XP, 7, and 8 desktops running locally within a hypervisor. The virtual desktop image is completely delivered to the hypervisor to allow for offline connectivity. Citrix XenClient is used to provide the Local VM FlexCast model.
- **On demand apps** – The On-Demand Apps FlexCast model does not provide users with a virtual desktop; instead Windows applications are centralized in the datacenter, and instantly delivered via a high-speed protocol (requires connection) or streamed (offline support) via [Microsoft App-V](#).

The following table compares the different FlexCast models available:

FlexCast Model Comparison

FlexCast Model	User Installed Apps	Image Delivery Technology	Virtual / Physical	Access	Desktop to User Ratio	
Hosted shared: Non-Persistent	No	MCS / PVS	Physical / Virtual	HDX	1:Many	
Hosted VDI:	Random / Non-Persistent	No	MCS / PVS	Virtual	HDX	1:Many
	Static / Non-Persistent	No	MCS / PVS	Virtual	HDX	1:1
	Static / Persistent	Yes	MCS / PVS	Virtual	HDX	1:1
Remote PC	Yes	Installed	Physical	HDX	1:1	
Streamed VHD	No	PVS	Physical	Local	1:1	
Local VM	Yes	XC	Virtual (XenClient)	Local	1:1	
On demand apps	No	MCS / PVS	Physical / Virtual	HDX	1:Many	

Each user group in the User Segmentation worksheet should be compared against the following table to determine which FlexCast Model should be assigned. Ensure that you update the FlexCast value for each user group in the worksheet.

Note: Any FlexCast decisions need to factor in latency. Please see the [Latency section](#) of the Design Phase.

FlexCast Model Capability Comparison

Segmentation Characteristic	Hosted Shared: Non Persistent	Random / Non-Persistent	Hosted VDI		Remote PC	Streamed VHD	Local VM: Persistent	Local VM: Non-Persistent	On Demand Apps
			Static / Non-Persistent	Static / Persistent					
Workload									
Light	•	o	o	o	o	o	o	o	•
Medium	•	•	o	o	o	o	o	o	•
Heavy	x	•	o	o	o	•	•	•	x
Mobility									
Local	•	•	o	o	o	•	o	o	•
Roaming Local	•	•	o	o	o	x	o	o	•
Remote	•	•	o	o	o	x	o	o	•
Mobile	x	x	x	x	x	x	•	•	x
Personalization									
None	•	•	o	x	x	•	o	o	•
Basic	•	•	o	x	x	•	o	o	•
Complete	x	x	x	•	o	x	•	x	x
Security									
Low	•	•	o	o	o	o	o	o	•
Medium	•	•	o	o	o	o	o	o	•
High	x	•	•	x	x	x	•	•	x
High + Audit	x	x	•	x	x	x	•	x	x
Desktop Loss Criticality									
Low	•	•	o	o	o	o	o	o	•
Medium	•	•	o	o	x	o	o	o	•
High	•	•	x	x	x	x	x	x	•

“•”: Recommended “o”: Viable “x”: Not Recommended

Citrix Consulting Tips for Success

1. **Lead with hosted shared/VDI** – As you can see in the FlexCast capability table above, the hosted VDI and hosted shared FlexCast models can be used in the majority of situations. The streamed VHD and local VM FlexCast models should only be used on an exception basis. By reducing the number of FlexCast models required, you will help to reduce deployment time and simplify management.
2. **Perfect match** – It may not be possible to select a FlexCast model which is a perfect match for your user group, for example you can't provide users with a desktop that is highly secure and offers complete personalization at the same time. In these situations, select the FlexCast model which is the closest match.
3. **Desktop loss criticality** – There are only three FlexCast models that meet the needs of a high criticality user group (backup desktops available) – none of which allow for complete personalization. If a high-criticality user group also requires the ability to personalize their desktop they could be provided with a pool of backup desktops (hosted shared, pooled, streamed) in addition to their primary desktop. Although these desktops would not include customizations made to their primary desktop, they would allow users to access core applications such as mail, Internet and Microsoft Office.

Step 4: Define the Applications

Once the users have been divided up in to groups the next step is to determine which applications they require. This is a two-step process:

1. **Application rationalization** – Help to simplify the application assessment by removing redundant applications from the inventory that were captured during the data capture.
2. **Link apps to users** – Use the results from the data capture process to map applications to user groups.

Application Rationalization

The number of applications identified during the inventory is often surprising, even for organizations that believe they have a high-level of control over applications. To help reduce complexity as well as overall time required, it's important to take the time to consolidate the list of applications. Start by arranging an application assessment meeting with all relevant application owners.

Note: Consolidated applications should be identified within the application assessment worksheet by selecting consolidated in the status column. Consolidated applications should not be removed from the spreadsheet so that the rationalization process can be reviewed within the organization.

The following guidelines will help to ensure that your application list is consolidated appropriately:

- **Multiple versions** – Different versions of the same application may have been identified during the inventory. There are various reasons for this, including an inconsistent patching or upgrade process, decentralized application management, limited licenses and situations where users require specific application versions for compatibility with other applications, macros and document

formats. Where possible, work with the application owners to reduce the number of versions required. The best practice is to standardize on a single version of each application, typically the latest.

- **Business applications** – Applications, which are not required by the business, should be removed from the application inventory to reduce resource requirements and to help simplify the overall project. Non-business related applications are typically found in an application inventory when users have been provided with the ability to install their own applications and typically include games, communication clients, screen savers, peripheral software and media players.
- **Legacy applications** – The inventory may identify legacy applications that have since been retired or that are no longer required within the business. These applications may not have been removed from the desktops because there is no established process to do so or because there are always more high-priority activities to complete. These applications should be consolidated during the rationalization stage of the application assessment.
- **Management applications** – The antivirus, application delivery, monitoring, inventory, maintenance and backup applications will be completely re-designed across the organization during the desktop virtualization project. These applications should also be consolidated during this stage.

Experience from the Field

Government: A government organization identified that there were 2,660 applications installed across their desktop estate. Most of which were installed by users with local administrative rights. By following the application rationalization recommendations above, it was possible to reduce the number of applications required to 160.

[Click here to provide feedback](#)

Link Apps to Users

Once the application rationalization process is complete, use the results from the data capture process to identify which applications will be required by each user group. The following characteristics should be identified for each application so that the right application delivery model can be selected during the design phase of the project:

- **Workload (T)** – Collecting application workload requirements will allow the virtualization infrastructure to be sized and an appropriate application delivery model to be selected. For example, resource intensive applications will not be delivered via a Hosted Shared Desktop because there is limited control over how the resources are shared between users. There are two classifications available in the user assessment worksheet:

Application Worksheet Characteristics

Workload	Requirement
Resource Intensive	Application requires 1GB+ of RAM or averages 50%+ CPU utilization.
None	The application is not resource intensive.

- **Technically challenging (B)** – An application should be classified as technically challenging if it is complex to set up, has extensive dependencies on other applications or requires a specialized configuration, for example an Electronic Medical Records (EMR) application. Technically challenging applications need to be identified during the application assessment because they are not generally appropriate for installation in to a base desktop image or delivery by application streaming. Delivering technically challenging applications as published applications will help to reduce the complexity of the base desktop image.
- **Works offline (B)** – Some user groups may require the ability to work offline. If so, it is important that the design can determine

which applications will work without a network connection and which ones will not. Applications that require backend infrastructure such as web and database servers are not typically available offline.

- **Peripheral connectivity (T)** – If applications require connectivity with peripheral devices, identify the interface required so that it can be made available to the application when it is run from a virtual session.

Note: Generic USB redirection is now supported in Hosted Shared Desktops delivered through XenApp 7.6. XenApp and XenDesktop 7.6 are also USB 3.0 ready. A future release of Citrix Receiver for Windows and Linux will include automatic device mapping and plug-and-play capabilities for USB 3.0 devices.

Note: USB redirection works for XenApp and XenDesktop 7.6 and above. For more information please refer to Citrix eDocs – [USB and client drive considerations](#).

- **Restricted access (B)** – Application access may need to be restricted due to insufficient licenses / resources and to protect sensitive data / tools. For example, applications with a limited number of licenses should not be installed in to a base image that is shared with unlicensed users. There are three restricted access categories in the application assessment workbook:

Restricted Access Characteristics

Restricted Access	Requirement
No	There are no security restrictions for the application and it can be accessed by any user within the organization
User Group	The application may be installed on a multi-user operating system but only a specific group of users should be provided with an icon
Virtual Machine	Application should only be installed on a virtual machine that is accessible by authorized users

(T) Technical Characteristic

(B) Business Characteristic

[Click here to provide feedback](#)

Step 5: Project Management Roadmap

Most companies don't have sufficient time or resources to migrate every user group in one go. As such, it is important that the user groups identified are prioritized so that the business receives the maximum value from their investment as soon as possible. To achieve this, you need to compare the business impact and time to value of each group:

- **Business impact** – Consider the impact that desktop virtualization will have on each user group and rank them accordingly. It is important that you double back here, and use the [business drivers](#) identified at the start of the project to make sure that you assign an appropriate ranking. Don't just assign ratings based on how highly the business values each user group; instead focus on which user groups offer the most benefit to the company after virtualization. It's a subtle but important difference.
- **Time to value** – For each user group, estimate how long it will take to implement the chosen FlexCast model based on the findings from the Capabilities Assessment. For example, you might find that a company already has a XenDesktop solution that can be leveraged to support those user groups that require a hosted VDI desktop resulting in a low time to value. Alternatively, the business might have no prior experience with XenClient resulting in a longer time to value. Compare application sets, user requirements and user numbers when differentiating between user groups that have been assigned the same FlexCast model.

Note: If there are limited skills available in-house to implement a chosen FlexCast model, consider hiring external resources so that Time to Value can be reduced for the associated user groups.

Representing this information in a graph provides an easy way to visualize the results:



When it comes to assigning the implementation order, start from the top left hand corner of the chart and work your way to the bottom right hand corner. This way you start with some quick wins that offer a high-level of value to the company.

Once the project roadmap has been created, update the project plan so that it incorporates the prioritized roadmap.

Experience from the Field

Utility company – A large utility company realized that there would be a long time to value for user groups that had been assigned with a hosted VDI FlexCast mode, because they had no prior experience or training with this technology. To address this concern, the utility company engaged with Citrix Consulting who provided Consultants with previous experience of successfully deploying XenDesktop.

Build the Right Team

Desktop virtualization is a fundamental change that requires close collaboration between various business and technical teams in order to be successful. For example, the virtualization and desktop teams need to work together to ensure that the virtual desktop

image meets user needs while also being optimized for the datacenter. Failure to build a cohesive project team that consists of the right roles and skill sets can negatively impact performance, availability, user experience and supportability while also increasing costs and risk.

The following tables identify the business and technical roles required during an enterprise virtual desktop deployment. Although the list may seem quite large, many of these roles are only required for a short time and multiple roles may be performed by a single person. The project manager and Citrix Architect are considered to be full time roles with other team members being brought in only when required. The project manager role is key to ensuring that the right people are involved in the project at the right time.

Business and Technical Roles

Role	Description	Example Responsibilities
Business Roles		
Project Sponsor	The project sponsor is a senior company executive who recognizes the benefits that desktop virtualization will bring to the business. The project sponsor role is often performed by the chief technology officer (CTO).	<p>Pre-project</p> <ul style="list-style-type: none"> Promote desktop virtualization within business Identify members of the steering committee <p>Secure funding</p> <ul style="list-style-type: none"> Assess Identify and prioritize key business drivers
Project Manager	The project manager directs the project team and is responsible for ensuring that project objectives are completed on time and within budget.	<p>All steps</p> <ul style="list-style-type: none"> Define key project milestones Create and update project plan Track progress against plan Track expenditure against budget Maintain issue and risk register Manage scope changes Create weekly project reports Brief steering committee on progress Organize project workshops and meetings Ensure project teams are synchronized Ensure pre-requisites are in place Creates change control requests
Business Manager	Depending on company structure and size, business managers oversee planning and performance at a department, region or company level. A business manager will understand the requirements necessary for their employees to be successful.	<p>Assess</p> <ul style="list-style-type: none"> Assist with application consolidation project Provide details on connectivity requirements of user group, including offline usage Provide details on risk tolerance of user group Identify requirements for peripherals <p>Deploy</p> <ul style="list-style-type: none"> Promote benefits of desktop virtualization Assist with coordinating the rollout
Business Continuity Manager	The business continuity manager ensures that an organization can continue to function after a disruptive event such as natural disaster, crime or human/computer error.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of the current business continuity plan <p>Design</p> <ul style="list-style-type: none"> Update business continuity plan to incorporate the new Citrix infrastructure <p>Deploy</p> <ul style="list-style-type: none"> Test business continuity plan
Test Manager	The test manager is responsible for ensuring that the test and user acceptance environments match the production environment as closely as possible. The test manager helps to reduce risk by ensuring that changes are fully tested before being implemented in production.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current testing infrastructure and processes <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design an appropriate testing infrastructure and test plan for new Citrix environment <p>Deploy</p> <ul style="list-style-type: none"> Ensure that testing design is implemented correctly and new Citrix infrastructure is fully tested before rollout

Business and Technical Roles

Role	Description	Example Responsibilities
Application Owners	An application owner is a subject matter expert on specific applications deployed within the business. Application owners are responsible for ensuring that problems with the applications are resolved and that upgrades/updates are performed without issue. Application owners are also responsible for managing support agreements with the application vendors.	Assess <ul style="list-style-type: none"> Assist with application consolidation project Identify application licensing limitations Provide details on security restrictions Provide details on application dependencies Provide location of backend resources Deploy <ul style="list-style-type: none"> Provide installation pre-requisites and install guide Assist Citrix team with installing and testing applications in VDI environment
Service Desk Manager	The service desk manager helps to improve productivity and end-user satisfaction by ensuring that production issues are logged, escalated and resolved in a timely manner. The service desk manager is also responsible for reporting on common issues, call volumes and service desk performance.	Assess <ul style="list-style-type: none"> Identify common issues with existing environment Provide details on support tools currently used Design <ul style="list-style-type: none"> Assist Citrix architect with designing a delegated administration model Participate in operations and support design workshops Work with training manager to identify training requirements Deploy <ul style="list-style-type: none"> Monitor helpdesk calls for rollout related issues
Training Manager	The training manager ensures that support staff and end-users are proficient with new areas of technology. The training manager also has responsibility for ensuring that the training plan is up-to-date and followed appropriately.	Assess <ul style="list-style-type: none"> Determine current skill set for support staff and end users Design <ul style="list-style-type: none"> Create training plan for support staff and end users Deploy <ul style="list-style-type: none"> Implement training plan for support staff and end users
Communications Manager	The communication manager is responsible for disseminating key information throughout the organization.	Design <ul style="list-style-type: none"> Work with project manager to create communications plan Deploy <ul style="list-style-type: none"> Relay benefits of desktop virtualization Inform users of key migration dates Ensure expectations are set accordingly
Technical Roles		
Citrix Desktop Architect	The Citrix architect will act as the design authority for all Citrix products and will liaise with other architects to ensure that the Citrix infrastructure is successfully integrated into the organization.	Assess <ul style="list-style-type: none"> Work with project sponsor and key stakeholders to identify and prioritize key business drivers Oversee user segmentation and app. assessment Map FlexCast models to user groups Perform capabilities assessment to determine current state of readiness Identify areas of risk and provides remedial actions Design <ul style="list-style-type: none"> Create Citrix design which includes hardware and storage estimates Coordinate with other architects to integrate Citrix infrastructure into organization Work with monitoring architect to ensure that Citrix environment is monitored appropriately Create operations and support design Create implementation and rollout design Create test plan Deploy <ul style="list-style-type: none"> Ensures that the Citrix environment is implemented in accordance with design Verifies that implementation passes test plan Ensures that the Citrix design is implemented correctly

Business and Technical Roles

Role	Description	Example Responsibilities
Active Directory Architect	Design authority on Microsoft Active Directory, including Organizational Units (OU) and Group Policy Objects (GPOs).	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current Active Directory architecture <p>Design</p> <ul style="list-style-type: none"> Work with the Citrix architect to design OU structure, group policies, permissions, service accounts, etc. for new Citrix environment Update Active Directory infrastructure design to reflect centralization of user data and accounts <p>Deploy</p> <ul style="list-style-type: none"> Ensure that Active Directory design is implemented correctly
Virtualization Architect	Design authority on server and desktop virtualization using Citrix XenServer, Microsoft Hyper-V or VMware vSphere.	<p>Assess</p> <ul style="list-style-type: none"> Provides Citrix architect with detailed understanding of current virtualization architecture <p>Design</p> <ul style="list-style-type: none"> Works with Citrix architect to design hardware, networking, storage, high availability, etc. for server and desktop virtualization Work with monitoring architect to ensure that virtualization environment is monitored appropriately <p>Deploy</p> <ul style="list-style-type: none"> Ensures that the virtualization design is implemented correctly
Network Architect	Design authority on networking, including routing, VLANs, DHCP, DNS, VPN and firewalls.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current networking architecture <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design physical network, virtual networks, routing, firewalls, quality of service, remote access, network optimization, etc. for new Citrix environment Work with monitoring architect to ensure that network is monitored appropriately <p>Deploy</p> <ul style="list-style-type: none"> Ensure that network design is implemented correctly
Desktop Architect	Design authority on Microsoft desktop operating systems, including Windows XP, Windows 7 and Windows 8.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current desktop environment <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design core desktop virtual image, core applications, desktop optimizations, etc. for new Citrix environment Work with monitoring architect to ensure that the virtual desktops are monitored appropriately <p>Deploy</p> <ul style="list-style-type: none"> Ensure that desktop design is implemented correctly
Storage Architect	Design authority on storage solutions, including direct-attached storage, storage-attached networks and network attached storage.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current shared storage environment <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design storage architecture, tiers, sizing, connectivity, etc. for new Citrix environment Work with monitoring architect to ensure that storage is monitored appropriately <p>Deploy</p> <ul style="list-style-type: none"> Ensure that storage design is implemented correctly
Backup Architect	Design authority on backup and recovery, including virtual machines, desktops, servers, user data and databases.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current backup architecture and processes <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect and disaster recovery architect to design backup architecture, process, schedule, retention, etc. for new Citrix environment <p>Deploy</p> <ul style="list-style-type: none"> Ensure that backup design is implemented correctly

Business and Technical Roles

Role	Description	Example Responsibilities
Application Packaging Architect	Design authority on packaging applications for deployment via the systems management team.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix Architect with detailed understanding of current application packaging process and status <p>Deploy</p> <ul style="list-style-type: none"> Ensure that all required applications are packaged according to design
Monitoring Architect	Design authority on monitoring, including hardware, network, servers, storage and security appliances.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current monitoring architecture and processes <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design monitoring architecture, metrics, alerts, etc. for new Citrix environment and supporting infrastructure <p>Deploy</p> <ul style="list-style-type: none"> Ensure that monitoring design is implemented correctly Provide regular reports on capacity and trends during rollout
Systems Management Architect	Design authority on systems management, including server/desktop build process, patching and automated application installation.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with a detailed understanding of the current systems management processes <p>Design</p> <ul style="list-style-type: none"> Works with Citrix architect to define server/desktop build process, patching and application delivery strategy for new Citrix environment <p>Deploy</p> <ul style="list-style-type: none"> Ensure that the systems management design is implemented correctly
Security Architect	Design authority on security, including desktops, servers, networks and VPNs.	<p>Assess</p> <ul style="list-style-type: none"> Provide Citrix architect with detailed understanding of current security policy <p>Design</p> <ul style="list-style-type: none"> Work with Citrix architect to design security standards for new Citrix environment, including authentication, encryption, port numbers, firewall rules, etc. <p>Deploy</p> <ul style="list-style-type: none"> Ensures that security design is implemented correctly

Design Quick Access Links

Design Overview	27	Resource Allocation	84
User Layer	27	Bandwidth Requirements	87
Endpoint Selection	27	Control Layer	92
Receiver Selection	30	Infrastructure Controllers	92
Access Layer	34	Resource Controllers	107
StoreFront	37	Image Controllers	123
NetScaler Gateway	43	Hardware Layer	140
Resource Layer	55	Hardware Sizing	140
Personalization	55	Hypervisors	146
User Profiles	55	Hyper-V 2008 R2	146
User Policies	60	Hyper-V 2012 R2	155
Printing	65	Hardware	155
Personal vDisk	72	Host Scalability	159
Applications	73	Networking	160
Images	80	Network Performance Tuning	166

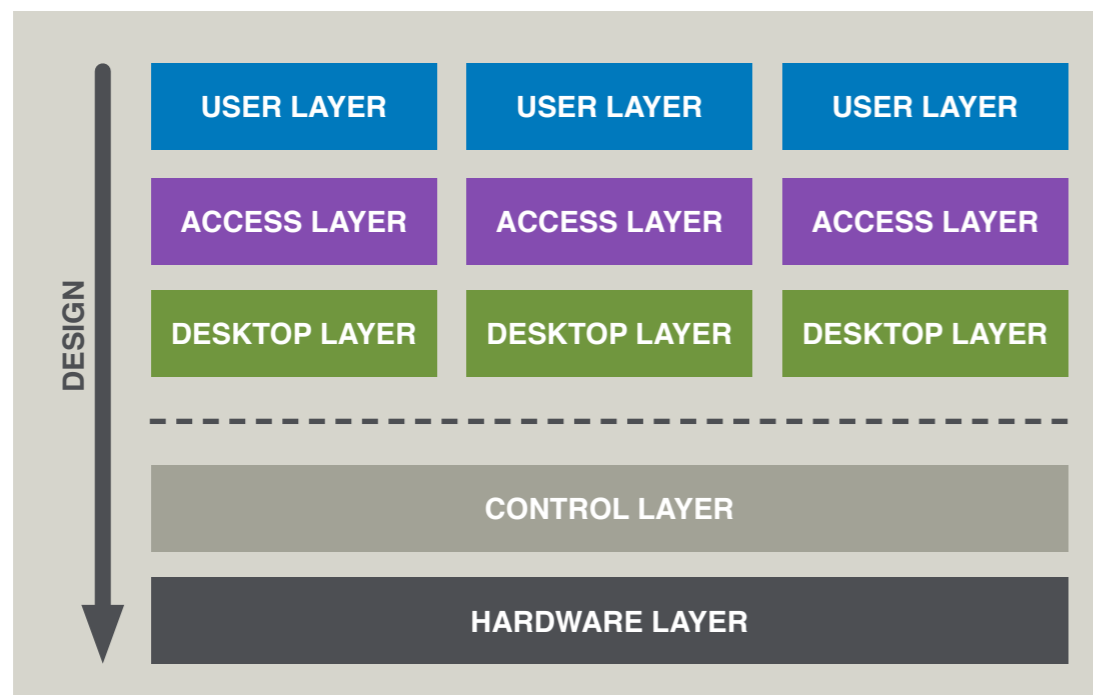
Design Quick Access Links

Hyper-V Storage	167
Virtual Machine Manager	169
VM Provisioning	171
High Availability	172
Monitoring	174
Backup and Recovery	175
Storage	176
Disaster Recovery	181

Design Overview

Designing a desktop virtualization solution is simply a matter of following a proven process and aligning technical decisions with organizational and user requirements. Without the standardized and proven process, architects tend to randomly jump from topic to topic, which leads to confusion and mistakes. The recommended approach focuses on working through five distinct layers:

Five-Layer Design Model



The top three layers are designed for each user group independently, which were identified during the [user segmentation section](#) of the assess phase. These layers define the user's virtual desktop and how users access their desktop. Upon completion of these three layers, the foundational layers (control and hardware) are designed for the entire solution.

This process guides the design thinking in that decisions made higher up impact lower level design decisions.

User Layer

The top layer of the design methodology is the user layer, which is defined for each unique user group.

The user layer appropriately sets the overall direction for each user group's virtualized environment. This layer incorporates the assessment criteria for business priorities and user group requirements in order to define effective strategies for endpoints and Citrix Receiver. These design decisions impact the flexibility and functionality for each user group.

Endpoint Selection

There are a variety of endpoints devices available, all with differing capabilities, including:

- Tablet based on Android or iOS
- Laptop
- Desktop PC
- Thin client
- Smartphone

The user's primary endpoint device must align with the [overall business drivers](#) as well as each user's role and associated requirements. In many circumstances, multiple endpoints may be suitable, each offering differing capabilities.

Decision: Endpoint Ownership

In most organizations, endpoint devices will be corporate owned and managed. However, more and more organizations are now introducing bring your own device (BYOD) programs to improve employee satisfaction, reduce costs and to simplify device management. Even if BYOD is a business priority, it does not mean

that every user should be allowed to use a personal device in the corporate environment.

Certain user requirements, which were identified during the [user segmentation](#), can greatly impact the suitability of personal devices:

- **Security** – Users requiring a high-level of security might not be able to bring a personal device into the secured environment for risk of data theft.
- **Mobility** – Users operating in a disconnected mode might not be able to use a personal device, as the local VM FlexCast model associated with this type of requirement can have specific hardware requirements, or special maintenance requirements. Additionally, the local operating system may be destroyed when the local VM FlexCast option is utilized.
- **Criticality** – Users with a high criticality rating might require redundant endpoints in the event of failure. This would require the user to have an alternative means for connecting in the event their personal device fails, likely making these users poor candidates for a BYOD program.
- **FlexCast model** – A personal device should not be recommended for user groups utilizing a client-hosted FlexCast model like streamed VHD, local VM or Remote PC. These FlexCast models typically require a specific hardware configuration or installation that will restrict device selection.

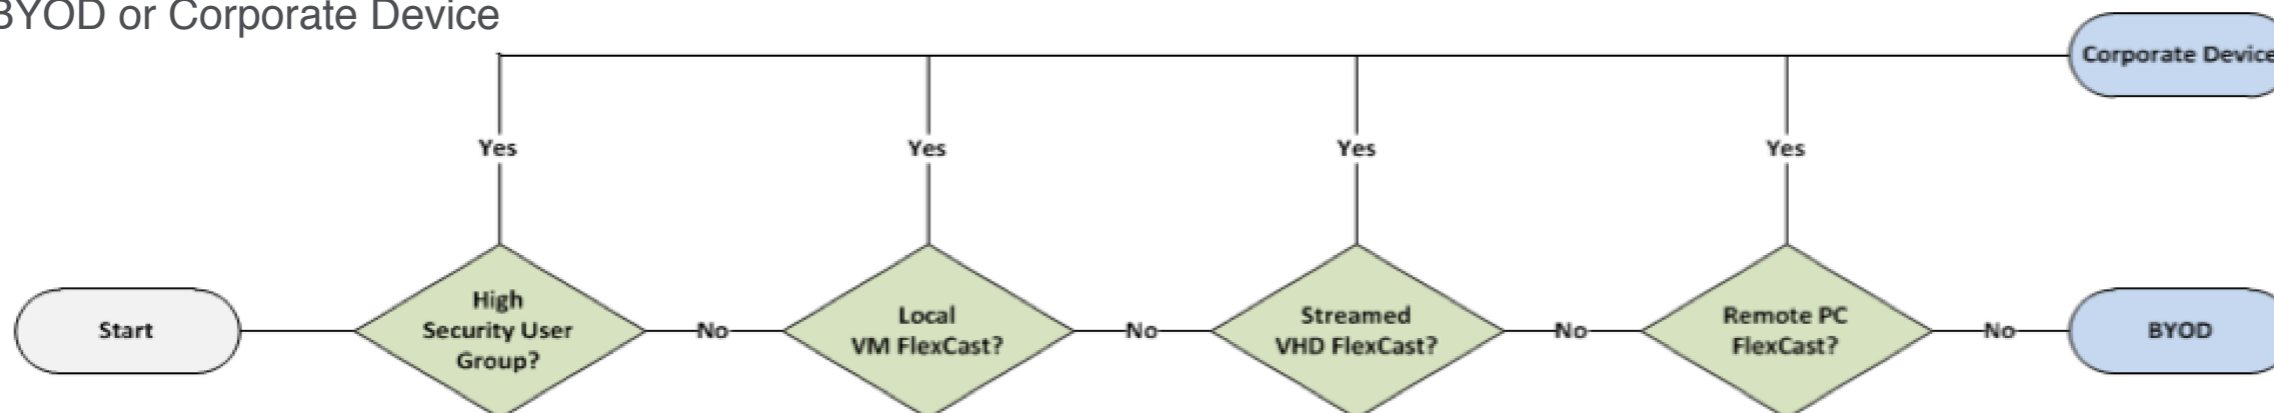
The diagram shown below provides guidance on when user owned devices could be used.

Decision: Endpoint Lifecycle

Organizations may choose to repurpose devices in order to extend refresh cycles or to provide overflow capacity for contract workers. Endpoints now offer more capabilities allowing them to have longer useful lifespans. In many cases, these hardware capabilities vastly outstrip the needs of a typical user. When coupled with the ability to virtualize application and desktop workloads, this provides new options to administrators such as repurposing existing workstations. These options go well beyond the simple three-year PC refresh cycle. However, the benefits of repurposing or reallocating a workstation should be balanced against the following considerations.

- **Minimum standards** – While cost factors of repurposing existing workstations may be compelling, certain minimum standards should be met to guarantee a good user experience. At a minimum, it is recommended that repurposed workstations have a 1GHz processor, 1GB of RAM, 16GB of free disk space and a GPU that is capable of supporting HDX features.
- **Business drivers** – Priorities underpin the success of any major project. Those organizations that have prioritized reducing

BYOD or Corporate Device



capital expenditure by means of prolonging the hardware refresh cycle can benefit from repurposing hardware. Conversely, if an organization's business drivers include reducing power consumption as part of an overall green initiative, purchasing newer endpoints may be beneficial in order to take advantage of the latest generation of power management capabilities available in the most modern devices.

- **Workload** – The type of work and FlexCast model for an end user can determine whether they are a good candidate for a repurposed endpoint, or may be better served with a new device. If the work performed by the individual involves locally installed applications, the individual may be best served by a new endpoint that offers the most powerful and recently updated processor and graphics architecture. However, if a user is largely performing tasks associated with virtualized applications that do not involve the latest multimedia capabilities such as webcams, VoIP and media redirection, then a repurposed workstation should be a viable alternative.

The following planning matrix outlines considerations when repurposing existing hardware:

Endpoint Procurement Criteria

Endpoint Provisioning Criteria	Repurpose Existing	Procure New
Capital restrained environment	•	
Heavy multimedia usage		•
High failure rate among existing desktops		•
High number of virtualized applications	•	
Low intensity workload	•	
Outmoded client-side feature set		•
Power consumption or green initiative(s)		•
Desire to prolong hardware refresh cycle	•	

Decision: Endpoint Form Factor

The capabilities of endpoints have grown along with efficiencies offered in thin client form factors. Even mid-range thin clients now

have graphics capabilities that allow utilization of HDX features such as multi-monitor support while offering management and power efficiency benefits. This expansion of capabilities has given IT administrators more options and flexibility than ever before.

Most organizations will likely deploy a mixture of fully featured clients as well as thin clients. It is important to focus on multiple user group attributed in order to best determine the type of endpoint that should be deployed:

Endpoint Selection

Endpoint	Mobility	FlexCast	Business Drivers	Security
Thin Clients	• Local • Roaming	• Hosted Shared • Hosted VDI • On Demand Apps	• Agility • Better desktop management • Increased security	High
Desktop PCs	• Local • Roaming	• Hosted Shared • Hosted VDI • Remote PC • Streamed VHD • On Demand Apps	• N/A – existing state	Medium
Desktop PC w/ XenClient	• Local • Roaming • Mobile	• Local VM	• Better desktop management • Increased security	High
Laptops	• Local • Roaming • Remote	• Hosted Shared • Hosted VDI • Remote PC • On Demand Apps	• BYOD • Work from anywhere	Low
Laptops w/ XenClient	• Local • Roaming • Remote	• Local VM	• Better desktop management • Increased security • Work from anywhere	High
Tablets	• Local • Roaming • Remote	• Hosted Shared • Hosted VDI • On Demand Apps	• BYOD • Work from anywhere	Low

Decision: Thin Client Selection

Thin client vendors now offer a range of operating system choices, including Windows Thin PC (based on Windows 7), embedded versions of Windows (XP, Windows 7 and Windows 8), Linux variants, as well as limited functionality clients that boot directly into a virtual desktop and offer a zero operating system footprint. The following factors should be considered during the selection of a

thin-client device:

- **User workload** – Windows Thin PC or limited functionality solutions such as Dell Wyse Zero clients should be tailored to true task workers or knowledge workers with limited needs. More capable thin devices such as Windows Embedded solutions can be provided to users that require more graphic capabilities and other intensive activities.
- **In-house expertise** – Many organizations have management toolsets already in place to handle thin client infrastructure such as retailers that may have kiosks. In-house expertise with a thin client management infrastructure can determine the selection of thin client platform. It is recommended that existing in-house expertise is leveraged, so long as the platform is capable of supporting a virtual desktop infrastructure implementation, as outlined on the [Citrix Ready site](#).
- **Licensing cost** – There are licensing considerations for most platforms. Windows thin PC and embedded versions incur immediate license costs related to Microsoft licensing, whereas a custom Linux solution may not. However, these costs must be closely balanced against additional add-on licensing that may be required for Linux devices, which are built into Windows. For example, various media codecs may require additional license expenditure in a Linux thin client context. For more information, please refer to the [Microsoft Partner Site](#).

Experience from the Field

Large systems integrator – A large systems integrator recommended that a customer deploy a single type of low-end, limited capability endpoint for all users. Upon deployment to production, users immediately complained that they received a poor user experience when viewing multimedia content over the WAN. At great cost, the systems integrator and customer re-assessed

the environment and chose to deploy endpoints that supported HDX MediaStream. The mistake caused a schism between systems integrator and the customer, resulting in lost time, capital and the end of a business relationship that was fostered over many years. It is critical that the endpoints assigned to each user group can support their requirements.

Receiver Selection

Citrix Receiver is an easy-to-install software client that provides access to applications, desktops and data easily and securely from any device, including smartphones, tablets, PCs and Macs.

The following section provides a series of design decisions that should be considered when deploying Citrix Receiver.

Decision: Receiver Type

While most organizations should simply deploy the latest Citrix Receiver compatible with their endpoint, it is important to recognize that there are certain differences between editions. The following table should be referenced to determine the most appropriate edition of Citrix Receiver for each user group. For the latest feature matrix, please refer to CTX104182 – [Receiver - Client Feature Matrix](#).

Decision: Initial Deployment

There are several deployment options available for delivering Citrix Receiver to an endpoint. Although it is usually a best practice to have a full version of Citrix Receiver deployed to an endpoint to provide the greatest level of functionality, it is important to consider fallback options such as the Java Client and the HTML5 Receiver for those situations where an installation of Citrix Receiver is simply not possible.

Experience from the Field

Furniture distributor – A furniture distributor maintains a configurator application for various furniture options. The configurator application is accessed via a limited functionality kiosk that is deployed at various furniture outlets, including small, independent retailers with little to no IT staff present. The kiosks are completely locked down in many situations, to the point where even the running of Java applications is limited. The kiosks do feature a modern browser (Google Chrome), and therefore, are able to utilize the HTML5 Receiver in order to provide access to the configurator application.

County government – A government IT organization provides services to all agencies operating in the county. A mixture of full desktops and iPads. Since the desktops are joined to the Active Directory domain, GPOs are utilized to deploy and configure Citrix Receiver. Mobile users accessing the Citrix environment via an iPad install and configure Receiver from the App Store. To allow for seamless provisioning, e-mail based discovery was configured. This allows users to configure Receiver for both internal and external access through NetScaler Gateway by entering in their e-mail address.

The following mechanisms are commonly used to deploy and update Citrix Receiver:

- **StoreFront** – If Citrix StoreFront is available, administrators can deploy Citrix Receiver via a Receiver for Web site. When deployed, a Receiver for Web site enables users to access StoreFront stores through a web page. If the Receiver for Web site detects that a user does not have a compatible version of Citrix Receiver, the user is prompted to download and install Citrix Receiver.
- **Internal download site** – Users may be prevented from downloading software from the Internet, even if they have permission to install applications. Administrators can create an

internal website for supported Citrix Receivers. Templates for such a site are available from the [Citrix Downloads site](#). When a user accesses a deployed site, it will detect the operating system on the user's device and guide the user to a local Citrix Receiver install suitable for their device.

- **Windows store** – The Citrix Receiver for Windows 8/RT is available from the Windows 8 store. This Receiver is available for ARM or Intel based Windows 8/RT devices only. It supports native Windows 8 style and gestures.
- **Mobile device** – Many mobile devices have unique methods of deploying applications. It simply may not be possible to deploy Citrix Receiver via traditional methods, therefore the following options are most likely to be selected for mobile devices:
 - **Markets and stores** – The latest supported Citrix Receiver is available for Android and iOS on the deployment portal of choice for these platforms. For Android devices version 2.2 and higher, this will be the Android Play platform. For iOS devices, this will be the Apple Store.
 - **Other mobile deployment methods** – Many mobile platforms offer proprietary methods of deploying software. For example, it is possible to utilize BlackBerry Enterprise Server to deploy the BlackBerry Citrix Receiver 2.2 to supported devices.
- **Master image** – Most organizations support a limited number of master desktop images, which are deployed to each business owned desktop, laptop, server, or virtual desktop. A common mechanism to ensure access to virtual desktops and applications is to include a supported version of Citrix Receiver in the master image. Subsequent updates to Citrix Receiver are handled either by utilizing the Citrix Receiver updater plug-in and Merchandising Server, enterprise software deployment tools, or manually.
- **Enterprise software deployment** – Many organizations employ an enterprise software deployment (ESD) tool. ESD tools can

be used to deploy Citrix Receiver to managed endpoints. ESD tools cannot be used to deploy Citrix Receiver to unmanaged endpoints, such as employee owned or partner devices.

- **Group policy** – Deploy and configure Citrix Receiver via Microsoft Group Policy. Sample start-up scripts that deploy and remove Citrix Receiver and Receiver Enterprise are available on Citrix XenApp and XenDesktop media:

Citrix Receiver and Plugins\Windows\Receiver\Startup_Logon_Scripts

- **Manual install** – All supported versions of Citrix Receiver are available from the [Citrix Receiver Download site](#). Upon landing on this site, client detection is performed and a platform and operating system specific link is provided to allow users to download an appropriate edition of Citrix Receiver. It is important to note that no configuration will be accomplished via this download, so email based discovery will need to be performed. This option is likely to be utilized in a BYOD environment.
- **Merchandising Server** – Citrix Merchandising Server can be used to deploy Citrix Receiver. In addition, Merchandising Server can be configured to update Citrix Receiver and all supported plug-ins, or a sub-set depending on the needs of an organization. To achieve this, Merchandising Server periodically communicates with the Citrix update service, which maintains a list of components and their updates. Merchandising Server is not recommended as a long-term solution because it will reach end of maintenance in August 2014 and end of life in August 2015. For more information, please refer to the Citrix support article – [Lifestyle Milestones for Citrix XenDesktop](#).

Selecting the appropriate deployment method is based on the type of Citrix Receiver selected. The table on the next page should be referenced to help identify the appropriate deployment options for Citrix Receiver.

Decision: Initial Configuration

Citrix Receiver must be configured in order to provide access to enterprise resources. The method of configuration varies by Citrix Receiver edition, the form factor of the device, and lastly the access method (local or remote) that is involved. Several methods may be viable for an organization. The method utilized is contingent on the resources (people, systems, time) available as well as larger organizational initiatives such as BYOD programs.

The following methods can be used to configure Citrix Receiver:

- **Email based discovery** – The latest releases of Citrix Receiver can be configured by entering an email address. Email based Discovery requires Citrix StoreFront as well as an SRV DNS record which points to the FQDN of the StoreFront server.

Note: Any DNS platform should support email-based discovery, however only Windows DNS has been explicitly tested.

For remote access, NetScaler Gateway must be utilized with the corresponding SRV record in DNS. A valid server certificate on the NetScaler Gateway appliance or StoreFront server must be present in order to enable email-based account discovery.

- **Group policy** – Microsoft Group Policy can be used to configure Citrix Receiver. Once startup scripts have been deployed for Citrix Receiver, edit the corresponding script and ensure there is a value for the following parameter: SERVER_LOCATION=Server_URL. The default value is blank. Provide the URL of the server running Citrix StoreFront. The URL must be in the format http://servername or https://servername.

Receiver Deployment Options

Deployment Option	Built into Base Image	Enterprise Software Deployment	Active Directory and Group Policy	Receiver for Web Site	Internal Downloads Site	Windows App Store	Mobile Market or App Store	Merchandising Server
Citrix Receiver for Windows	•	•	•	o	o			1
Citrix Receiver Enterprise for Windows	•	•	•	o	o			
Citrix Receiver for Mac				o				•
Citrix Receiver for Windows 8/RT	o	o		•		•		
Citrix Receiver for Linux	3	3		3	o			
Citrix Receiver for Android							•	
Citrix Receiver for iOS							•	
Citrix Receiver for BlackBerry		2					•	
HTML 5 Receiver				•				
Online Plug-in	•	•						•
ShareFile Plug-Ins	•							•
Offline Plug-In	•							•
Single Sign-On Plugin	•							•
NetScaler Gateway Plug-in	•							•
CloudBridge Acceleration Plug-in	•							•
Desktop Lock	•							

• - Recommended o - Available as an option

1 - Option for VDI hosts 2 - Proprietary to BlackBerry 3 - Possible, but not extensively tested

- **Provisioning file** – For environments running StoreFront, it is possible to provide users with a provisioning file that contains store information. Provisioning files are exported from the StoreFront console. The file is saved with a “*.cr” extension and can then be placed on a shared network resource, a Receiver for Web site, or other web based resource. The file can then be launched from an endpoint, which automatically configures Citrix Receiver to use the store.
- **Manually** – If allowed, it is usually possible to configure Citrix Receiver manually. This method should be reserved for administrators or users that have advanced knowledge.

Decision: Updates

Citrix Receiver and plug-ins are in active development. As such, periodic updates are released that provide enhanced functionality or address user issues. As with any actively developed product, the latest version of these products should be deployed to the endpoints so that users benefit from the latest functionality. There are multiple methods available to update Citrix Receiver and, if applicable, associated plug-ins.

- **Citrix.com update service** – Citrix Receiver can be configured to use Citrix.com as its source for updates. This configuration, which is set inside the StoreFront administration panel, allows receiver and plugin updates to be downloaded automatically from Citrix.com. Since the updates will be downloaded from an external host, this is inefficient approach when a large number of users at the same site request an update.
- **Merchandising Server** – For platforms that support the Citrix Receiver updater plug-in (Citrix Receiver, Citrix Receiver Enterprise, and Citrix Receiver Mac). The Citrix Receiver updater plug-in communicates with Merchandising Server and identifies package updates for deployment.
- **Enterprise software deployment** – ESD tools provide a viable

alternative to Merchandising Server with receiver updater. Additional thought must be given to updating unmanaged devices and endpoints outside of the corporate firewall, which is a solution that Merchandising Server can address.

- **Manual updates** – When no automated solution is available, manual methods can be used to update Citrix Receiver. Whether deployed on Receiver for Web site, StoreFront, an internal Citrix Receiver site, or an external site, these options will require user involvement in updating Citrix Receiver and the associated plug-ins. Due to the involved nature of manual updates coupled with the opportunity for a user mistake, this option should only be considered as a last resort.

Access Layer

The second layer of the design methodology is the access layer, which is defined for each user group.

Creating an appropriate design for the access layer is an important part of the desktop virtualization process. This layer handles user validation through authentication and orchestrates access to all components necessary to establish a secure virtual desktop connection.

The access layer design decisions are based on the mobility requirements of each user group as well as the endpoint devices used.

Decision: Authentication Point

Before a user connects to a virtual resource, they must first authenticate. The place of authentication is often determined by the user group’s mobility requirements, which were defined during the user segmentation process. There are two authentication points available in XenDesktop 7:

- **StoreFront** – Citrix StoreFront provides authentication and

resource delivery services for Citrix Receiver, enabling centralized enterprise stores to deliver desktops, applications and other resources.

- **NetScaler Gateway** – NetScaler Gateway is an appliance providing secure application access and granular application-level policy controls to applications and data while allowing users to work from anywhere.

The following table lists preferred authentication points according to user group mobility requirements:

Preferred Authentication Point

User Group's Mobility Requirement	Preferred Authentication Point
Local	StoreFront
Roaming Local	StoreFront
Remote	NetScaler Gateway
Mobile	NetScaler Gateway

Authentication for user groups with a mobility requirement of remote or mobile may occur directly on StoreFront where required. For example, DMZ security policies may prohibit access from the NetScaler Gateway to the domain, which is required to support SmartCard client certificate authentication. Access to StoreFront for authentication may then be delivered via a NetScaler SSL_BRIDGE virtual server, which provides a conduit for https traffic. Typically, the virtual server would be hosted alongside a NetScaler Gateway on the same NetScaler configured to provide HDX Proxy access to the virtual desktop environment. Although such a use case may sometimes be necessary, the recommended best practice is to authenticate external users via NetScaler Gateway.

Decision: Authentication Policy

Once the authentication point has been identified, the type of authentication must be determined. The following options are the primary methods available:

- **StoreFront** – Supports a number of different authentication methods, although not all are recommended depending on the user access method, security requirements and network location.
 - **User name and password** – Requires users to logon directly to the site by entering a user name and password.
 - **Domain pass-through** – Allows pass-through of domain credentials from users' devices. Users authenticate to their domain-joined Windows computers and are automatically logged on when they access their stores.
 - **NetScaler Gateway pass-through** – Allows pass-through authentication from NetScaler Gateway. Users authenticate to NetScaler Gateway and are automatically logged on when they access their stores.
 - **Smart card** – Allows users to authenticate using smart cards and PINs through Citrix Receiver for Windows and NetScaler Gateway. To enable smart card authentication, user accounts must be configured either within the Microsoft Active Directory domain containing the StoreFront servers or within a domain that has a direct two-way trust relationship with the StoreFront server domain. Multi-forest deployments involving one-way trust or trust relationships of different types are not supported.
 - **Unauthenticated access** – Allow users to access applications and desktops without presenting credentials to StoreFront or Citrix Receiver. Local anonymous accounts are created on demand on the Server VDA when sessions are launched. This requires a StoreFront store configured for unauthenticated access, a Server OS based VDA, and a XenApp 7.6 Delivery Group configured for unauthenticated users.
- **NetScaler Gateway** – The NetScaler Gateway supports several authentication methods. The list below includes those primarily used in virtual desktop environments. Each may be used individually, but are often combined to provide multi-factor

authentication.

- **LDAP** – The lightweight directory access protocol (LDAP) is used to access directory information services such as Microsoft Active Directory. NetScaler Gateway uses LDAP to authenticate users and extract their group membership information.
- **RADIUS (token)** – Remote authentication dial in user service (RADIUS) is a UDP based network security protocol that provides authentication, authorization and accounting. A network access server (NetScaler Gateway in this case) forwards credentials to a RADIUS server that can either check the credentials locally, or check them against a directory service. The RADIUS server could then accept the connection, reject the connection, or challenge and request a second form of authentication such as a token.
- **Client certificate** – Users logging on to a NetScaler Gateway virtual server can also be authenticated based on the attributes of a client certificate presented to the virtual server. Client certificates are usually disseminated to users in the form of smartcards or common access cards (CACs) that are read by a reader attached to each user's device.

The authentication type for a user group is often determined based on security requirements as well as the authentication point used. The table shown below helps define the appropriate solution for each user group based on the level of security required:

Experience from the Field

Retail – A small private retail company provides virtual desktop

users with access to non-sensitive data such as marketing catalogs and email. They are not required to adhere to security regulations such as Sarbanes Oxley. Therefore, LDAP authentication has been implemented based on user name and password.

Financial – A medium financial enterprise provides their virtual desktop users with access to confidential data such as banking transaction records. They are governed by security regulations such as the Statement on Accounting Standards (SAS) 70 and are required to utilize multi-factor authentication for remote access users. LDAP authentication has been implemented based on user name and password along with RADIUS authentication using tokens.

Government – A large federal institution provides virtual desktop users with access to highly confidential data such as private citizens' personal records. They are subject to regulation by Department of Defense (DOD) security standards. LDAP authentication has been implemented based on user name and password, along with Client Certificate authentication using CAC cards.

Decision: StoreFront or Web Interface

Web Interface and StoreFront are two different solutions, whose feature sets overlap in many areas, but also offer a variety of distinct features. Therefore it is very important for organizations to review the capabilities of each product against their requirements. In general, it is strongly recommended to build new solutions based on StoreFront. Since new features will not be added to Web Interface, this chapter will focus on StoreFront only.

Note: Web Interface 5.4 support has been extended for XenDesktop

Authentication Policy Guidance

Initial Authentication Point	Mobility	LDAP User Name and Password	Pass-through	LDAP + Token	LDAP + Smartcard	Token + Smartcard
StoreFront	Local & Roaming Local	Medium	Medium	High	High	High
NetScaler Gateway	Remote	Low	N/A	High	High	High

7.6 & XenApp 7.6. For more information please view the Citrix Product Matrix.

While StoreFront goes beyond Web Interface in many areas, StoreFront 2.5 and 2.6 does not support all features of Web Interface.

Note: With StoreFront 2.0 and higher, it is no longer necessary to store user subscription data in Microsoft SQL database.

The following table outlines the Web Interface features that are not available in StoreFront 2.6:

Web Interface Features not Supported by StoreFront 2.6

Area	Feature
Deployment options	NetScaler Integration – StoreFront is deployable as an application
	NetScaler Nav UI-access to web link and documents
	Simple Multi-Tenant Administrator Admin
	Simple Data Back up
	Multiple Servers on same IIS Server
	Deployment on Work Group (StoreFront must be deployed on a domain)
User Experience	Users settings in Web Interface
	Client proxy settings configuration
	Offline Apps (Users cannot access offline applications or App-V sequences through Receiver for Web sites. Native Receiver is required)
	Compact/Low graphics Mode and embedding
Authentication	Authentication via XML Services (StoreFront does direct AD auth)
	Active Directory Federation Services (ADFS) 1.0 integration
	Account self-service (SSPR) – reset/unlock with security questions
	Explicit NDS password
	Deep portal integration (SharePoint, customer portal)
Other Features	Settings per location (IP subnet)
	Client proxy settings configuration
	Offline Apps (users cannot access offline applications or App-V sequences through Citrix Receiver for Web Sites. Native Citrix Receiver is required)
	Compact/low graphics mode and embedding
	Editable error messages for both web and PNA access
	Web Interface for SharePoint

StoreFront

Citrix StoreFront, the successor to Citrix Web Interface, authenticates users to XenDesktop, XenApp, and App Controller (SaaS Apps) resources. StoreFront enumerates and aggregates available desktops and applications into stores that users access through Citrix Receiver for Windows, iOS, Android, Win8/RT or Receiver for Web sites. StoreFront is an integral component of XenDesktop 7.x and can be used with XenApp 5.0/XenDesktop 5.5 and higher deployments. StoreFront is essential for managing multi-site XenDesktop deployments. For more information on StoreFront, see the Citrix eDocs – [About StoreFront](#).

Decision: Unauthenticated Access

Unauthenticated access allows users to access XenApp published desktops and applications via Citrix StoreFront without having to provide Active Directory domain credentials. Unauthenticated access offers a fast logon experience and is generally used with public or kiosk workstations, or applications with built-in user management features.

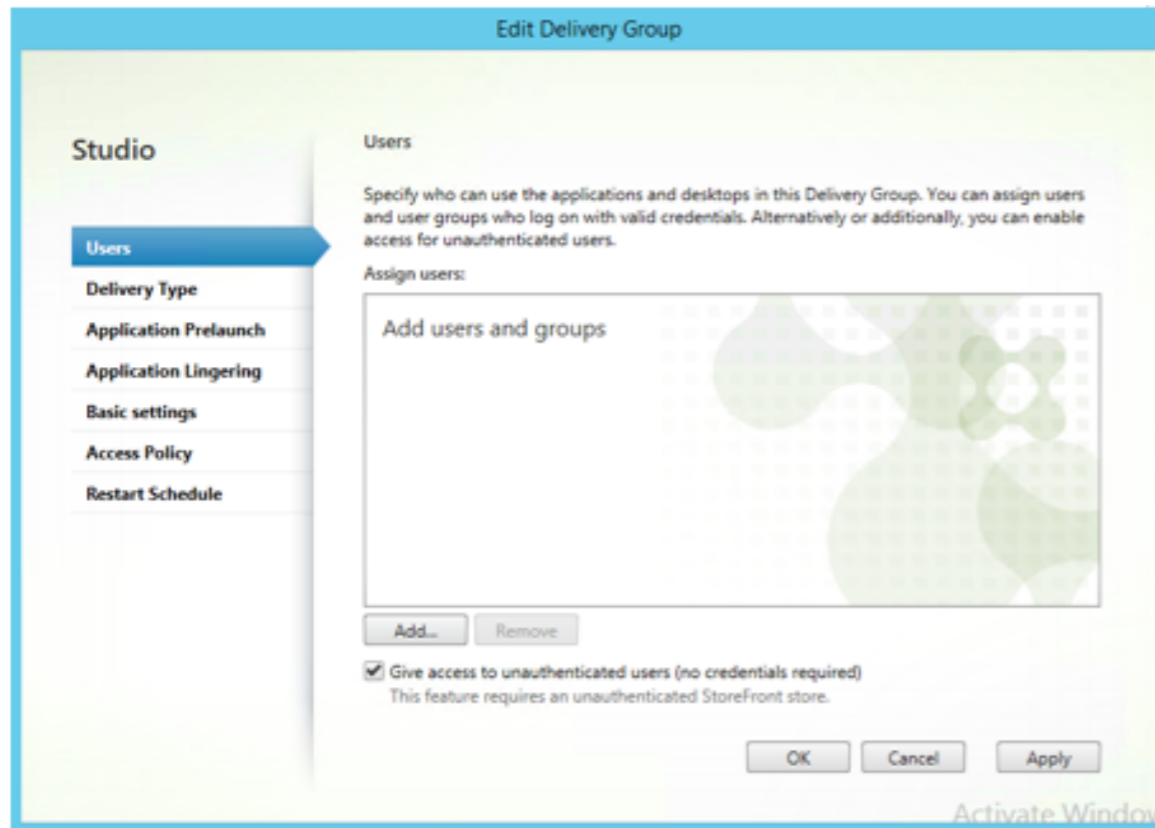
Building a XenApp environment with unauthenticated access requires the following components:

- XenApp 7.6 Delivery Controller
- StoreFront 2.6 store that has been configured for unauthenticated users
- Virtual Delivery Agent running on Windows Server 2008 R2 or higher
- A client with Citrix Receiver installed

When a XenApp unauthenticated access session is launched, a local user account becomes associated with the session. When the session logs off, the local user account is returned to the pool to be used by another connection. The local accounts are typically named AnonXYZ, where XYZ is a unique 3-digit value.

Unauthenticated access is enabled in Citrix Studio when specifying the users or user groups allowed access to applications and/or desktops in the Delivery Group.

Unauthenticated Access



The server VDA will create the anonymous local accounts on demand up to the maximum specified in the registry or 99 if no maximum is provided. This number can be changed by editing the value for the following registry key:

`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Citrix\MaxAnonymousUsers`

The Anonymous Users' profiles are reset after each session ends. For considerations and requirements for unauthenticated access please refer to Citrix eDocs – [Manage users in a Delivery Group](#).

Experience from the Field

A hospital is using XenApp to deliver their EMR application to users. ThinClient devices on stationary and mobile carts are being used by doctors and nurses to capture and retrieve patient data. Unauthenticated access has been configured to prevent medical staff from having to authenticate to the domain as well as the EMR application.

Decision: High Availability

If the server hosting StoreFront or the respective web service is unavailable, users will not be able to launch new virtual desktops, published applications or manage their subscriptions. Therefore at least two StoreFront servers should be deployed to prevent this component from becoming a single point of failure. By implementing a load balancing solution, users will not experience an interruption in their service. Options include:

- **Hardware load balancing** – An intelligent appliance, which is capable of verifying the availability of the StoreFront service and actively load balance user requests appropriately. Citrix NetScaler is a great example of a hardware load balancer. Citrix NetScaler is an ideal load balancer, coming pre-configured with StoreFront health checks.
- **DNS Round Robin** – Provides rudimentary load balancing across multiple servers without performing any checks on availability. If a StoreFront server becomes unavailable, DNS round robin would still route users to the failed server. Because of this, DNS round robin is not recommended by Citrix.
- **Windows network load balancing** – A Windows service capable of performing rudimentary checks to verify the server is available but cannot determine the status of individual services. This can cause users to be forwarded to StoreFront servers which are not able to process new requests. The user would then not be able to launch applications in their session.

The following figure (on the next page) shows a typical StoreFront deployment using Citrix NetScaler, operating as a load balancer for the environment. External users authenticate and gain access to StoreFront with the help of a NetScaler Gateway. NetScaler will also authenticate internal users as well.

Decision: Delivery Controller High Availability and StoreFront

To provide users with desktops and applications, StoreFront must be configured with the IP address or DNS name of at least one Controller in each XenDesktop and XenApp site. For fault tolerance, multiple controllers should be entered for each site and/or farm specified. StoreFront will automatically failover to the second server in the list if the first server becomes unavailable (active/passive). For large deployments or environments with a high logon load an active distribution of the user load (active/active) is recommended. This can be achieved by means of a load balancer with built-in XML monitors and session persistency, such as Citrix NetScaler.

Decision: Security - Inbound Traffic

Communications from the web browser or Receiver and StoreFront server include user credentials, resource sets and session initialization files. Remote traffic is routed over networks outside the datacenter boundaries or on completely untrusted connections (such as the Internet). Therefore Citrix strongly recommends that this traffic is encrypted using SSL encryption. Remote traffic can be proxied via NetScaler Gateway in order to provide a secure connection.

Internal StoreFront servers may use a public or private (domain) certificate. Certificates must be installed on the StoreFront server and the NetScaler. If StoreFront is load balanced by a NetScaler and is using SSL encryption, the following table explains where the certificates must be installed.

Certificate Locations

Certificate	Type of certificate	Location of certificate
Public Certificate	Public	StoreFront server
	Intermediate	NetScaler
Private (Domain) Certificate	Private (Domain)	StoreFront server
	Root	NetScaler

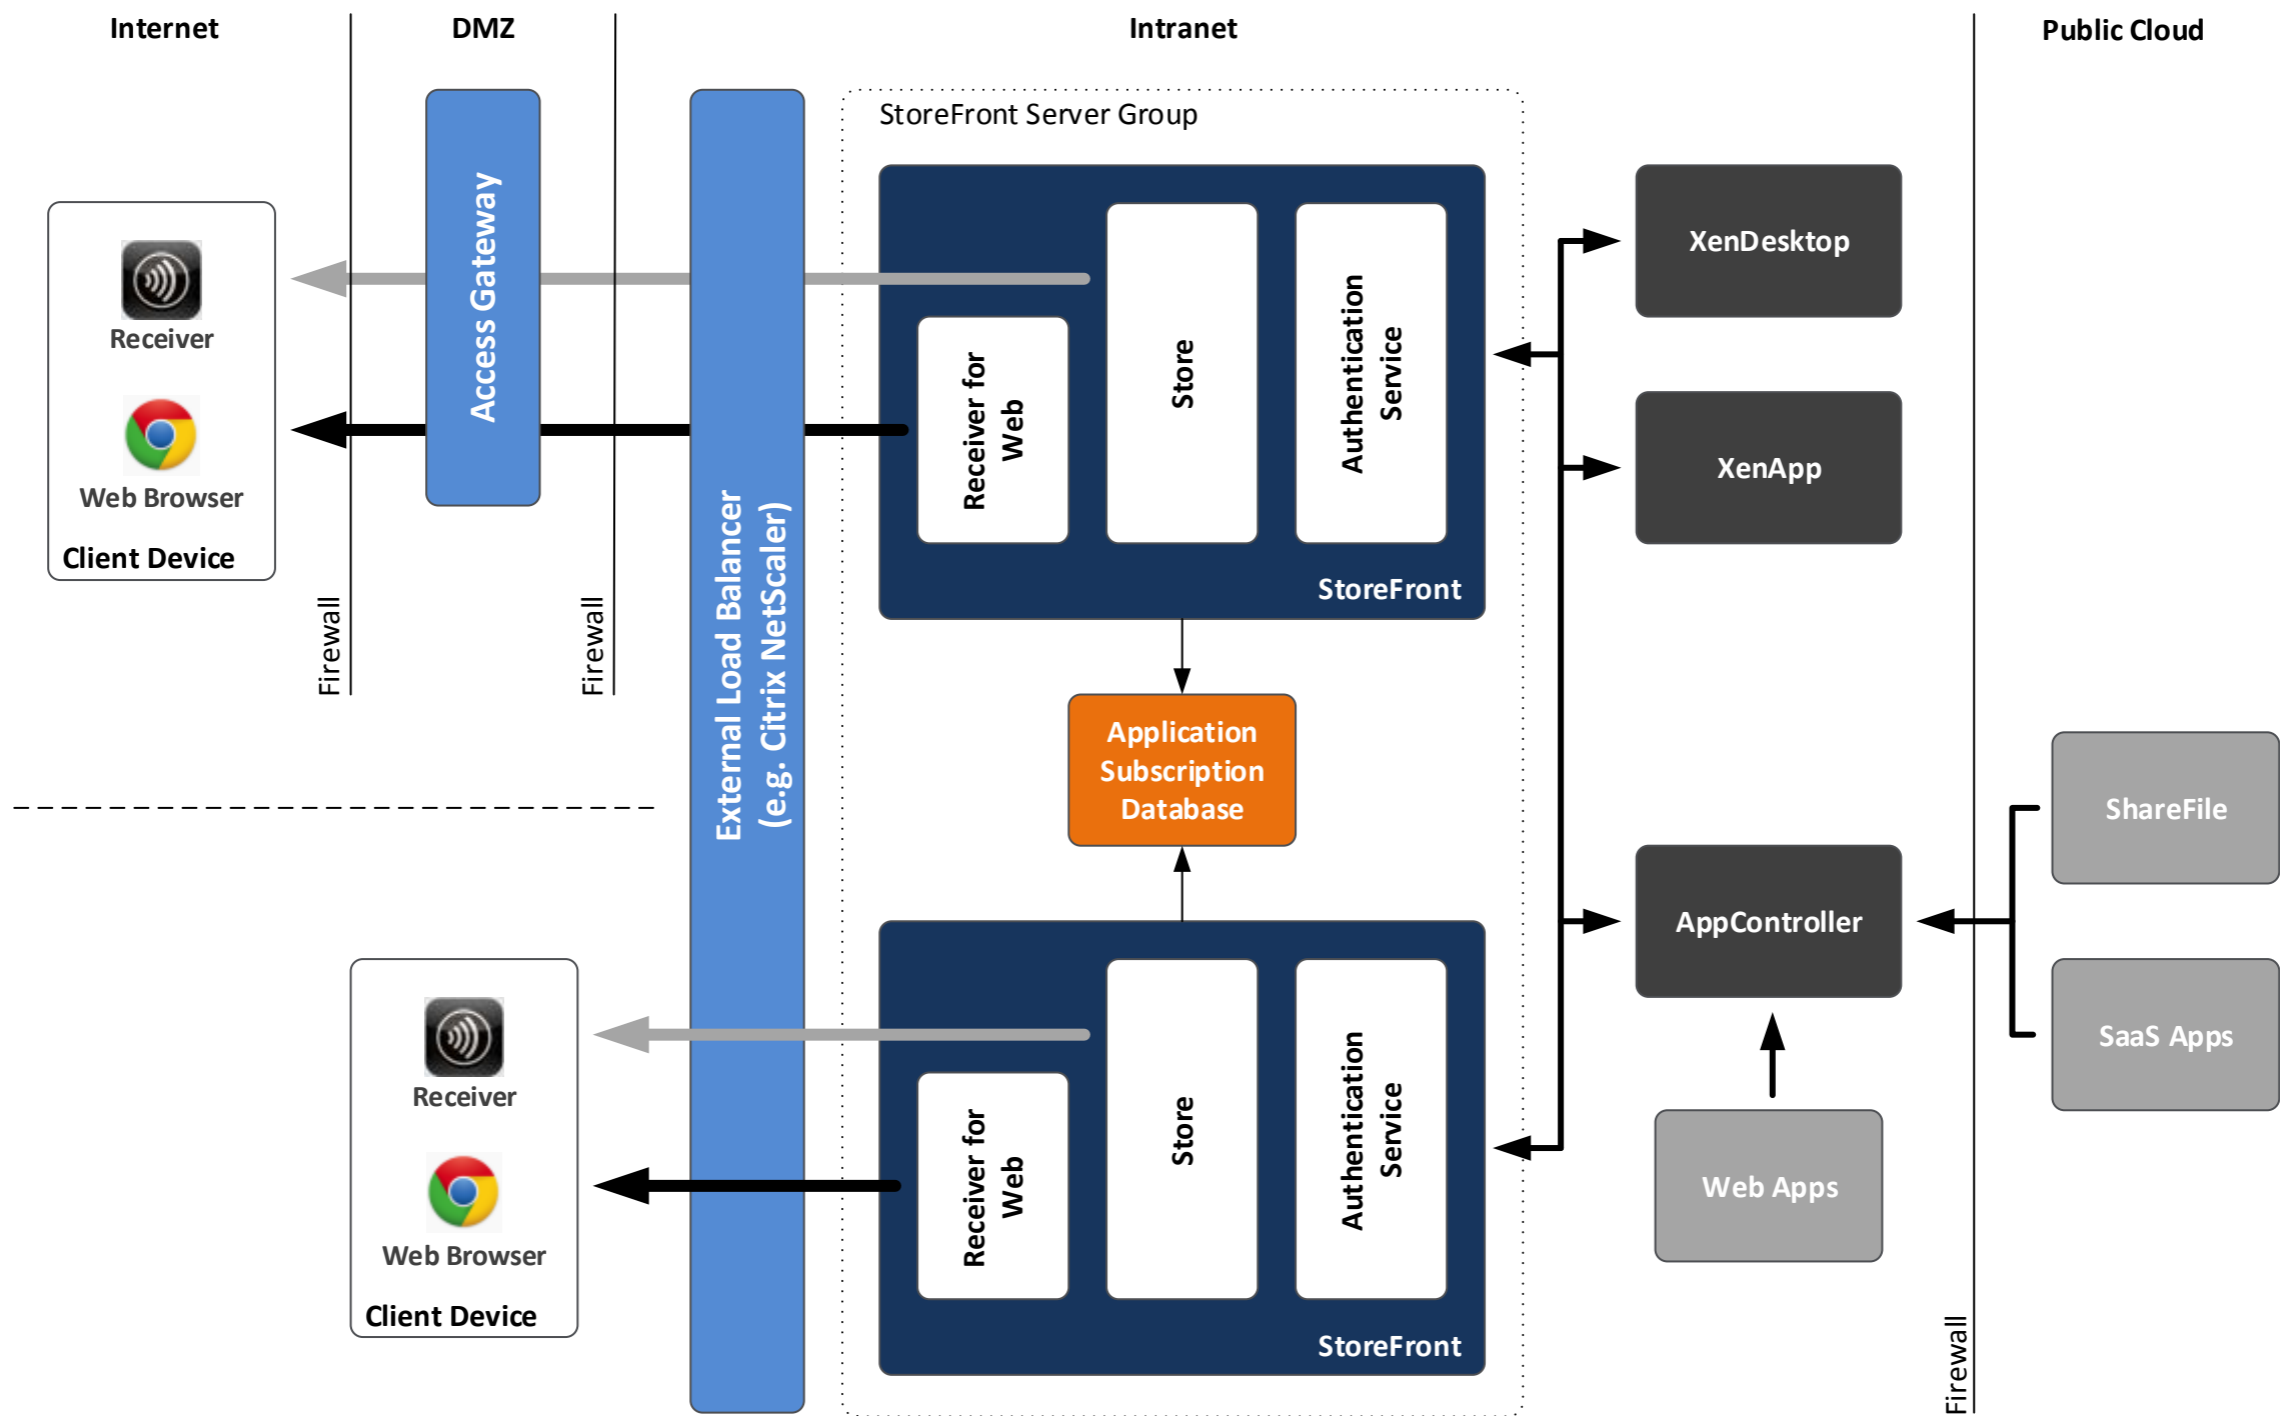
When using a mobile device (phones and tablets), that are not joined to the domain, a root certificate needs to be installed on each device. Alternatively, it may be easier to use a public certificate on an internal StoreFront to allow such devices on an internal network to easily connect.

Note: By default, Citrix Receiver requires SSL encryption to connect to StoreFront. This means email-based account discovery or the manual configuration of a StoreFront store in Citrix Receiver will fail unless a valid and trusted SSL certificate has been imported on to the StoreFront server and external load balancer (if applicable).

Decision: Security – Backend Traffic

User credentials are sent between StoreFront and the XenApp Controllers, XenDesktop Controllers and the App Controller virtual appliance. For example, in a typical session with a XenDesktop Controller, the StoreFront server passes user credentials to the Citrix XML Service for user authentication and the Citrix XML Service returns resource set information. By default, a HTTP connection is used pass the information between the StoreFront server and the XML service. The XML data transferred is sent in clear text, with the exception of passwords, which are transmitted using obfuscation. For environments with high security standards, it is recommended to encrypt the traffic between the StoreFront servers and the XML service by enabling SSL. Since the backend traffic is contained between infrastructure controllers and does not interact with user devices, it is easy to encrypt this traffic with a certificate from a private certificate authority.

Typical StoreFront Architecture



For more information, please refer to Citrix eDocs:

- [Use the SSL Relay](#) (XenApp 6.5 & Below Only)
- [Use SSL on XenDesktop & XenApp 7.5 Controllers](#)
- [Secure your StoreFront environment](#)

Decision: Routing Receiver with Beacons

Citrix Receiver uses beacon points (websites) to identify whether a user is connected to an internal or external network. Internal users are connected directly to resources while external users are connected via Citrix NetScaler Gateway. It is possible to control what a user sees by restricting applications due to which beacon they have access to.

The internal beacon should not be a site that is resolvable externally. By default, the internal beacon is the StoreFront service URL. The external beacon can be any external site that will produce an http response. Citrix Receiver continuously monitors the status of network connections (for example, link up, link down or change of the default gateway). When a status change is detected, Citrix Receiver will first verify that the internal beacon points can be accessed before moving on to check the accessibility of external beacon points. StoreFront provides Citrix Receiver with the http(s) addresses of the beacon points during the initial connection process and provides updates as necessary.

It is necessary to specify at least two highly available external beacons that can be resolved from public networks.

Decision: Resource Presentation

StoreFront displays resources differently than Web Interface. Instead of having all accessible resources appear on the home screen, first time users are invited to choose (subscribe) to the resources they want to regularly use after they logon. Before a user can launch an application, they must first choose which resources should be placed on their home screen. This approach, deemed

“Self-Service”, allows users to restrict the resources that they see on their home screen to the ones that they use on a regular basis. The resources chosen by every user for each store are recorded by the subscription store service so that they can be displayed on the Citrix Receiver home screen from any device that the user connects from.

Administrators should determine which applications should always be displayed to users on their home screen or the featured tab. These applications will vary for each deployment and should be defined during the Assess Phase of a Citrix Assessment. In general, these applications are common applications such as Microsoft Word and any other applications that every user in an environment may need. StoreFront can filter/present these resources using Keywords.

The following table explores the Keyword options:

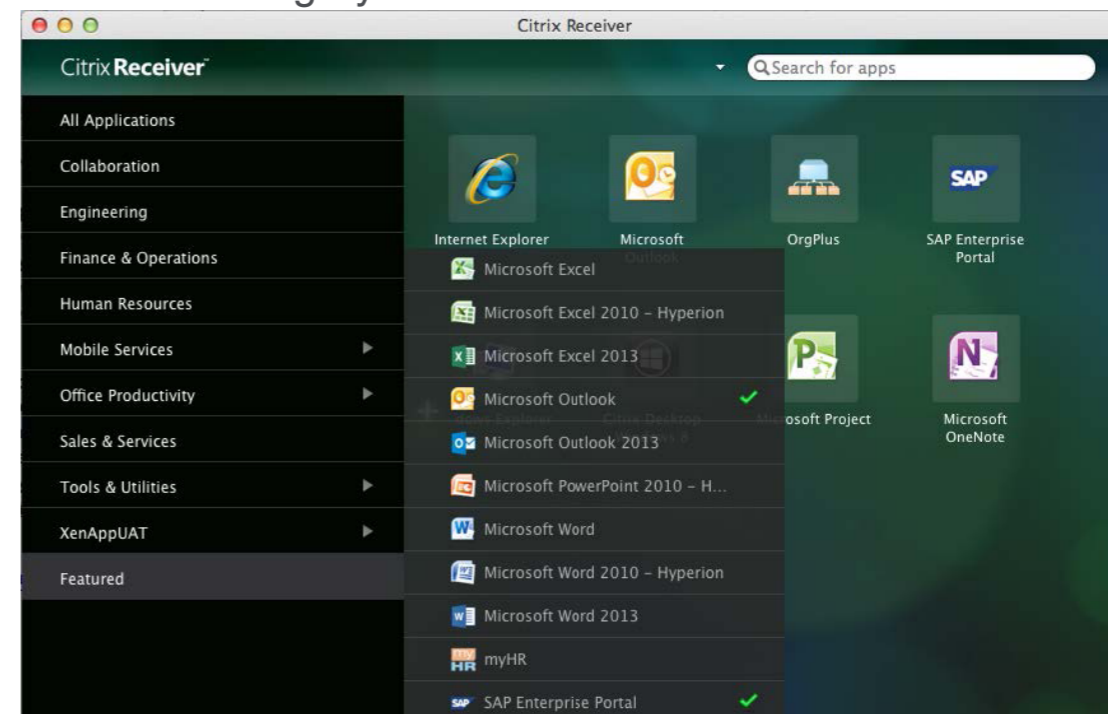
Keywords for Application Delivery

Area	Feature
Auto	Automatically subscribes all users of a store to an application. When users log on to the store, the application is automatically provisioned without users needing to manually subscribe to the application.
Mandatory	New in StoreFront 2.5, the Mandatory keyword will make applications automatically be subscribed to users of the store. However, users will not have the option to remove the application. This setting is useful when creating a core set of applications which must always be presented to all users.
Featured	Advertise applications to users or make commonly used applications easier to find by listing them in the Receiver Featured list.
Prefer	Specify a locally installed application should be used instead of an application available in Receiver. This feature is referred to as Local App Access. Before installing an application on a user's computer, Receiver searches for the specified patterns to determine if the application is installed locally. If it is, Receiver subscribes the application and does not create a shortcut. When the user starts the application from the Receiver window, Receiver starts the locally installed (preferred) application. If a user uninstalls a preferred application outside of Receiver, the application is unsubscribed during the next Receiver refresh. If a user uninstalls a preferred application from the Receiver window, Receiver unsubscribes the application but does not uninstall it.
Treat As App	By default, XenDesktop and XenApp hosted shared desktops are treated like other desktops by Receiver for Web sites. By using the keyword "TreatAsApp", the desktop will be displayed in the application views of Receiver for Websites rather than the desktop views. Users are required to subscribe before they can access the desktop.
Primary	When in a multi-site deployment, using this keyword ensures that an application is delivered from a designated site, no matter where the user is located. If an application is available from both sites, with the same name, the application from the secondary site will only be displayed if the application is not available from the primary site.
Secondary	A same property as the "Primary" keyword, except it designates an application in the secondary site.

Using Keywords is a very simple way to prepopulate a user's home screen. To add applications to the home screen, add KEYWORDS:Auto to the application or desktop description in XenApp or XenDesktop. Another option that can be used to

organize resources is using the keyword KEYWORDS:Featured. Unlike the Auto keyword, which places certain resources on the home screen, the Featured keyword only places resources in the Featured category (as shown below):

Featured Category in Citrix Receiver



The resource will also appear in another category if a Client Application folder has been specified.

In addition the string KEYWORDS:prefer="application" can be used to specify that the locally installed version of an application should be used in preference to the equivalent delivered instance if both are available.

For more information please refer to Citrix eDocs – [Optimize the user experience](#) and [Configure application delivery](#).

Decision: Scalability

The number of Citrix Receiver users supported by a single StoreFront server depends on the resources assigned and level of user activity:

StoreFront Scalability

StoreFront deployment	CPU usage	Simultaneous activities
<ul style="list-style-type: none"> Standalone deployment 4 CPUs 4 GB RAM Heavy Usage* 	75%	<ul style="list-style-type: none"> 291 per second
	90%	<ul style="list-style-type: none"> 375 per second
<ul style="list-style-type: none"> Cluster StoreFront deployment 2 Nodes each with: <ul style="list-style-type: none"> 4 CPUs 4 GB RAM Heavy Usage* 	75%	<ul style="list-style-type: none"> 529 per second
	90%	<ul style="list-style-type: none"> 681 per second

For an optimal user experience, Citrix recommends that no more than ten XenDesktop, XenApp, App Controller and VDI-in-a-Box deployments are aggregated into a single store.

Synchronization Database – If users connect to multiple StoreFront servers within an environment, their personalized settings (application subscriptions) are immediately stored on the StoreFront server and replicated to other StoreFront servers. The Synchronization Database on each StoreFront server needs to be large enough to accommodate user and application subscriptions, about 3 KB per user per app.

Decision: Multi-site App Synchronization

Moving between multiple locations would benefit from synchronization of their application subscriptions between multiple deployments. For example, a user based in Location 1 can log on to the StoreFront deployment in Location 1, access the store, and subscribe to some applications. If same user would travel to Location 2 and accessed a similar store provided by the Location 2 StoreFront deployment, the user would have to re-subscribe to all of their applications once again.

Each StoreFront deployment maintains details of users' application subscriptions separately by default. By configuring subscription synchronization between the stores, it is possible to ensure that

users only need to subscribe to applications in one location. The interval for the subscription synchronization will vary from site to site due to the distance between sites. The recommended interval should be less than the time needed for a user to travel from one site to another. For more information on how to use PowerShell with StoreFront 2.5, please refer to the Citrix eDocs article – [To configure subscription synchronization](#).

NetScaler Gateway

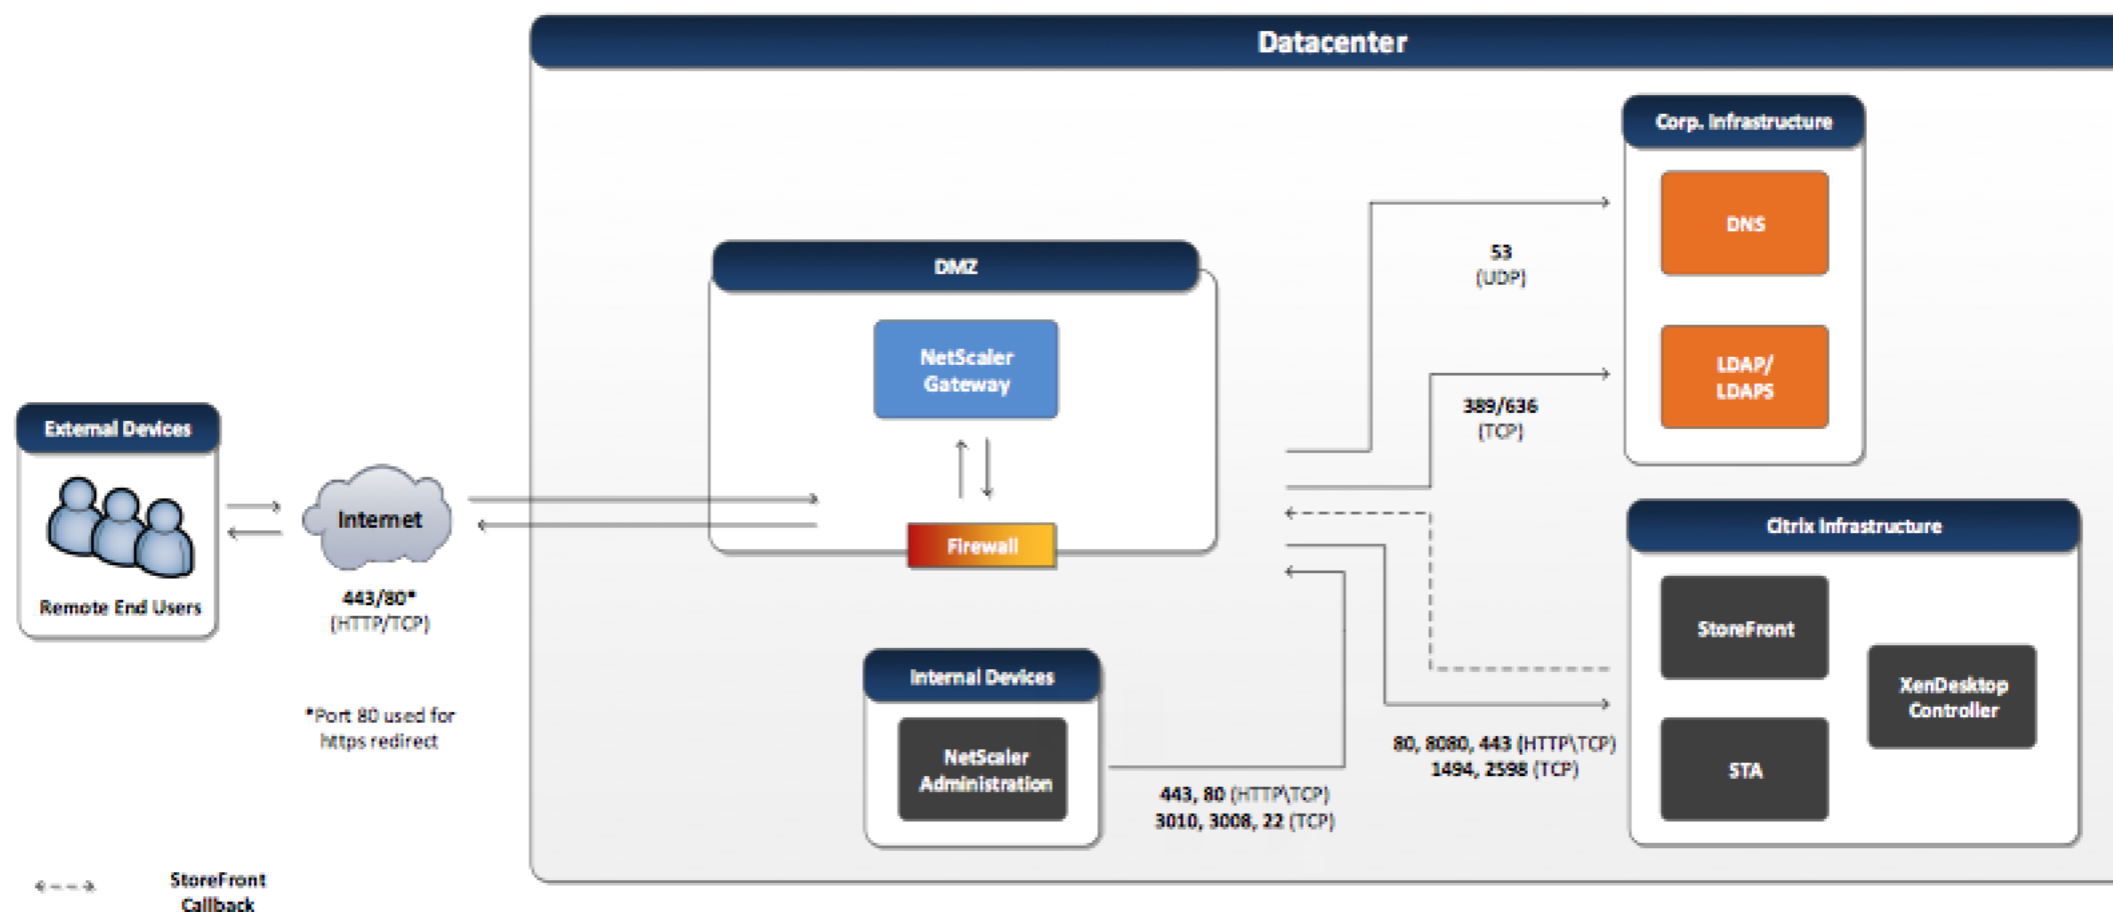
User groups utilizing NetScaler Gateway as their authentication point have additional access layer design decisions that must be considered. These design decisions are not applicable for non-NetScaler Gateway authentication points.

Decision: Topology

Selection of the network topology is central to planning the remote access architecture to ensure that it can support the necessary functionality, performance and security. The design of the remote access architecture should be completed in collaboration with the security team to ensure adherence to corporate security requirements. There are two primary topologies to consider, each of which provides increasing levels of security:

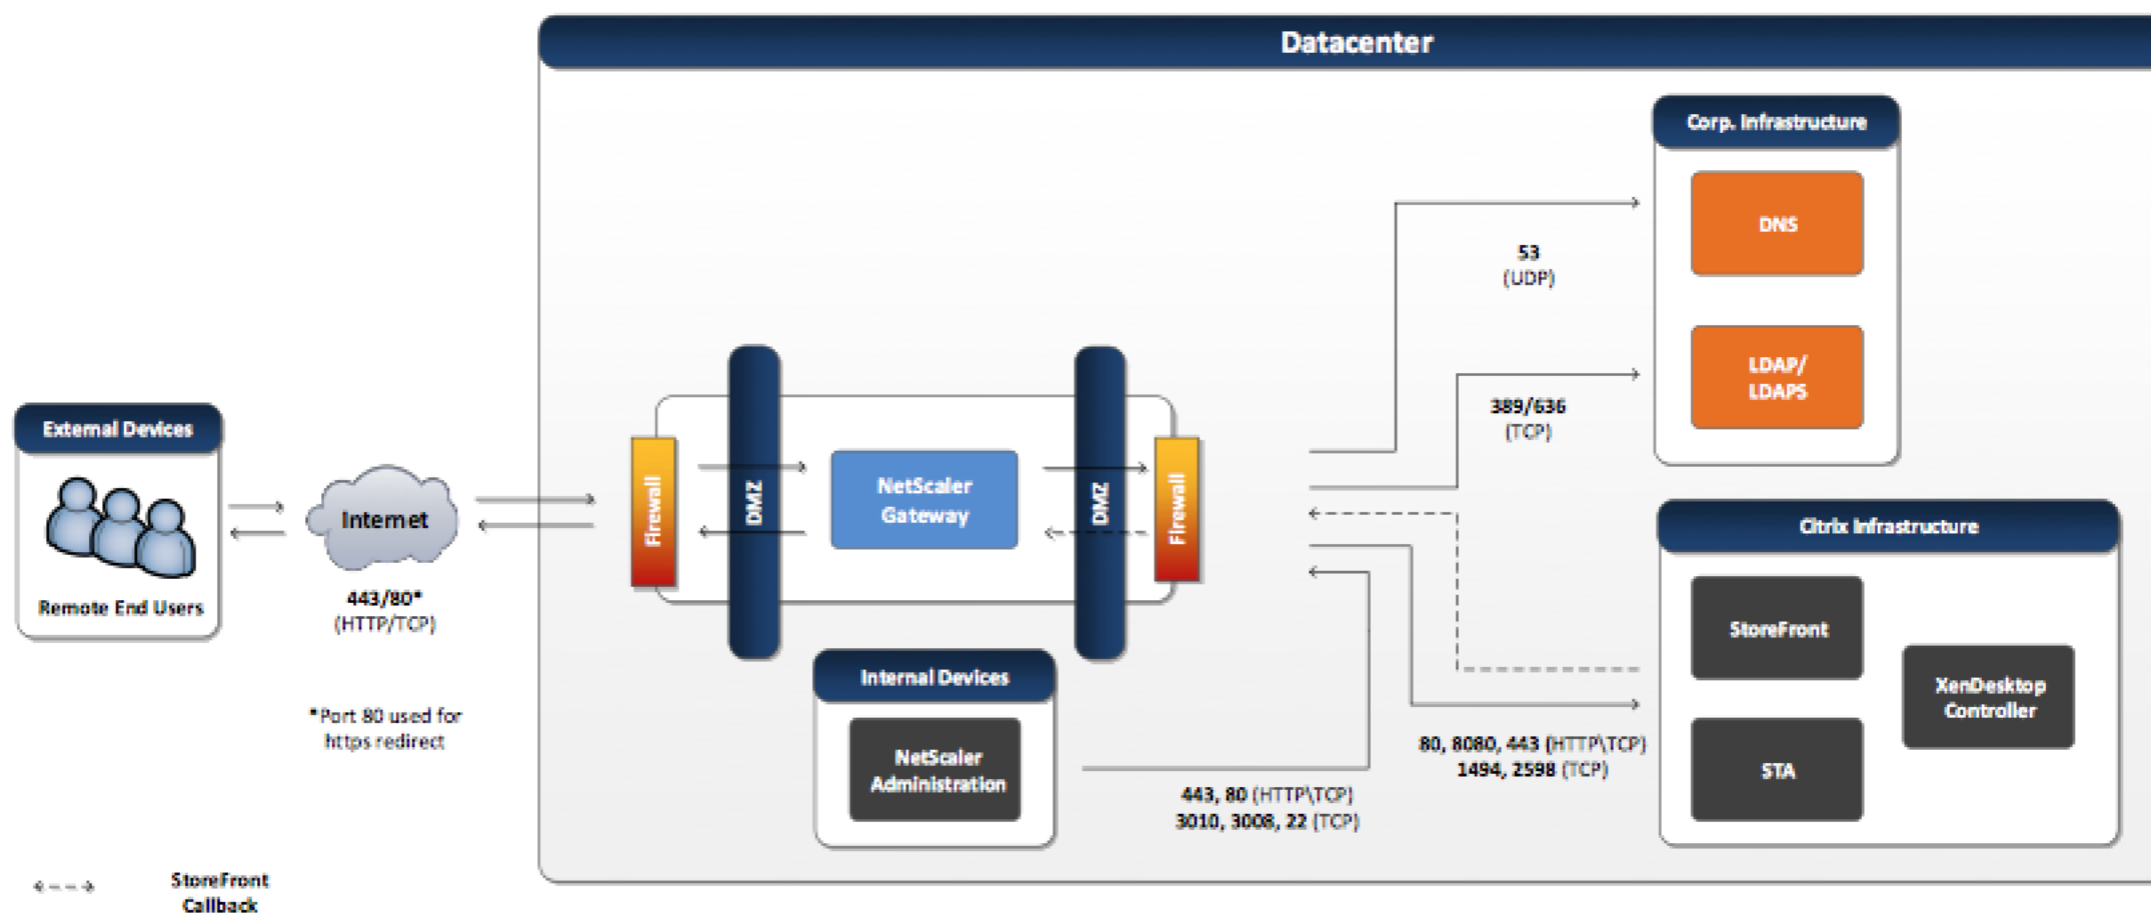
- **1-Arm (normal security)** - With a 1-arm topology, the NetScaler utilizes one physical or logical bonded interface, with associated VLAN and IP subnet, to transport both fronted traffic for users and backend traffic for the virtual desktop infrastructure servers and services.

1-Arm Topology



- **2-Arm (high security)** – With a 2-arm topology, the NetScaler Gateway utilizes two or more physically or logically bonded Transport of the frontend traffic for users is directed to one of these interfaces. The frontend traffic is isolated from backend traffic, between the virtual desktop infrastructure servers and services, which is directed to a second interface. This allows the use of separate demilitarized zones (DMZs) to isolate frontend and backend traffic flows along with granular firewall control and monitoring.

2-Arm Topology



Decision: High Availability

If the NetScaler Gateway is unavailable, remote users will not be able to access the Citrix environment. Therefore at least two NetScaler Gateway servers should be deployed to prevent this component from becoming a single point of failure.

When configuring NetScaler Gateway in a high availability pair, the secondary NetScaler Gateway monitors the first appliance by sending periodic messages, also called a heartbeat message or health check, to determine if the first appliance is accepting connections. If a health check fails, the secondary NetScaler Gateway tries the connection again for a specified amount of time until it determines that the primary appliance is not working. If the secondary appliance confirms the health check failure, the secondary NetScaler Gateway takes over for the primary NetScaler Gateway.

Each NetScaler Gateway appliance must be running the same version of the NetScaler Gateway software and have the same license. For additional considerations when configuring a NetScaler Gateway high availability pair please reference Citrix eDocs – [How High Availability Works](#).

Decision: Platform

In order to identify an appropriate NetScaler platform to meet project requirements, the key resource constraints must be identified. Since all remote access traffic will be secured using the secure sockets layer (SSL, transported by Hypertext Transfer Protocol (HTTP) in the form of HTTPS), there are two resource metrics that should be targeted:

- **SSL throughput** – The SSL throughput is the gigabits of SSL traffic that may be processed per second (Gbps).
- **SSL transaction per second (TPS)** – The TPS metric identifies how many times per second an Application Delivery Controller (ADC) may execute an SSL transaction. The capacity varies

primarily by the key length required. TPS capacity is primarily a consideration during the negotiation phase when SSL is first setup and it is less of a factor in the bulk encryption / decryption phase, which is the majority of the session life. While TPS is an important metric to monitor, field experience has shown that SSL Throughput is the most significant factor in identifying the appropriate NetScaler Gateway.

For more information, please refer to the Citrix white paper – [Best Practices for Implementing 2048-bit SSL](#).

The SSL bandwidth overhead average is often considered negligible relative to the volume of virtual desktop traffic and is not typically accounted for as part of required SSL Throughput. However making provisions for SSL bandwidth will help ensure the total throughput estimated is sufficient. The fixed bandwidth added to packet headers can vary according to the encryption algorithms used and the overall percentage of bandwidth may vary widely according to packet size. Ideally, the overhead should be measured during a proof of concept or pilot. However, in the absence of such data incrementing the workload bandwidth by 2% is a reasonable rule of thumb. Therefore, to determine the SSL throughput required by a NetScaler platform, multiply the maximum concurrent bandwidth for a datacenter by 1.02:

$$\text{SSL Throughput} = \text{Maximum Concurrent Bandwidth} * 1.02$$

For example, assuming 128Mbps maximum concurrent bandwidth, the appropriate NetScaler model can be determined as follows:

$$\sim 130\text{Mbps} = 128\text{Mbps} * 1.02$$

The SSL throughput value should be compared to the throughput capabilities of various NetScaler platforms to determine the most appropriate one for the environment. There are three main platform groups available, each of which provides broad scalability options.

- **VPX** – A NetScaler VPX device provides the same full functionality as hardware NetScaler. However, NetScaler VPXs can leverage ‘off the shelf’ servers for hosting and are suitable for small to medium sized environments.
- **MDX** – A NetScaler MDX is the hardware version of the NetScaler devices. The MDX device is more powerful than the virtual NetScaler and can support network optimizations for larger scale enterprise deployments.
- **SDX** – A NetScaler SDX is a blend between the virtual and physical NetScaler devices. An SDX machine is a physical device capable of hosting multiple virtual NetScaler devices. This consolidation of devices aids with reducing required shelf space and device consolidation. NetScaler SDXs are suitable for handling network communications for large enterprise deployments.

SSL throughput capabilities of the NetScaler platforms may be found in the [Citrix NetScaler data sheet](#). Therefore, based on the example calculation above, a NetScaler MPX 5500 appliance would be sufficient to handle the required load. However, actual scalability will depend on security requirements. NetScaler SSL throughput decreases with the use of increasingly complex encryption algorithms and longer key lengths. Also, this calculation represents a single primary NetScaler. At a minimum, N+1 redundancy is recommended which would call for an additional NetScaler of the identical platform and model.

Note: The Citrix NetScaler data sheet typically represents throughput capabilities under optimal conditions for performance. However, performance is directly affected by security requirements. For example, if the RC4 encryption algorithm and a 1k key length are used, a VPX platform may be able to handle more than 500 HDX proxy connections. However, if a 3DES encryption algorithm and 2k key length are used (which are becoming more common), the throughput may be halved.

[Click here to provide feedback](#)

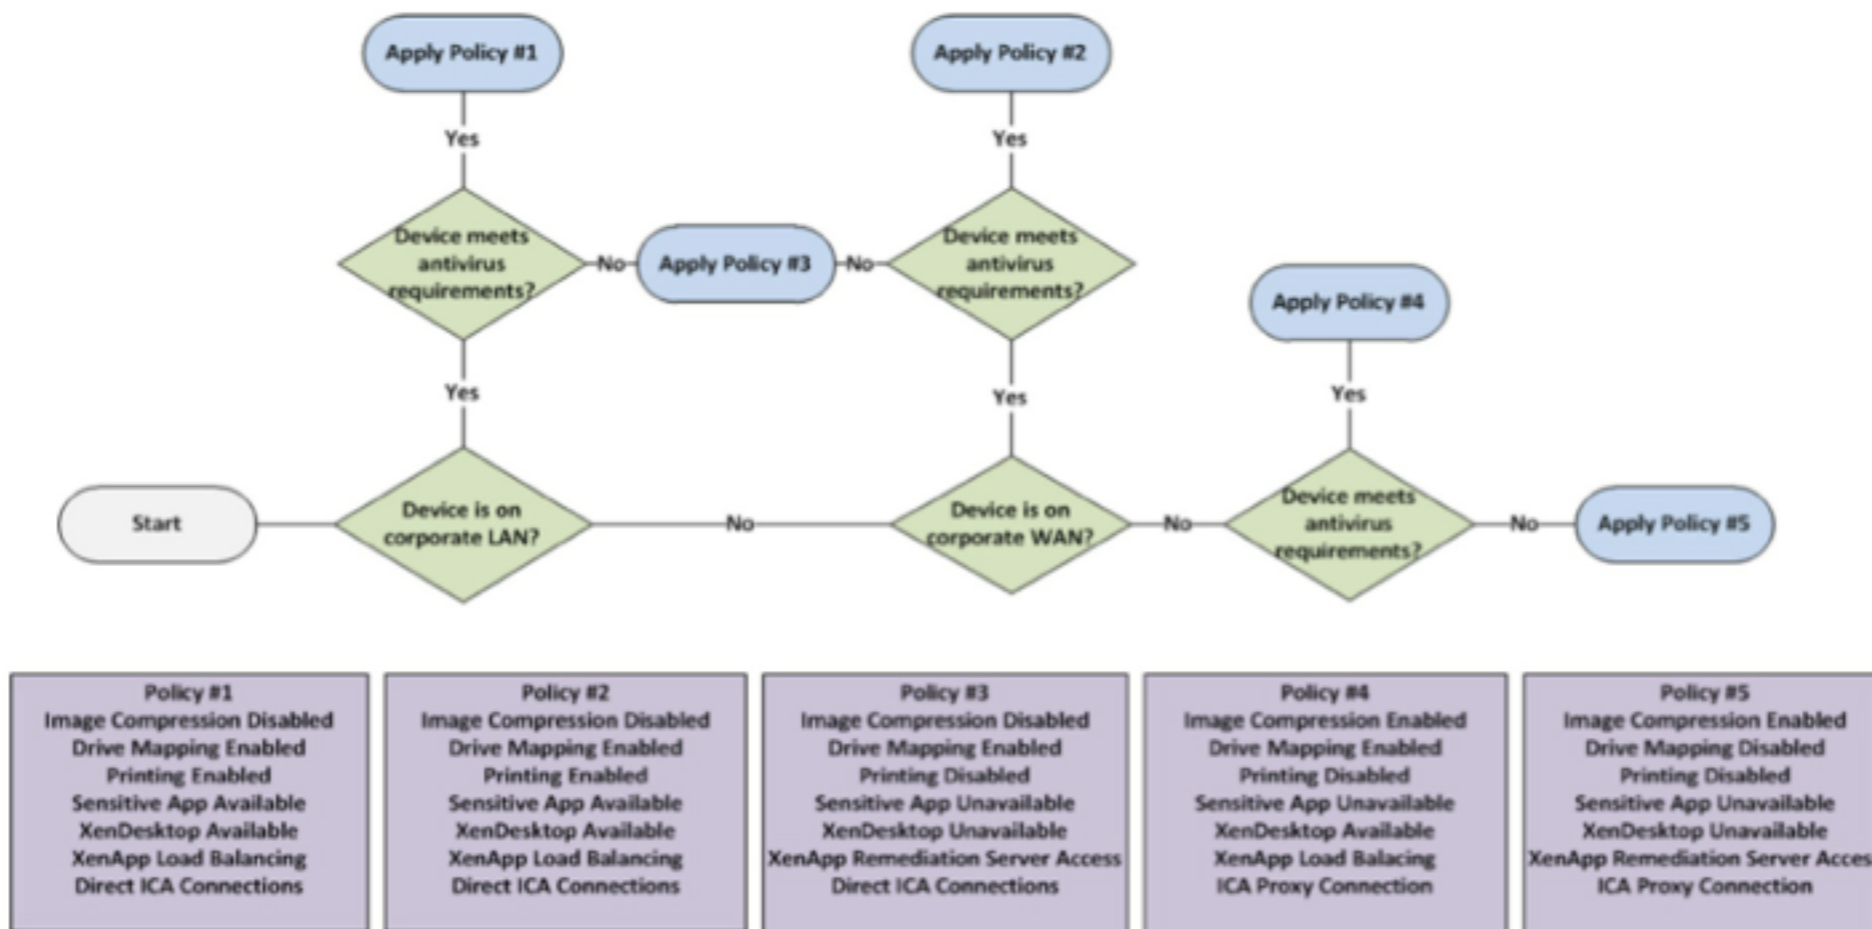
Decision: Pre-Authentication Policy

An optional pre-authentication policy may be applied to user groups with NetScaler Gateway as their authentication point (this design decision is not applicable for non-NetScaler Gateway authentication points). Pre-authentication policies limit access to the environment based on whether the endpoint meets certain criteria through Endpoint Analysis (EPA) Scans.

Pre-authentication access policies can be configured to test antivirus, firewall, operating system, or even registry settings. These policies are used by XenDesktop to control features such as clipboard mapping, printer mapping and even the availability of specific applications and desktops. For example, if a user device does not have antivirus installed, a filter can be set to hide sensitive applications.

The figure below, provides an overview of how multiple policies can be used to customize the features of a virtualization resource:

Simplified SmartAccess Decision Logic



Experience from the Field

Retail – A small private retail company uses EPA to scan for the presence of updated antivirus definitions prior to allowing access.

Financial – A medium financial enterprise uses EPA scans of the Domain SID to verify that users are members of the enterprise domain prior to allowing Access.

Government – A large federal institution uses EPA to scan endpoint devices to ensure that a specific certificate (or set of certificates) has been installed on the device prior to allowing access.

Decision: Session Policy

User groups with NetScaler Gateway as their authentication point must have corresponding session policies defined. Session policies are used to define the overall user experience.

Organizations create sessions policies based on the type of Citrix Receiver used. For the purpose of session policy assignment, devices are commonly grouped as either non-mobile (such as Windows OS based), or mobile (such as iOS or Android). Therefore a decision on whether to provide support for mobile devices, non-mobile devices, or both should be made based on client device requirements identified during the [assess](#) phase.

To identify devices session policies, include expressions such as:

- **Mobile devices** – The expression is set to REQ.HTTP.HEADER User-Agent CONTAINS Citrix Receiver which is given a higher priority than the non-mobile device policy to ensure mobile devices are matched while non-mobile devices are not.
- **Non-mobile devices** – The expression is set to ns_true which signifies that it should apply to all traffic that is sent to it.

An alternative use of session policies is to apply endpoint analysis expressions. These session policies are applied post authentication yet mimic the previously mentioned pre-authentication policies. Use of session policies is an option to provide a fallback scenario to

endpoints that do not meet full security requirements such read-only access to specific applications.

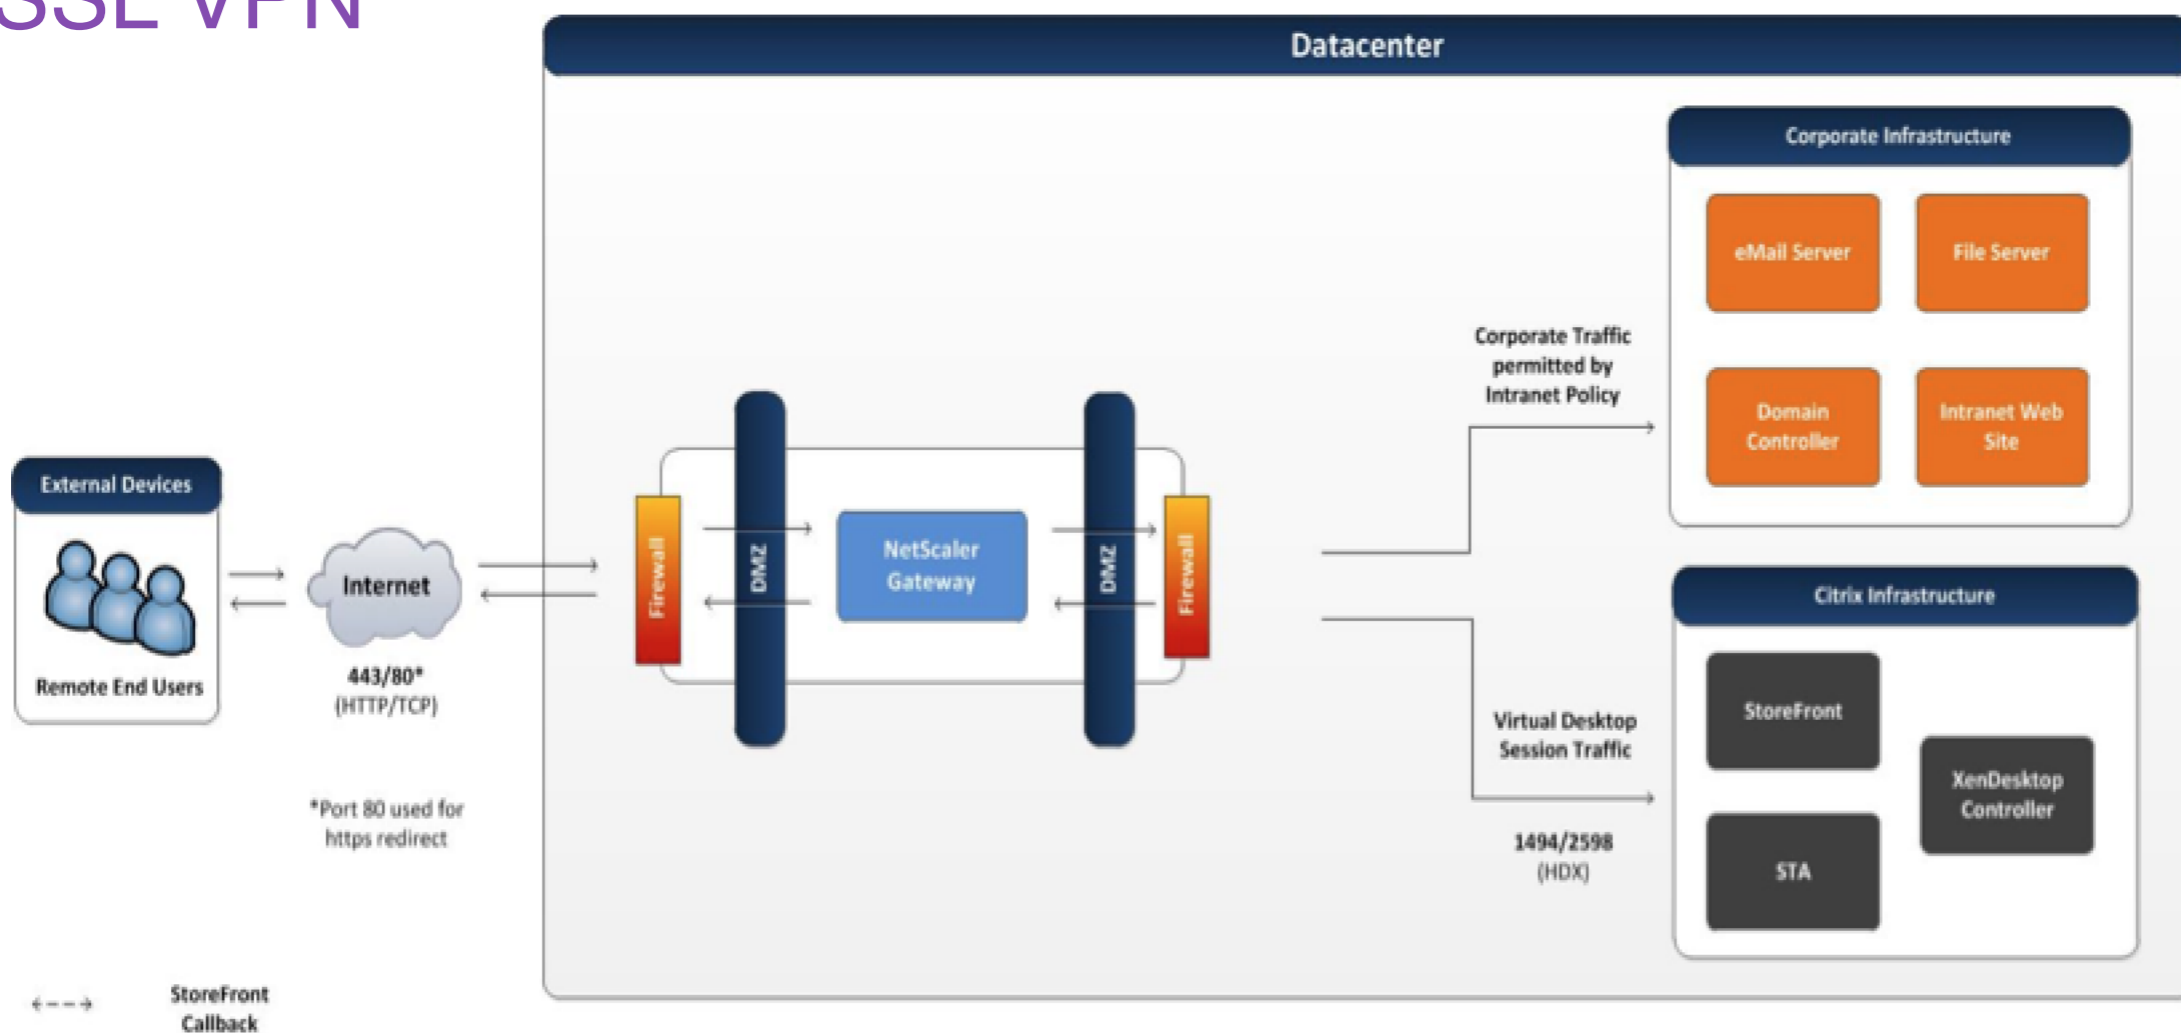
Decision: Session Profile

Each session policy must have a corresponding session profile defined. The session profile defines details required for the user group to gain access to the environment. There are two primary forms of session profiles that determine the access method to the virtual desktop environment:

- **SSLVPN** – Users create a virtual private network and tunnel all traffic configured by IP addresses through the internal network. The user's client device is able to access permitted intranet resources as if it were on the internal network. This includes XenDesktop sites and any other internal traffic such as file shares or intranet websites. This is considered a potentially less secure access method since network ports and routes to services outside of the virtual desktop infrastructure may be opened leaving the enterprise susceptible to risks that may come with full VPN access. These risks may include denial of service attacks, attempts at hacking internal servers, or any other form of malicious activity that may be launched from malware, trojan horses, or other viruses via an Internet based client against vulnerable enterprise services via routes and ports.

Another decision to consider when SSLVPN is required is whether to enable split tunneling for client network traffic. By enabling split tunneling, client network traffic directed to the intranet by Citrix Receiver may be limited to routes and ports associated with specific services. By disabling split tunneling, all client network traffic is directed to the intranet, therefore both traffic destined for internal services as well as traffic destined for the external services (Internet) traverses the corporate network. The advantage of enabling split tunneling is that exposure of the corporate network is limited and network bandwidth is conserved. The advantage of disabling split tunneling is that client traffic may be monitored or controlled through systems such as web filters or intrusion detection systems.

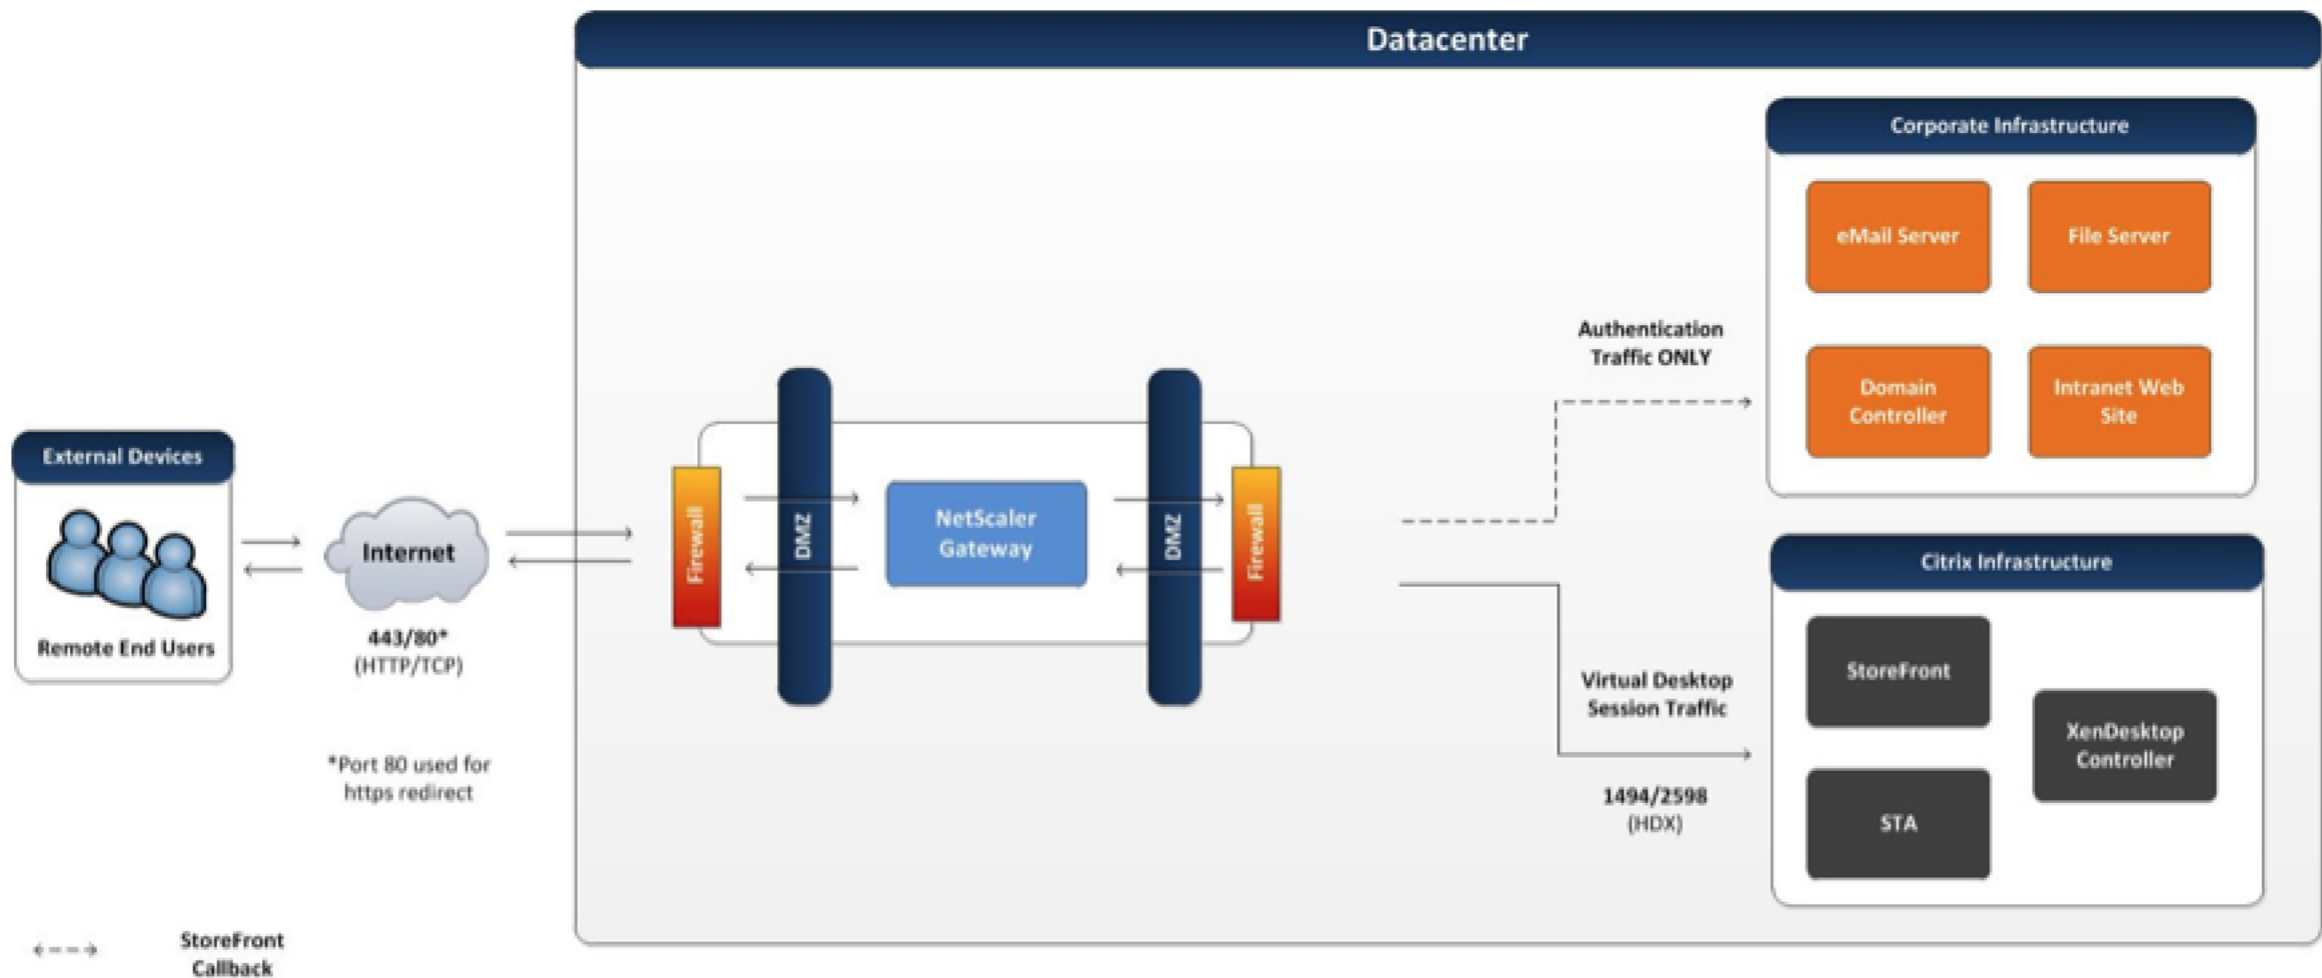
SSL VPN



- **HDX proxy** – With HDX Proxy, users connect to their virtual desktops and applications through the NetScaler Gateway without exposing internal addresses externally. In this configuration, the NetScaler Gateway acts as a micro VPN and only handles HDX traffic. Other types of traffic on the client's endpoint device, such as private mail or personal Internet traffic do not use the NetScaler Gateway.

Based on the endpoint and Citrix Receiver used, a decision must be made as to whether this method is supported for each user group. HDX Proxy is considered a secure access method for remote virtual desktop access since only traffic specific to the desktop session is allowed to pass through to the corporate infrastructure. Most Citrix Receivers support HDX Proxy and it is the preferred method:

HDX Proxy



Decision: Preferred Datacenter

Enterprises often have multiple active datacenters providing high availability for mission critical applications. Some virtual desktops or applications may fall into that category while others may only be accessed from a specific preferred datacenter. Therefore, the initial

NetScaler Gateway that a user authenticates to in a multi-active datacenter environment may not be within the preferred datacenter corresponding to the user's virtual desktop resources. StoreFront is able to determine the location of the user's assigned resources and, direct the HDX session to those resources.

There are static and dynamic methods available to direct HDX sessions to their virtual desktop resources in their primary datacenter. The decision regarding which method to select should be based on the availability of technology to dynamically assign sites links such as Global Server Load Balancing (GSLB) along with the network assessment of intranet and Internet bandwidth as well as Quality of Service (QoS) capabilities.

Note: For more information on configuring the static and dynamic methods of GSLB, please refer to Citrix eDocs – [Configuring GSLB for Proximity](#).

- **Static**

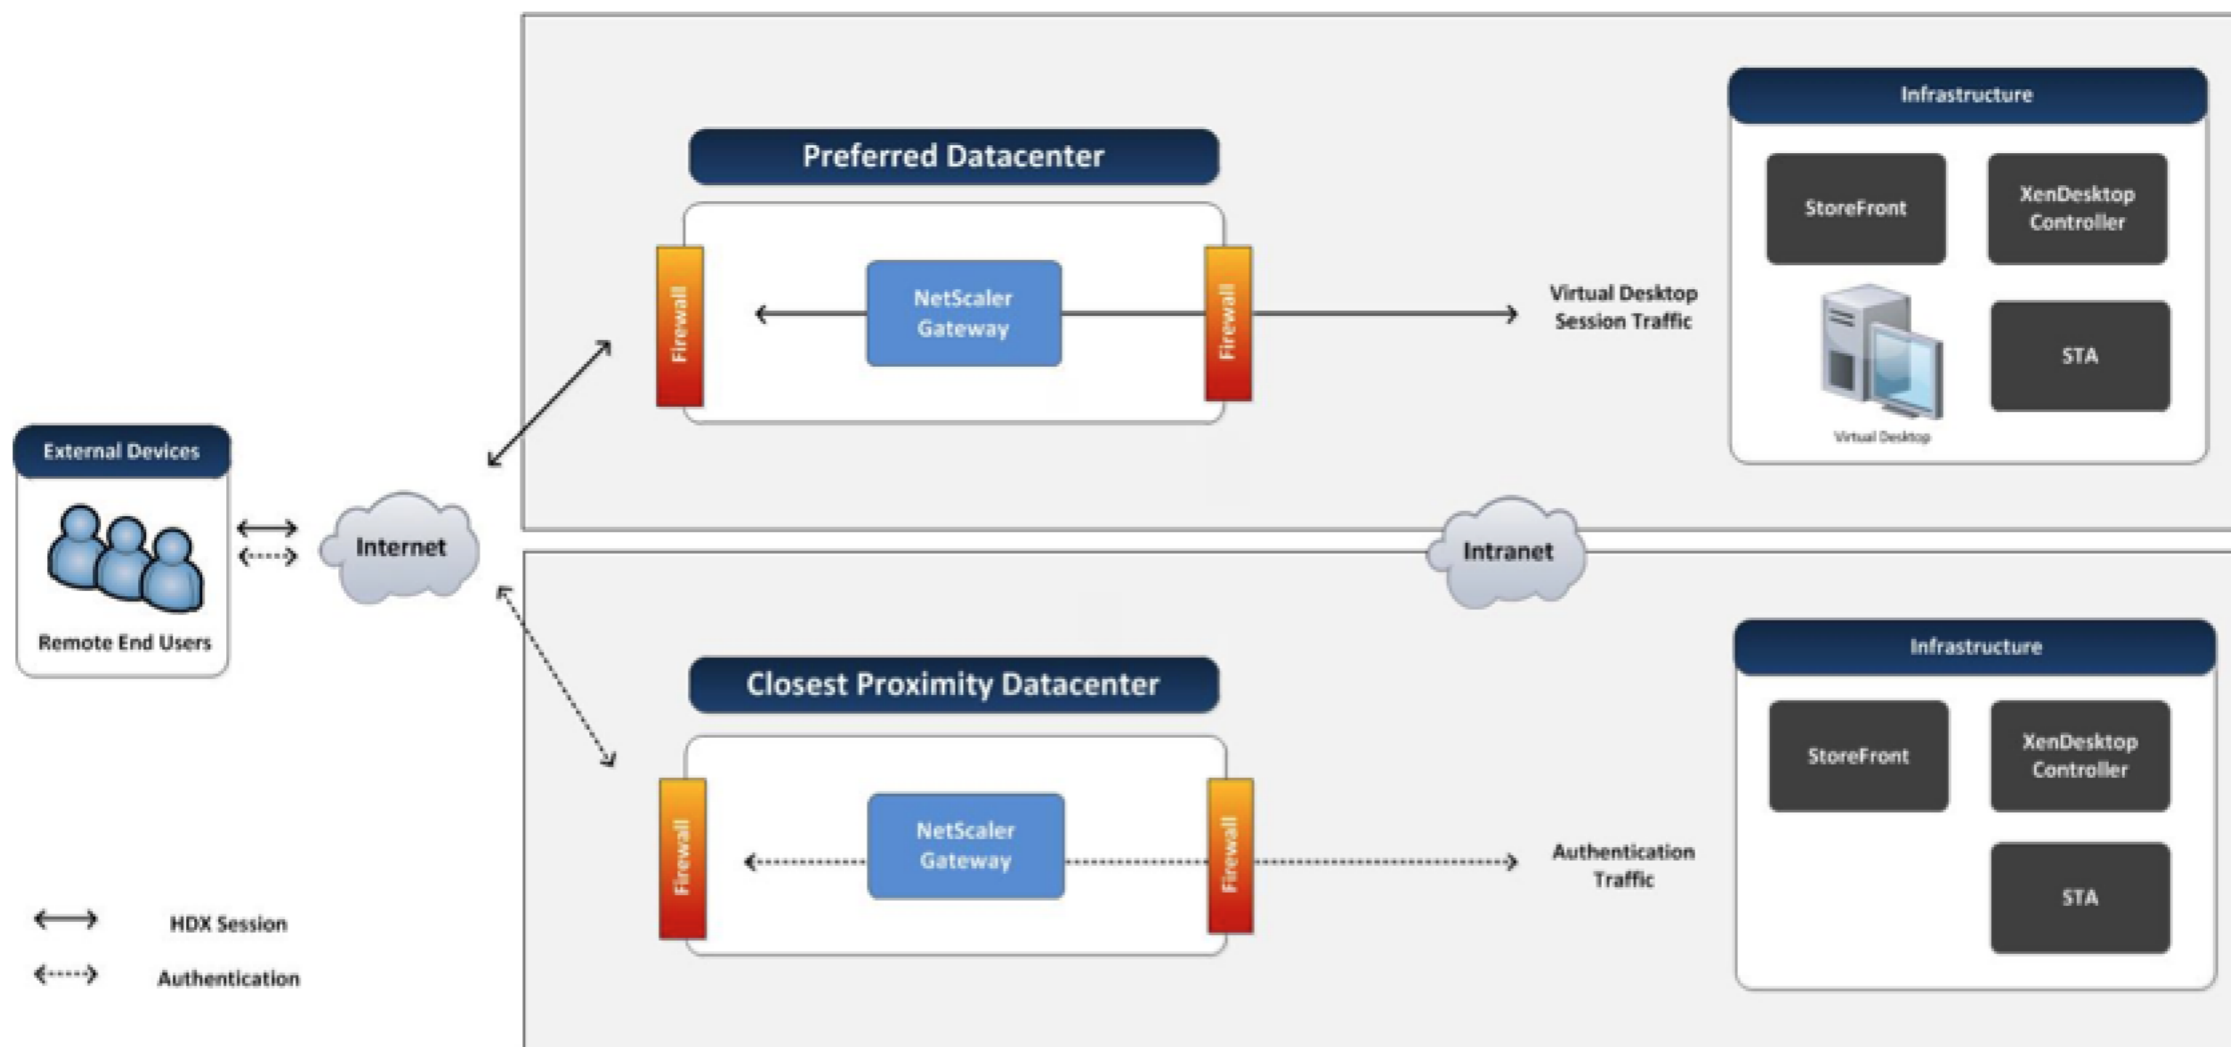
Direct – The user may be given a FQDN mapped to an A record that is dedicated to the primary datacenter NetScaler Gateway(s) allowing them to access their virtual desktop directly wherever they are in the world. This approach eliminates a layer of complexity added with dynamic allocation. However, it also eliminates fault tolerance options such as the ability to access the virtual desktop through an alternative intranet path when a primary datacenter outage is limited to the Internet access infrastructure.

- **Dynamic**

Internet – For most dynamic environments, the initial datacenter selected for authentication is the one closest to the user. Protocols such as GSLB dynamic proximity calculate the least latency between the user's local DNS server and the NetScaler Gateway. Thereafter, the HDX session may be redirected through a NetScaler Gateway to the preferred datacenter by assignment in StoreFront accordingly. However, a significant portion of the HDX session would likely be forced to travel over the best effort Internet without QoS guarantees.

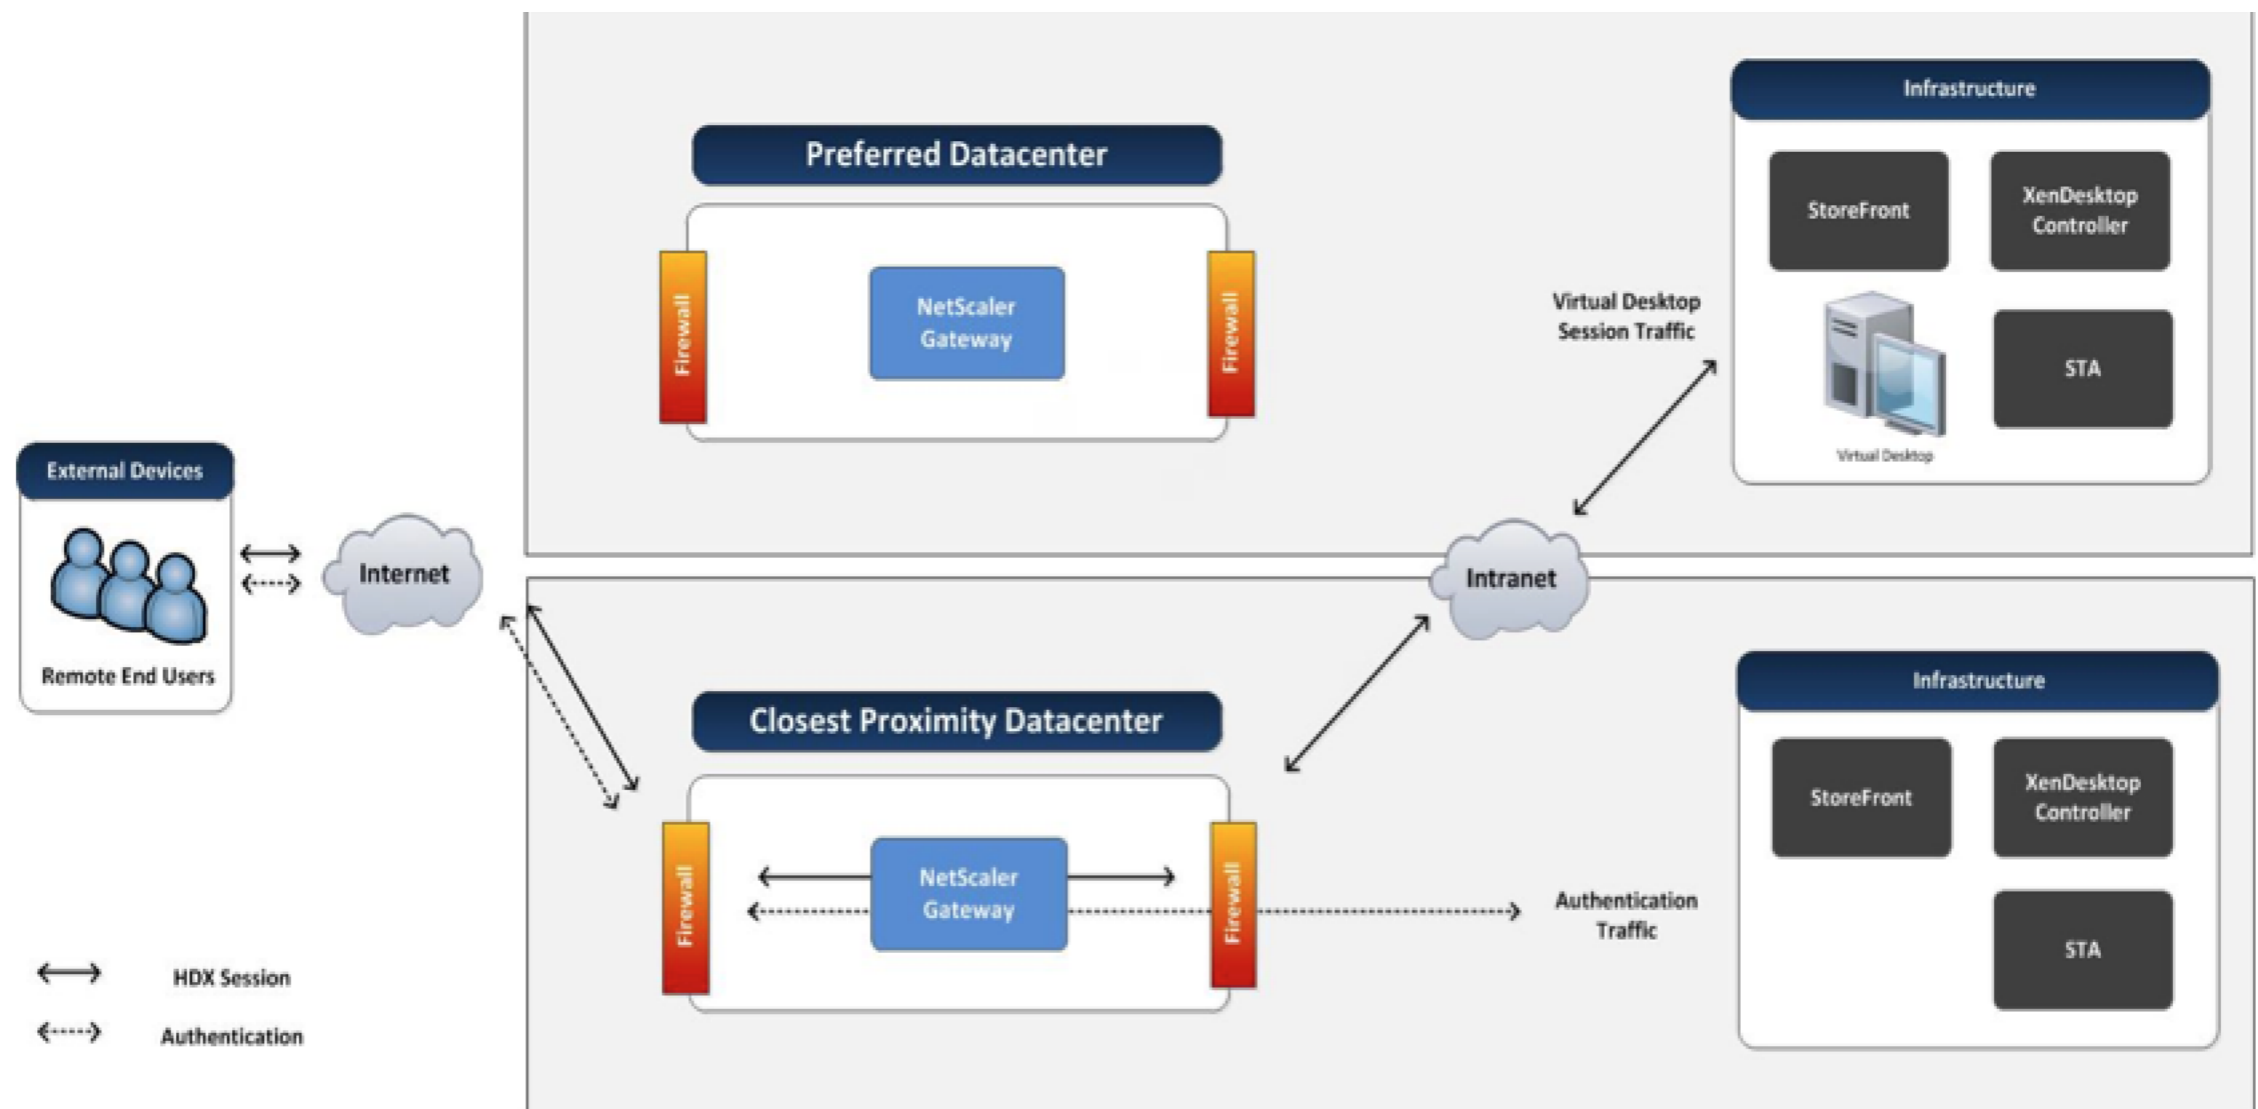
For example, a user with a dedicated desktop in the United States, traveling in Europe may be directed to a NetScaler Gateway hosted in a European datacenter. However, when the user launches their desktop, an HDX connection will be established to the virtual desktop via a NetScaler Gateway hosted in the preferred datacenter in the United States.

Internet Connection



- *Intranet* – Alternatively, the HDX session over which their virtual desktop will be delivered may be transported across the Enterprise WAN and proxied by the same NetScaler Gateway where they authenticated. The advantage of this approach is that a significant portion of the HDX traffic would traverse the Enterprise network where quality of service may be used to provide an optimal user experience. Therefore, this is the recommended dynamic method provided the network assessment identifies that sufficient bandwidth is available or that the project allows for it to be provisioned and that an appropriate QoS infrastructure is in place or may be configured.

Intranet Connection



Resource Layer

The resource layer is the third layer of the design methodology and the final layer focused specifically on the user groups.

The overall user acceptance of the solution is defined by the decisions made within the resource layer. Personalization, applications and overall desktop image design play a pivotal role in how well the desktop is aligned with the user group's requirements, which were identified within the [user data capture](#) and [application data capture](#) sections of the assess phase.

Personalization

User Profiles

A user's profile plays a critical role in determining how successful the user experience is within a virtual desktop or virtual application scenario. Even a well-designed virtual desktop solution can fail if users are frustrated due to lengthy logon times or lost settings.

The user profile solution chosen must align with the personalization characteristics of the user group captured during the [assess phase](#) as well as the FlexCast model selected.

Decision: Profile Type

This section provides an overview on the different profile types available and provides guidance on the optimal user profile for each FlexCast model.

- **Local profiles** – Local profiles are stored on each server or desktop operating system and are initially created based on the default user profile. Therefore, a user accessing these resources would create an independent profile on each system. Users are able to retain changes to their local profile on each individual

system, but changes are only accessible for future sessions on that system. Local profiles require no configuration; if a user logging into a server or desktop operating system does not have a profile path administratively defined, a local profile is created by default.

- **Roaming profiles** – Roaming profiles are stored in a centralized network repository for each user. Roaming profiles differ from local profiles in that the information in the profile (whether it is a printer, a registry setting, or a file stored in the documents folder) can be made available to user sessions accessed from all systems in the environment. Configuring a user for a roaming profile requires an administrator to designate the user's profile path (for virtual desktops) or terminal server profile path to a particular network share. The first time the user logs on to a server or desktop operating system, the default user profile is used to create the user's roaming profile. During logoff, the profile is copied to the administrator-specified network location.
- **Mandatory profiles** – Mandatory profiles are typically stored in a central location for many users. However, the user's changes are not retained at logoff. Configuring a user for a mandatory profile requires an administrator to create a mandatory profile file (NTUSER.MAN) from an existing roaming or local profile and assign users' with a terminal services profile path. This can be achieved by means of Microsoft Group Policy, customizing the user properties in Active Directory or Citrix Profile Management.
- **Hybrid profiles** – Hybrid profiles combine a robust profile core (a mandatory profile or a local default profile) with user specific registry keys or files that are merged during logon. This technique enables administrators to tightly control which changes are retained and to keep the user profiles small in size. Furthermore, hybrid profiles address the last write wins issue using mature queuing techniques that automatically detect and prevent simultaneous writes that could potentially overwrite changes made in another session. Thus minimizing, user

frustration resulting from lost profile changes when accessing multiple servers or virtual desktops simultaneously. In addition, they capture and record only the changes within the profile, rather than writing the entire profile at logoff. A good example of a hybrid profile solution is Citrix Profile Management, which will be discussed in detail within this chapter.

The following table compares the key features of each profile type:

Profile Type Capability Comparison

Feature	Local	Roaming	Mandatory	Hybrid
Central management / roams with user	●	●	● ¹	●
User settings are stored persistently	●	●	●	●
Granular configuration	●	●	●	●
Logon performance and stability enhancements	●	●	●	●

¹ When configured as Mandatory Roaming

● Functionality available ● Optional ● Functionality not available

In order to select the optimal profile type for each user group it is important to understand their personalization requirements in addition to the FlexCast model assigned. The following table provides guidance on selecting the optimal user profile type:

Profile Type Selection

Feature	Local	Roaming	Mandatory	Hybrid
User setting persistence required (personalization characteristic: basic / complete)				
Hosted VDI – random	●	●	●	●
Hosted VDI – dedicated / static with PVD	●	●	●	●
Hosted shared	●	●	●	●
XenClient	●	●	●	●
User setting persistence not required or not desired (personalization characteristic: none)				
Hosted VDI – Random	●	●	●	●
Hosted VDI – dedicated / static with PVD	●	●	●	●
Hosted shared	●	●	●	●
XenClient	●	●	●	●

● Recommended ● Viable ● Not Recommended ● Recommended for users who use a single virtual desktop only ● Recommended for users who use more than one virtual desktop

[Click here to provide feedback](#)

Decision: Folder Redirection

While redirecting profile folders, such as user documents and favorites, to a network share is a good practice to minimize profile size, architects need to be aware that applications may frequently read and write data to profile folders such as AppData, causing potential issues with file server utilization and responsiveness. It is important to thoroughly test profile redirection before implementation in production to avoid these issues. Therefore, it is important to research profile read / write activities and to perform a pilot before moving to production. Microsoft Outlook is an example of one application that regularly performs profile read activities as the user signature is read from the user profile every time an email is created.

The following table provides general recommendations to help identify the appropriate folders to redirect:

Folder Redirection Matrix

Folder	Local	Roaming	Mandatory	Hybrid
Application Data	●	●	●	●
Contacts	●	●	●	●
Desktop	●	●	●	●
Downloads	●	●	●	●
Favorites	●	●	●	●*
Links	●	●	●	●
My Documents	●	●	●	●*
My Music	●	●	●	●
My Pictures	●	●	●	●
My Videos	●	●	●	●
Saves Games	●	●	●	●
Searches	●	●	●	●
Start Menu	●	●	●	●

● Recommended ● Optional ● Not Recommended ● Recommended for backup purposes

*** Recommended for simplified backup / restore and data sharing between profiles

Decision: Folder Exclusion

Excluding folders from being persistently stored as part of a roaming of hybrid profile can help to reduce profile size and logon times. By default Windows excludes the Appdata\Local and Appdata\LocalLow folders, including all subfolders, such as History, Temp and Temporary Internet Files. In addition, the downloads and saved games folders should also be excluded.

Decision: Profile Caching

Local caching of roaming or hybrid user profiles on a server or virtual desktop is default Windows behavior and can reduce login times and file server utilization / network traffic. With profile caching, the system only has to download changes made to the profile. The downside of profile caching is that it can consume significant amounts of local disk storage on multi-user systems, such as a hosted shared desktop server.

Citrix recommends not deleting locally cached profiles for the following scenarios:

For optimizing logon times:

- **Hosted VDI** – Dedicated / Existing / Physical
- **Hosted VDI** – Static with PVD
- **Hosted VDI** – Remote PC
- **Local VM**

For optimizing logoff times:

- Non-persistent virtual desktops and Hosted Shared Desktop servers, with reboot on logoff or daily reboot.

Citrix recommends deleting locally cached profiles on logoff for the following scenarios to avoid the proliferation of stale profiles:

- Hosted shared provided by persistent hosted shared desktop servers.

- Hosted VDI pooled without immediate reboot on logoff.

Configuring the “Delay before deleting cached profiles” Citrix policy sets an optional extension to the delay before locally cached profiles are deleted at logoff. Extending the delay is useful if a process keeps files or the user registry hive open during logoff. This can also reduce logoff times for large profiles.

Decision: Profile Permissions

For security reasons, administrators, by default, cannot access user profiles. While this level of security may be required for organizations that deal with very sensitive data, it is unnecessary for most environments and can complicate operations and maintenance. Therefore, consider enabling the “Add the Administrators security group to roaming user profiles” policy setting. The configuration of this policy should be aligned with the security characteristics of the user groups captured during the assess phase. For more information on the permissions required for the file share, please refer to Microsoft TechNet - [Deploying Roaming Profiles](#).

Decision: Profile Path

Determining the network path for the user profiles is one of the most significant decisions during a user profile design process. In general it is strongly recommended to leverage a redundant and high performance file server or NAS device.

There are three topics that must be considered for the profile share:

- **Performance** – File server performance will effect logon times. For large virtual desktop infrastructures, a single file server cluster may not be sufficient to handle periods of peak activity. In order to distribute the load across multiple file servers, the file server address and share name will need to be adjusted.
- **Location** – User profiles are transferred over the network by means of the SMB protocol, which does not perform well on high-

latency network connections. Furthermore, WAN connections are typically bandwidth constrained, which can add additional delay to the profile load process. Therefore, the file server should be located in close proximity to the servers and virtual desktops to minimize logon times.

- **Operating system platforms** – User profiles have a tight integration with the underlying operating system and it is not supported to reuse a single user profile on different operating systems or different platforms like 64-Bit (x64) and 32-Bit (x86). For more information, please refer to the Microsoft knowledge base article KB2384951 – [Sharing 32 and 64-bit User Profiles](#). Windows 2008 and Windows Vista introduced a new user profile structure, which can be identified by .V2 profile directory suffix, which makes older user profiles incompatible with newer operating systems such as Windows 2012, 7 and 8. In order to ensure that a separate profile is used per platform, the profile directory has to be adapted.

There are two methods that can be used to address these challenges that are based on Windows built-in technology:

- **User object** – For every user object in Active Directory, an individual profile path, which contains file server name and profile directory, can be specified. Since only a single profile path can be specified per user object, it is not possible to ensure that a separate profile is loaded for each operating system platform.
- **Computer group policy or system variables** – The user profile path can also be configured by means of computer specific group policies or system variables. This enables administrators to ensure that a user profile is dedicated to the platform. Since computer specific configurations affect all users of a system, all user profiles will be written to the same file server. To load balance user profiles across multiple servers dedicated XenDesktop delivery groups have to be created per file server.

Note: Microsoft does not support DFS-N combined with DFS-R for

actively used user profiles. For more information, please refer to the Microsoft articles:

- [Information about Microsoft support policy for a DFS-R and DFS-N deployment scenario](#)
- [Microsoft's Support Statement Around Replicated User Profile Data](#)

When using Citrix Profile Management, a third option is available to address these challenges:

User object attributes and variables – Citrix Profile Management enables the administrator to configure the profile path by means of a computer group policy using attributes of the user object in Active Directory to specify the file server dynamically. In order to achieve this, three steps are required:

1. Create a DNS alias (for example, fileserver1) which refers to the actual file server
2. Populate an empty LDAP attribute of the user object (for example, I or UID) with the DNS Alias
3. Configure Citrix Profile Management by means of GPO to use a profile path which refers to the LDAP attribute (for example, If attribute UID is used the profile path becomes \\#UID#\Profiles\profiledirectory)

In addition, Citrix Profile Management automatically populates variables to specify the profile path dynamically based on the operating system platform. Valid profile management variables are:

- **!CTX_PROFILEVER!** – Expands to v1 or v2 depending on the profile version.
- **!CTX_OSBITNESS!** – Expands to x86 or x64 depending on the bit-level of the operating system.
- **!CTX_OSNAME!** – Expands to the short name of the operating system, for example Win7

By combining both capabilities of Citrix Profile Management, a

fully dynamic user profile path can be created, which can be load balanced across multiple file servers and ensure profiles of different operating system platforms are not mixed. An example of a fully dynamic user profile path is shown below:

```
\\#UID#\profiles$\%USERNAME%.%USERDOMAIN%\!CTX_  
OSNAME!!CTX_OSBITNESS!
```

Decision: Profile Streaming

Note: The following design decision only applies to those environments that use Citrix Profile Management.

With user profile streaming, files and folders contained in a profile are fetched from the user store (file server) to the local computer when a user accesses them. During the logon process, Citrix Profile Management immediately reports that the profile load process has completed reducing profile load time to almost zero.

Citrix recommends enabling profile streaming for all scenarios. If it is desired to keep a local cached copy of the user profile, it is recommended to enable the “Always Cache” setting and configure a size of 0. This ensures that the user profile is downloaded in the background and enables the system to use this cached copy going forward.

Note: Profile streaming is not required and does not work with the personal vDisk feature of Citrix XenDesktop. Even if explicitly enabled by means of Group Policy, the profile streaming setting is automatically disabled.

Decision: Active Write Back

Note: The following design decision only applies to those environments that use Citrix Profile Management.

By enabling the active write back feature, Citrix Profile Manager detects when an application has written and closed a file and copies the file back to the network copy of the profile during idle periods. In scenarios where a single user leverages multiple

virtual desktops or hosted shared desktops simultaneously, this feature can be tremendously beneficial. However, Citrix Profile Management does not copy any registry changes back to the network, except during an ordered logoff. As such, there is a risk that the registry and files may get out of alignment on provisioned systems, where locally cached profile information is wiped upon reboot. Therefore, it is recommended to disable active write back functionality for non-persistent Provisioning Services or Machine Creation Services scenarios.

Decision: Configuration Approach

Note: The following design decision only applies to those environments that use Citrix Profile Management.

Citrix Profile Management can be configured by means of an “.ini” file, Microsoft Group Policy and Citrix Policy (Citrix Profile Management 5.0 only). While each option offers the same configuration settings, Group Policy is recommended because it allows administrators to perform Windows and Citrix profile configurations from a single point, minimizing the tools necessary for profile management.

Note: With Citrix Profile Management 5.0, the desktop type is automatically detected and Citrix Profile Management policies set accordingly. For more information, please refer to Citrix eDocs – [How automatic configuration works](#).


Decision: User Profile Replication Between Datacenters

While having an active/active datacenter on a network level is easily accomplished with GSLB, the replication of user data makes having a fully active/active deployment complex in most situations. To have an active/active configuration where users are not statically assigned to a specific datacenter, will require users to have no form of personalization requirements. This will limit the user’s ability to make any configuration changes and will not allow them to create

any documents or persistent data. The exception to this is when a high-speed low latency connection such as dark fibre is available between datacenters. This will let resources in both locations can point to the same file server allowing for a true active/active solution. Also, an active/active configuration can be accomplished when applications are used that rely solely on a backend database that is actively replicated between datacenters and do not store any data in the user profile.

For redundancy and failover purposes, user data such as Windows profiles and documents should be synchronized between datacenters. Although it is recommended to replicate user data between datacenters, the replication would be an active/passive configuration. This means the data can only be actively consumed from a single datacenter. The reason for this limitation is the distributed file locking method inside Windows that only allows a single user to actively write to a file. Therefore, active/active replication of user data is not supported. Any supported configuration consists of a one-way replication of data that is active in a single datacenter at any point in time.

For example, the figure below describes a scenario where user data is passively replicated from Datacenter A to Datacenter B. In this example, File Server A is the primary location for user data in Datacenter A and File Server B is the primary location in Datacenter B. One-way replication of the user data occurs for each file server to allow for the user data to be available in the opposite datacenter if a failover occurs. Replication technologies such as Microsoft DFS can be configured to mirror user profiles and documents to a file server in another datacenter. DFS Namespaces can also be used to have a seamless path for the location of the user data.



User Policies

Citrix policies provide the basis to configure and fine tune XenDesktop environments, allowing organizations to control connection, security and bandwidth settings based on various combinations of users, devices or connection types.

When making policy decisions it is important to consider both Microsoft and Citrix policies to ensure that all user experience, security and optimization settings are considered. For more information on specific Windows related policies, please refer to the Microsoft white paper - [Group Policy Settings Reference for Windows and Windows Server](#), specifically settings related to Windows Server 2008 R2, Windows 7, Windows Server 2012 and Windows 8. For a list of all Citrix-related policies, please refer to the [Citrix Policy Reference](#) spreadsheet.

Decision: Preferred Policy Engine

With XenDesktop 7 organizations have the option to configure Citrix

policies via Citrix Studio or through Active Directory group policy using Citrix ADMX files, which extend group policy and provide advanced filtering mechanisms.

Using Active Directory group policy allows organizations to manage both Windows policies and Citrix policies in the same location, and minimizes the administrative tools required for policy management. Group policies are automatically replicated across domain controllers, protecting the information and simplifying policy application.

Citrix administrative consoles should be used if Citrix administrators do not have access to Active Directory policies. Architects should select one of the above two methods as appropriate for their organization's needs and use that method consistently to avoid confusion with multiple Citrix policy locations.

Decision: Policy Integration

When configuring policies, organizations often require both Active Directory policies and Citrix policies to create a completely configured environment. With the use of both policy sets, the resultant set of policies can become confusing to determine. In some cases, particularly with respect to Windows Remote Desktop Services (RDS) and Citrix policies, similar functionality can be configured in two different locations. For example, it is possible to enable client drive mapping in a Citrix policy and disable client drive mapping in a RDS policy. The ability to use the desired feature may be dependent upon the combination of RDS and Citrix policy. It is important to understand that Citrix policies build upon functionality available in Remote Desktop Services. If the required feature is explicitly disabled in RDS policy, Citrix policy will not be able to affect a configuration as the underlying functionality has been disabled.

In order to avoid this confusion, it is recommended that RDS policies only be configured where required and there is no corresponding policy in the XenDesktop configuration, or

the configuration is specifically needed for RDS use within the organization. Configuring policies at the highest common denominator will simplify the process of understanding resultant set of policies and troubleshooting policy configurations.

Decision: Policy Filtering

Once policies have been created, they need to be applied to groups of users and/or computers based on the required outcome. Policy filtering provides the ability to apply policies against the requisite user or computer groups. With Active Directory based policies, a key decision is whether to apply a policy to computers or users within site, domain or organizational unit (OU) objects. Active Directory policies are broken down into user configuration and computer configuration. By default, the settings within the user configuration apply to users who reside within the OU at logon, and settings within the computer configuration are applied to the computer at system startup, and will affect all users who logon to the system. One challenge of policy association with Active Directory and Citrix deployments revolves around three core areas:

- **Citrix specific computer policies** – Citrix servers and virtual desktops often have computer policies that are created and deployed specifically for the XenDesktop environment. Applying these policies is easily accomplished by creating separate OU structures for the servers and the virtual desktops. Specific policies can then be created and confidently applied to only the computers within the OU and below and nothing else. Based upon requirements, virtual desktops and servers may be further subdivided within the OU structure based on server roles, geographical locations or business units.
- **Citrix specific user policies** – When creating policies for XenDesktop there are a number of policies specific to user experience and security that are applied based on the user's connection to the Citrix environment. However, the user's account could be located anywhere within the Active

Directory structure, creating difficulty with simply applying user configuration based policies. It is not desirable to apply the Citrix specific configurations at the domain level as the settings would be applied to every system any user logs on to. Simply applying the user configuration settings at the OU where the Citrix servers or virtual desktops are located will also not work, as the user accounts are not located within that OU. The solution is to apply a loopback policy, which is a computer configuration policy that forces the computer to apply the assigned user configuration policy of the OU to any user who logs onto the server or virtual desktop, regardless of the user's location within Active Directory. Loopback processing can be applied with either merge or replace settings. Using replace overwrites the entire user GPO with the policy from the Citrix server or virtual desktop OU. Merge will combine the user GPO with the GPO from the Citrix server or desktop OU. As the computer GPOs are processed after the user GPOs when merge is used, the Citrix related OU settings will have precedence and be applied in the event of a conflict. For more information, please refer to the Microsoft TechNet article - Understand [User Group Policy Loopback Mode](#).

- **Active Directory policy filtering** – In more advanced cases, there may be a need to apply a policy setting to a small subset of users such as Citrix administrators. In this case, loopback processing will not work, as the policy should only be applied to a subset of users, not all users who logon to the system. Active Directory policy filtering can be used to specify specific users or groups of users to which the policy is applied. A policy can be created for a specific function, and then a policy filter can be set to apply that policy only to a group of users such as Citrix administrators. Policy filtering is accomplished using the security properties of each target policy.

Citrix policies created using Citrix Studio have specific filter settings available, which may be used to address policy-filtering situations that cannot be handled using group policy. Citrix policies may be

applied using any combination of the following filters:

Citrix Policy Filters

Filter Name	Filter Description	Scope
Access Control	Applies a policy based on access control conditions through which a client is connecting. For example, users connecting through a Citrix NetScaler Gateway can have specific policies applied.	User settings
Citrix CloudBridge	Applies a policy based on whether or not a user session was launched through Citrix Cloud-Bridge.	User settings
Client IP Address	Applies a policy based on the IPv4 or IPv6 address of the user device used to connect the session. Care must be taken with this filter if IPv4 address ranges are used in order to avoid unexpected results.	User settings
Client Name	Applies a policy based on the name of the user device used to connect the session	User settings
Delivery Group	Applies a policy based on the delivery group membership of the desktop running the session.	User settings
Delivery Group Type	Applies a policy based on the type of machine running the session. For example, different policies can be set depending upon whether a desktop is pooled, dedicated or streamed.	User User and computer settings
Organizational Unit	Applies a policy based on the OU of the desktop running the session.	User User and computer settings
Tag	Applies a policy based on any tags applying to the desktop running the session. Tags are strings that can be added to virtual desktops in XenDesktop environments that can be used to search for or limit access to desktops.	User User and computer settings
User or Group	Applies a policy based on the Active Directory group membership of the user connecting to the session.	User settings

Decision: Policy Precedence

With the tree-based structure of Active Directory, policies can be created and enforced at any level in the tree structure. As such, it is important to understand how the aggregation of policies, known as policy precedence flows in order to understand how a resultant set of policies is created. With Active Directory and Citrix policies, the precedence is as follows:

Policy Precedence

Policy Precedence	Policy Type
Processed first (lowest precedence)	Local server policies
Processed second	Citrix policies created using the Citrix administrative consoles
Processed third	Site level AD policies
Processed fourth	Domain level AD policies
Processed fifth	Highest level OU in domain
Processed sixth and subsequent	Next level OU in domain
Processed last (highest precedence)	Lowest level OU containing object

Policies from each level are aggregated into a final policy that is applied to the user or computer. In most enterprise deployments, Citrix administrators do not have rights to change policies outside their specific OUs, which will typically be the highest level for precedence. In cases where exceptions are required, the application of policy settings from higher up the OU tree can be managed using “block inheritance” and “no override” settings. Block inheritance stops settings from higher-level OUs (lower precedence) from being incorporated into the policy. However, if a higher-level OU policy is configured with no override, the block inheritance setting will not be applied. Given this, care must be taken in policy planning, and available tools such as the “Active Directory Resultant Set of Policy” tool or the “Citrix Group Policy Modeling” wizard should be used to validate the observed outcomes with the expected outcomes.

Decision: Baseline Policy

A baseline policy should contain all common elements required to deliver a high-definition experience to the majority of users within the organization. A baseline policy creates the foundation for user access, and any exceptions that may need to be created to address specific access requirements for groups of users. It should be comprehensive to cover as many use cases as possible and should have the lowest priority, for example 99 (a priority number of “1” is the highest priority), in order to create the simplest policy structure

possible and avoid difficulties in determining the resultant set of policies. The unfiltered policy set provided by Citrix as the default policy may be used to create the baseline policy as it is applied to all users and connections. In the baseline configuration, Citrix policies should be enabled with default settings in order to clearly identify the policies applied, and to avoid confusion should default settings change over time.

Citrix Policy templates can be used to configure Citrix policies to effectively manage the end-user experience within an environment and can serve as an initial starting point for a baseline policy. Templates consist of pre-configured settings that optimize performance for specific environments or network conditions. The built-in templates included in XenDesktop are shown below:

XenDesktop 7 Built-in Policy Templates

Built-in Templates	
High definition user experience	Includes settings for providing high quality audio, graphics, and video to users.
High server scalability	Includes settings for providing an optimized user experience while hosting more users on a single server.
Optimized bandwidth for WAN	Includes settings for providing an optimized experience to users with low bandwidth or high latency connections.
Security and control	Includes settings for disabling access to peripheral devices, drive mapping, port redirection, and Flash acceleration on user devices.

For more information on Citrix policy templates, please refer to Citrix eDocs – [Manage Citrix Policy Templates](#).

A baseline policy configuration should also include Windows policies. Windows policies reflect user specific settings that optimize the user experience and remove features that are not required or desired in a XenDesktop environment. For example, one common feature turned off in these environments is Windows update. In virtualized environments, particularly where desktops and servers may be streamed and non-persistent, Windows update creates processing and network overhead, and changes made by the update process will not persist a restart of the virtual desktop or

application server. Also in many cases, organizations use Windows software update service (WSUS) to control Windows updates. In these cases, updates are applied to the master disk and made available by the IT department on a scheduled basis.

In addition to the above considerations, an organization's final baseline policy may include settings specifically created to address security requirements, common network conditions, or to manage user device or user profile requirements:

- **Security policies** – Security policies address policy decisions made to enforce corporate security requirements on the XenDesktop environment. Requirements pertaining to data security and access can be controlled by the correct application of security policy. Users can be allowed to read and write to local or removable media, connect USB devices such as storage devices, smart phones, or TWAIN compliant devices, or cut and paste from the local system based on security requirements. Organizations can also enforce encryption and authentication requirements through security related Citrix policies. Architects should consider the most appropriate level of security and add the policy settings to the baseline policy set, and then address security exceptions through additional policy sets.
- **Connection-based policies** – Connection based policy considerations are used to develop a policy solution that creates the best user experience based on the network environment through which end-users access the network infrastructure. Latency and available bandwidth will determine how to best provide access to audio and video over the HDX connection, providing the best quality experience based on the available resources. Image quality and compression, audio quality and video frame rates can be adjusted based on the connection quality to utilize the bandwidth and network performance appropriately. Multi-stream ICA features can be utilized in concert with network Quality of Service (QoS) to provide an optimized experience for multimedia, input and display and

printing requirements. As with security policies, architects should consider the appropriate base network configuration and add the settings to the initial baseline configuration. Additional network requirements can be dealt with by creating additional higher-level policies to override baseline configurations.

- **Device-based policies** – Device based policy configuration deals with the management of specific device requirements such as tablets and smart phones within an organization. Citrix has created a set of policies to optimize the experience of tablets and smart phones when connecting to XenDesktop environments, allowing these devices to use location services and to customize the user interface where appropriate. Multimedia specific features, such as Windows media and Flash redirection will automatically drop back from client side redirection to server side rendering of media content if the device does not support it; therefore no specific configuration is required to address these features with tablets, or with other devices such as thin clients that may not support these features. Another consideration for device based policy configuration revolves around the security requirements for bring your own devices (BYOD). These elements, such as the need to allow or prohibit local access to hard drives or removable devices, should be addressed through security policy settings.
- **Profile-based policies** – User profiles play a critical role in determining how successful the user experience is within a virtual desktop or virtual application scenario. User profile management can be a key player in mitigating the risks of lengthy logon times or lost settings, providing a consistent user experience across multiple devices, and providing users with their specific data and settings in a virtualized environment. With Citrix Profile Management, policies control two important aspects of user profiles; folder redirection, handled through AD group policy, and Citrix Profile Management settings through Citrix policy. There is more to configuring Citrix Profile Management than simply

turning the features on via Citrix policy. Architects must consider the correct folder redirection configuration for their environment, as well as configuring Citrix policy settings for folder exclusions. Settings for profile streaming and active write back must also be carefully considered based on the size of the profile and whether the operating system is persistent or non-persistent. Profile management policies should be included in the baseline policy if they are to be applied across all users in an organization.

These areas need to be addressed both in the default baseline policy configuration, as well as in any additional policy sets created to address exceptions or additional needs.

Note: Baseline policies are provided in the [Appendix for Profile Management, Microsoft Windows, and Folder Redirection](#). A Citrix baseline policy is provided in the [Citrix Policy Reference spreadsheet](#).

Printing

Citrix XenApp and Citrix XenDesktop support a variety of different printing solutions. In order to plan and successfully implement the proper printing solution it is important to understand the available technologies as well as their benefits and limitations.

Decision: Provisioning Printers

The process of creating printers at the printers at the start of a XenApp or XenDesktop session is called provisioning. There are two types of printer provisioning available:

- **Static** – A predetermined printer or set of printers are created in every session. The same printer or set of printers are provisioned each time, and do not vary according to policies.
- **Dynamic** – The printer or set of printers are provisioned according to policies and may vary with each session based on changes in the policy, user location, or the IP subnet the user

is connected to. This type of provisioning is best suited for larger environments and environments where users can roam and access their XenApp or XenDesktop sessions from different locations.

Dynamic printer provisioning provides more flexibility over static provisioning. However many organizations may choose to deploy a combination of the two. For example, there may be a set of floor printers that everyone in the office is mapped to, and a select group of printers that only members of a specific Active Directory group gets mapped to.

Auto-creation

Auto-creation is a form of dynamic provisioning that attempts to create some or all of the available printers on the client device at the start of a user session. This includes locally attached printers as well as network-based printers. Having all printers available in the XenApp or XenDesktop session may not be necessary for all users, therefore this process can be controlled through a Citrix policy setting called “Auto-create client printers”. This policy has four options:

- **Auto-create all client printers** – This is the default setting.
- **Do not auto-create client printers** – Turns off printer auto-creation when the user logs on. Users can add printers manually during the session unless the “Client printer redirection” policy is set to prohibited.
- **Auto-create the client’s default printer only** – Automatically create only the printer selected as the client’s default printer. Users can change their local default printer and have the change carried through dynamically to the XenApp or XenDesktop session.
- **Auto-create local (non-network) client printers only** – Automatically create only printers directly connected to the client device through LPT, COM, USB, TCP or another port. Network printers are printers mapped to a print server and

use SPOOLSS over RPC, and will not be created.

Auto-creating all client printers can increase the session logon time as each printer is enumerated during the logon process. In contrast, auto-creating the client's default printer only may be too constraining for users that often print to multiple printers. If users want to print to a different printer and the "auto-create the client's default printer only" policy is enabled, they would have to change their local default printer each time to print to a different printer. The Citrix Universal Printer is an alternative to both options by providing access to all the client's printers but only maps one printer during the XenApp or XenDesktop session. The [Citrix Universal Printer](#) is discussed later in this chapter.

Note: The "Auto-create client printers" policy requires the "Client printer redirection" policy to be enabled.

Session Printers

Session printers are a set of network-based printers assigned to users through a Citrix policy at the start of each session. Session printers can be static or dynamic depending on how the policy is configured. The same set of session printers can be created for each session, or may vary by filtering the session printer policy by IP subnet. Filtering by the subnet is also referred to as proximity printing which is covered in more detail in the [Printer Selection](#) section of this document.

Note: Session printers may be assigned using the "Session Printer" policy or the "Printer Assignments" policy. The "Session Printer" policy is intended to be used to set default printers for a farm, site, large group, or OU. The "Printer Assignments" policy is used to assign a large group of printers to multiple users. If both policies are enabled and configured, the session printers will be merged into a single list.

Decision: Managing Print Drivers

Managing print drivers in XenApp and XenDesktop can be a tedious

task, especially in large environments with hundreds of printers. In XenApp and XenDesktop there are several methods available to assist with print driver management.

Automatic Installation

When connecting a printer within a XenApp or XenDesktop session, a check is made to see if the required print driver is already installed in the operating system. If the print driver is not there the native print driver will be installed automatically if one exists, otherwise the Citrix Universal Print Driver will be used.

If users roam between multiple end points and locations, this can create inconsistency across sessions since users may access a different hosted resource every time they connect. When this type of scenario occurs, troubleshooting printing problems and maintenance of print drivers can become very challenging since every hosted resource may have different sets of print drivers installed.

To ensure consistency and simplify support and troubleshooting across XenApp and XenDesktop sessions, enabling automatic installation of print drivers is not recommended.

Manual Installation

When adding a printer within a XenApp or XenDesktop session, and the native print driver is not available and the Citrix Universal Print Driver is not suitable, the drivers can be installed manually.

Manual print driver installation has the same challenges that automatic installation does. Many different print drivers can potentially be installed on different resources creating inconsistency within the environment. To prevent this from occurring, Citrix administrators should install the print drivers on the master image, or stage the print drivers in the image. For more information on how to stage print drivers, please see [Staging Printer Drivers in an Image with PnPUtil](#).

Citrix Universal Print Driver

The Citrix Universal Printer Driver (UPD) is a device independent print driver, which has been designed to work with most printers. The Citrix Universal Printer Driver (UPD) simplifies administration by reducing the number of drivers required on the master image. The Citrix UPD consists of two components:

- **Server component** – The Citrix UPD is installed as part of the XenApp or XenDesktop VDA installation. When a print job is initiated the driver records the output of the application and sends it, without any modification to the end-point device of the user via HDX.
- **Client component** – The Citrix UPD is installed as part of the Citrix Receiver installation. It fetches the incoming print stream for the XenApp or XenDesktop session and forwards it to the local printing sub-system where the print job is rendered using the device specific printer drivers.

The Citrix UPD supports the following print formats:

- **Enhanced Metafile Format (EMF) (default)** – EMF is the 32-bit version of the Windows Metafile (WMF) format. The EMF format is good when printing graphics because the dimensions of the graphic are maintained regardless of the printer's print resolution. The EMF driver is usually perceived as being "faster" because printing can begin after the first page of the print job has been spooled. The EMF driver can only be used by Windows based clients.
- **XML Paper Specification (XPS)** – The XPS driver uses XML to create a platform-independent "electronic paper" similar to Adobe's PDF format. XPS print jobs tend to have a smaller footprint, but is usually perceived to be "slower" than EMF because the print jobs do not begin printing until the printer receives the last page of the job. XPS can deliver a better print performance than EMF, if the application supports the Windows Presentation Foundation (WPF) format, and the printer supports

the XPS format. If either is not supported then the print job will be converted to the EMF format.

- **Printer Command Language (PCL5c and PCL4)** – PCL is a printing protocol developed originally by Hewlett-Packard for inkjet printers. It is used for printing basic text and graphics and is widely supported on HP LaserJet and multifunction peripherals.
- **PostScript (PS)** – PostScript is a computer language that can be used for printing text and vector graphics. The driver is widely used in low-cost printers and multifunction peripherals.

The PCL and PS drivers are best suited when using non-Windows based devices such as a Mac or UNIX client.

Note: The order in which Citrix UPD attempts to use the drivers can be changed using the "Universal driver preference" policy.

The Citrix UPD is compatible with many different print devices and provides a consistent user experience since the same universal driver is installed across the servers and client devices. This is especially useful in large environments with many different types of printers. A drawback to the Citrix UPD however, is that it may not support all of the advanced features available in some printers. There may also be compatibility issues with printers designed to work with specific applications. For example, check printers may not print checks with the proper alignment and formatting unless the manufacturer's specific drivers are used.

Print Driver Mapping

Print driver mapping is the process of overriding the print driver name provided by the client and using the print driver provided by the server. Print driver mapping is useful in situations where the print driver on the client is named differently than the print driver on the server, (for example, "HP LaserJet 4L" versus "HP LaserJet 4") thereby causing the XenApp or XenDesktop session to fail to recognize that the drivers are the same.

This method can reduce the number of print drivers installed on a XenApp server or the XenDesktop master images by allowing administrators to substitute an equivalent driver instead of installing additional drivers. For example, since the HP LaserJet 4 driver is compatible with printers from the HP LaserJet 4 family, the Citrix administrator may choose to have all users printing to HP LaserJet 4x printers to map the HP LaserJet 4 driver instead of installing all of the different drivers available for the various HP LaserJet 4x models.

Print driver mapping can also be used in the following situations:

- The server print drivers are newer than the print drivers on the client.
- The Citrix Universal Print Driver is not compatible with the printer.
- A Print driver is known to cause issues with an application.

Note: Test the behavior of the printers in detail since some printing functionality and formatting may only be available with a specific driver.

Citrix Universal Print Server

The Citrix Universal Print Server (UPServer) extends XenApp and XenDesktop universal printing support to network printing. The Universal Print Server is comprised of the following components:

- **Server component** – The Citrix UPServer component is installed on a Windows-based print server. It retrieves the print data and forwards it to the respective printer by means of the Citrix UPServer.
- **Client component** – The client component is installed on the base image or XenApp server. It receives the EMF or XPS based print stream from the Citrix UPD and forwards it to the print server. Both the print commands and print data are sent using their own respective ports. Print commands are sent over TCP port 8080 and print data is sent over TCP port 7229 by default (though these ports are configurable by policy).

All connected network printers will leverage the Citrix Universal Print Server automatically when the server and client components have been installed and configured successfully.

Note: The “Universal Print Server enable” policy needs to be configured in order to use the UPServer feature.

The Universal Print Server is recommended for remote print server scenarios or environments with numerous network-based printers. Since the print drivers are installed on the print server instead of the XenApp or XenDesktop hosts, this greatly simplifies print driver management.

Decision: Printer Selection

Administrators must plan for how users will select the printers they print to. Users may be allowed to select any printer on the network, none, or just a select group. A combination of methods may be used to provide the most robust printing environment.

Auto-Created and Session Printers

Selecting an auto-created or session printer is the simplest method since it does not require the user to find and add a printer to their XenApp or XenDesktop session. A drawback to this method is that the list of available printers is fixed and extra effort is required to add other printers to the session. Also, depending on how the printer auto-creation policy is configured, all of the user’s printers may not appear in the session. Printing to an auto-created or session printer is best suited for users that print to the same local or network-based printers in every session.

Citrix Universal Printer

The Citrix Universal Printer is a generic printer object that is auto-created at the start of a session and is not linked to a printing device. When using the Citrix Universal Printer it is not required to enumerate the available client printers during logon, which can greatly reduce resource usage and decrease user logon times.

By default the Citrix Universal Printer will print to the client's default printer, however the behavior can be modified to allow the user to select any of their compatible local or network-based printers.

Note: The Citrix Universal Printer leverage the Citrix Universal Print Driver, therefore it is only compatible with Windows based client devices.

The Citrix Universal Printer is best suited for the following scenarios:

- The user requires access to multiple printers both local and network-based which may vary with each session.
- The user's logon performance is a priority and the Citrix policy "Wait for printers to be created" must be enabled due to application compatibility.
- The user is working from a Windows based device or thin client.

Manual Added Printers

Users at times may need to add printers to their XenApp or XenDesktop session. Allowing users to manually add printers gives them the flexibility to select printers by convenience. The drawback to manually adding network-based printers is that it requires the users to know the network name or path of the printers. There is also a chance that the native print driver is not installed in the operating system and the Citrix Universal Print Driver is not compatible, thereby requiring the user to seek administrative assistance.

Manually adding printers is best suited in the following situations:

- Users do not have any Citrix printer auto-creation policies or session printers assigned.
- Users roam between different locations using the same client device (i.e. laptop, tablet).

Proximity Printing

Proximity printing is presenting printers at the start of the XenApp or XenDesktop session based on the Citrix "Sessions Printers" policy

which has been filtered by IP subnet. The network printers created under this policy may vary based on where the user's session is initiated.

Note: To enable proximity printing, the Citrix policy "Default Printer" must also be enabled and a policy filter based on geographic location (i.e. IP address, client's workstation name) must be set.

Proximity printing is recommended in situations where:

- Users roam between different locations using the same client device (i.e. laptop, tablets).
- Thin Clients are used, which do not have the ability to connect to network-based printers directly.

Note: If a user roams with a device while the XenDesktop or XenApp session is active (for example, a laptop and the session is not disconnected while the user roams), proximity printing may not assign the nearest printer until the session is restarted. For example, if an office building has a printer per floor and a laptop user roams from an office to a conference room on the same floor, proximity printing does not need to be refreshed. However, if the user roams to a different floor, the XenApp or XenDesktop session will need to be restarted for proximity printing to display the nearest printer.

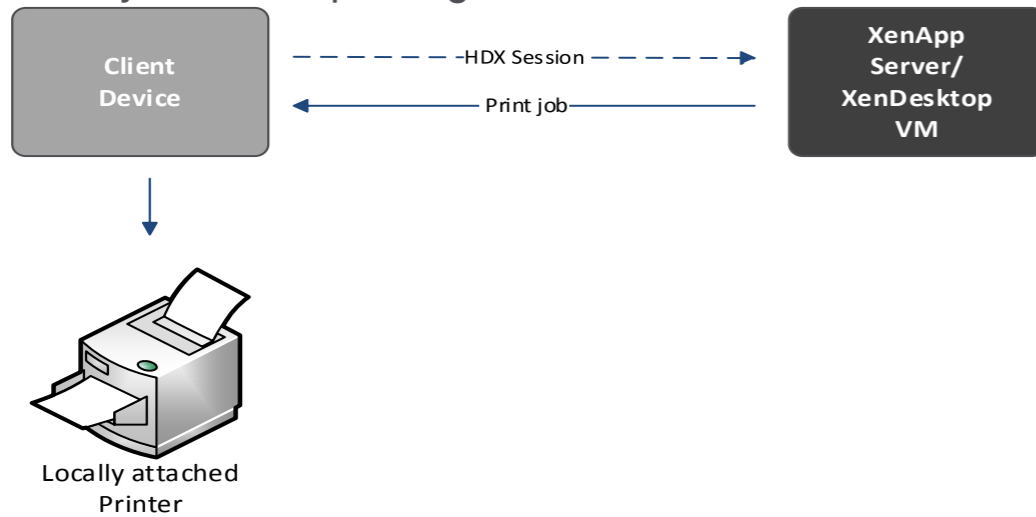
Decision: Print Job Routing

XenApp and XenDesktop print jobs can be routed along two different paths: through the client device or through a network print server.

Client Device Routing

Client devices with locally attached printers (printers attached through USB, LPT, COM, TCP, etc.) will route print jobs directly from the client device to the printer. The ICA protocol is used to send and compress the data.

Locally Attached printing



Print Server Routing

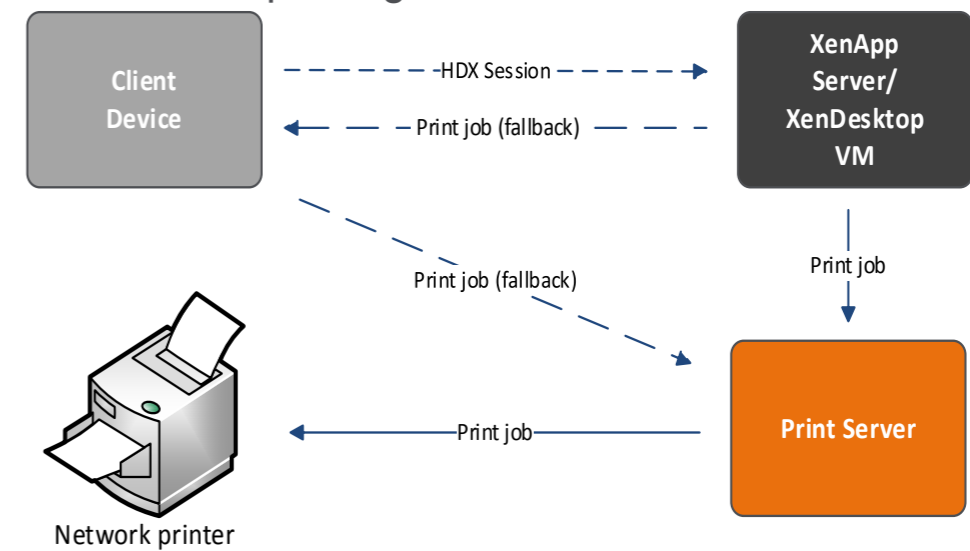
By default, print jobs sent to network-based printers will be routed from the XenApp or XenDesktop session to the print server.

However, the print job will take a fallback route through the client device when any of the following conditions are true:

- The XenApp or XenDesktop session cannot contact the print server.
- The print server is on a different domain without a trust established.
- The native print driver is not available on the master image used by the XenApp servers and XenDesktop virtual machines.

This fallback behavior can cause significant network overhead if not configured correctly.

Network-based printing



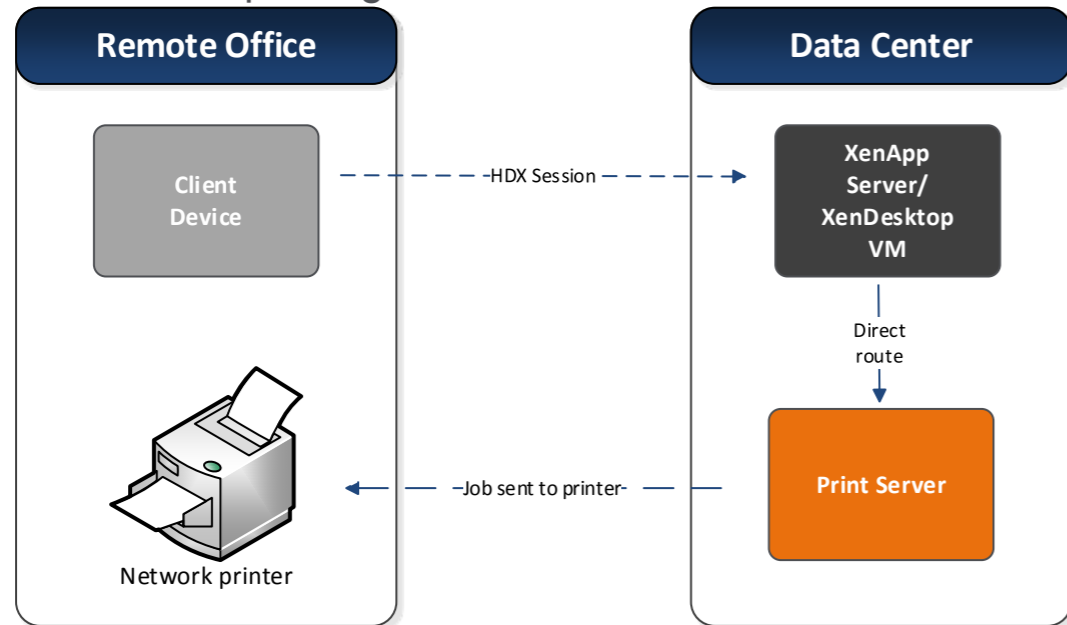
Routing print jobs through a network print server, is ideal for fast local networks, but is not optimal if the XenApp servers, XenDesktop virtual machines, the print server and printers are not on the same LAN. Many packets are exchanged between the XenApp or XenDesktop session and the print server. The end user will experience latency while the print job is spooling over the WAN. The traffic generated by the print job is not compressed and is treated as regular network traffic. The print traffic could consume a significant amount of bandwidth, and negatively impact the experience of other WAN users.

To print over the WAN Citrix recommends routing jobs through the client device so that the ICA protocol compresses the jobs and enables the administrator to limit the maximum consumable bandwidth. Another option is to implement Quality of Service rules to prioritize ICA traffic and help to ensure a good in-session user experience. This option is recommended for Thin Client devices that do not have printing capabilities.

Note: To force network print jobs to route through the client device disable the Citrix policy "Direct connections to print servers".

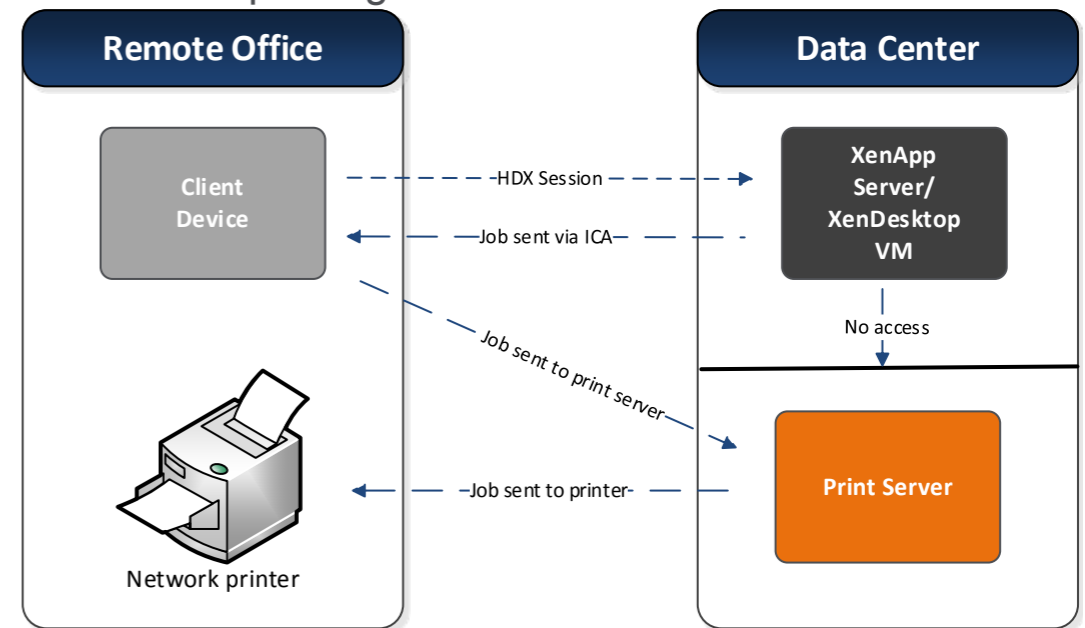
Consider the scenario of a user working in a remote office sending a print job to a network-based printer in the same office. The virtual desktop is hosted in the data center along with the print server. A direct route printing path is taken from the virtual desktop to the print server, and the print job crosses the WAN once to the network printer in the remote office.

Direct route printing



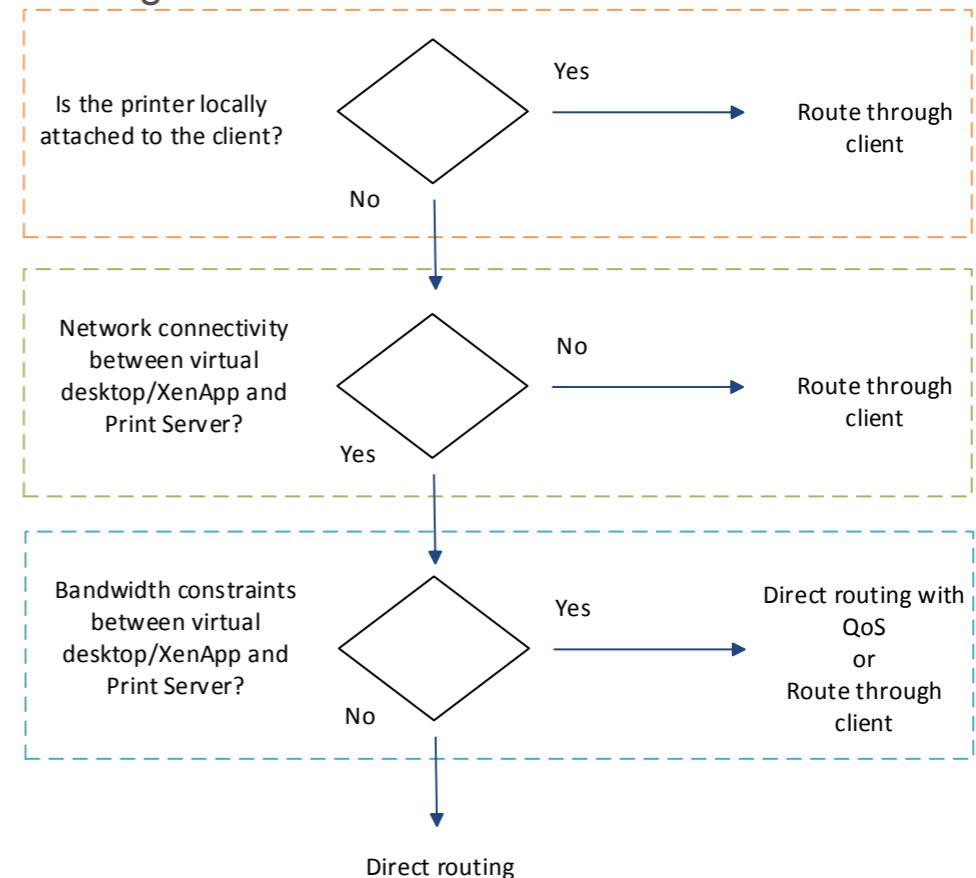
If the virtual desktop is not able to reach the print server directly, the fallback path is taken which routes the print job via ICA through the client device to the print server. The print job in this scenario crosses the WAN three times, but since it is traveling over ICA the print job sent from the virtual desktop to the client device is compressed.

Client route printing



The following flowchart summarizes the print routing decisions:

Printing decision flowchart



Experience from the Field

A print media company leverages Thin Clients and Windows-based workstations at the company headquarters. Network based printers are placed throughout the building (one per floor). Windows print servers reside in the datacenter and manage the network printers. XenDesktop and XenApp servers also reside in the datacenter.

A remote branch office has a few Windows workstations with locally attached printers.

Some employees working from home use Mac OS based computers with a direct attached printer.

Three different print strategies are applied:

Headquarters - *A Citrix Universal Print Server is used for printing within the XenApp and XenDesktop session. Native print drivers are not required on the Windows based workstations. A session printer policy is configured per floor which connects the floor printer as the default printer. The policies are filtered based on the subnet of the thin client for proximity printing.*

Quality of Service (QoS) policies are implemented. Inbound and outbound network traffic on ports TCP 1494 and TCP 2598 are prioritized over all other network traffic. This will prevent HDX user sessions from being impacted by large print jobs.

Branch Office - *Since all branch users work on Windows based workstations, auto-created client printers in conjunction with the Citrix Universal Printer Driver are used. Since the print job is delivered over ICA, the print data is compressed which saves bandwidth. The Citrix Universal Printer Driver ensures all printers connected to the client can be used within the XenApp or XenDesktop session without concern of the printer model used.*

Home Office - *Since the Mac is not compatible with EMF or XPS based Universal Print Drivers, the "HP Color LaserJet 2800 Series PS" driver is used to map the locally attached printer.*

Personal vDisk

Unlike traditional VDI deployments involving pooled desktops, where users lose their customizations and personal applications when the administrator alters the base VM, deployments using personal vDisks retain those changes. This means administrators can easily and centrally manage their base VMs while providing users with a customized and personalized desktop experience. Personal vDisks provide this separation by redirecting all changes made on a user's VM to a separate disk (the personal vDisk). The content of the personal vDisk is blended at runtime with the content from the base VM to provide a unified experience. In this way, users can still access applications provisioned by their administrator in the base VM.

Personal vDisks can be used with Provisioning Services or Machine Creation Services. Generally, the use of personal vDisks is evaluated when user groups require the ability to personalize their virtual desktop. This could include a need to use a variety of departmental applications or general personalization that is beyond what is available in the user profile. However, there is no defined limitation of Citrix personal vDisk technology to these specific use cases. It is entirely possible to utilize personal vDisks in scenarios that may not maximize flexibility or virtual desktop density, yet are entirely appropriate to the enterprise environment.

The need and use of personal vDisk technology must align with the [personalization characteristics](#) of the user group captured during the assess phase as well as the FlexCast model selected.

Note: Citrix personal vDisk technology loads a set of Kernel mode drivers very early in the Windows boot process. By and large, these Phase 1 drivers allow most applications to work in a seamless manner with a personal vDisk attached. Exceptions to this are applications that operate or install components prior to the loading of the Citrix personal vDisk drivers. Specifically, certain device driver installers or antivirus software may not work properly if

installed within the personal vDisk.

Decision: Size

While many factors play a part in deciding the size of a Personal vDisk, there are certain basic requirements that should be followed:

- **Minimum size** – 3GB
- **Maximum size** – *Undefined*

Note: The Personal vDisk can be expanded but cannot be shrunk.

Beyond system requirements, the following major factors influence what size to allocate for a user's Personal vDisk:

- **Anticipated growth** – Most organizations should size personal vDisks well above the minimum requirement of 3GB and should factor in each user group's [workload requirement](#). As a general rule of thumb, the following recommendations provide a good starting point:

Light user – 7GB

Medium user – 10GB

Heavy user – 15GB

- **Storage technologies** – The use of storage technologies such as thin provisioning and de-duplication can significantly reduce storage requirements for personal vDisks allowing larger sizes to be specified:

Thin provisioning – Administrator can use thin provisioning to present more storage space to the virtual machines than is actually available on the storage repository.

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.

- **Folder redirection or cloud data service** – In many smaller environments, utilizing personal vDisk as a profile management solution is common. However, the limitations present in

distributed profiles quickly present themselves when this option is selected. The first step many organizations utilize towards a robust profile solution is to utilize Microsoft special folder redirection in order to redirect user data to a network share. Other organizations employ a cloud based data service such as Citrix ShareFile. Personal vDisks can co-exist with either of these deployment options. Depending on the implementation mechanism, the presence of either of these technologies can significantly reduce the initial sizing of personal vDisks for users.

Experience from the Field

*Citrix – Analysis of Citrix's internal deployment of XenDesktop virtual desktops featuring personal vDisk reveal that most vDisks are utilized at a rate of 36.74%. The average size of a PvD is initially 10 GB and has only been expanded for a small number of power users. This size includes user profiles as well as any data, applications and other aspects accumulated in the course of eight months of usage. In this environment that is *NOT* utilizing a profile management solution or roaming profiles, vDisks populate a few GB up front with user profile data then are slowly accreting changes up to the aforementioned level over time.*

Applications

Choosing an appropriate application delivery method helps improve scalability, management and user experience. Based on the outcome of the [application assessment](#), there are several application delivery methods that can be considered:

- **Installed on image** – The application is part of the base desktop image. The install process involves dll, exe and other files being copied to the image drive as well as registry modifications.
- **Installed on personal vDisk** – The install process is similar to installed on image, except that during the installation process the application files and registry changes are automatically redirected to the user's personal vDisk, which is physically

separate but logically connected to the base desktop image.

- **App streaming** – The application is profiled and delivered to the desktops across the network on-demand. Application files and registry settings are placed in a container on the virtual desktop (or in persistent storage associated with the virtual desktop) and are isolated from the base operating system and each other, which helps to address compatibility issues.
- **Hosted** – The application is installed on a multi-user server and users access the application remotely using the Citrix HDX protocol.
- **VM hosted** – The application is installed on a virtual machine and users run the application remotely using the Citrix HDX protocol.

This section outlines the design decisions pertaining to application delivery. While it is important to understand the technical reasons for delivering an application in one way or another, it is also important to consider that the most optimal way to deliver an application might not be the best way, contingent on factors such as infrastructure

cost and complexity, experience with a given solution, or alignment with the overall design. The final decision on application delivery methods must be linked to the **business drivers** determined during the assess phase and based on technical requirements, capabilities and business processes.

Decision: Application Delivery Method

It is unlikely that a single delivery method will meet all requirements; multiple methods will likely be required to meet the needs of users within the environment. Furthermore, the methods described above are viable for application delivery, but they create different effects on the virtual desktop in terms of operation, resource impact and user experience as shown in the table shown below.

Architects should consider the organization's preference for managing application delivery or multiple images in determining the best approach for application delivery. The table at the bottom of the page provides high-level options for application delivery in a virtual desktop environment

Impact of Application Delivery Methods

	Installed	Installed with PVD	Streamed	Hosted	VM Hosted App
Description	Applications are installed directly on the OS Disk	Applications are installed on personal vDisk (hosted VDI only)	Executed locally, but not installed. streamed on first use	RDS-based applications are installed on a server or desktop and executed remotely	Non-RDS-based applications are installed on a virtual machine and executed remotely
User access	Every user authorized to access the desktop image	Individual users who have installed the application	Only authorized users	Only authorized users	Only authorized users
Updates	Update to base OS image	Update to individual applications	Update to application profile	Update to application on hosted desktop or server	Update to application on virtual machine
Limitations	Operating System dependencies (Windows 8 vs. Windows 7)	Applications that start early in the Windows boot phase (for example, antivirus)	Applications with device driver or those that start services at boot	Compatibility with multi-user windows, Windows Server Operating System	Only one HDX connection to each VM host
Resource usage	Desktop	Desktop	Desktop	Hosting servers or desktops	Hosting desktop

Potential Application Delivery Strategies

	Base Applications	Wide Use Applications	Individual/ Departmental	Resource Intensive	Technically Challenging
Description	Common apps/utilities used by all users	Applications used by large groups of users	Applications installed by individuals or department managed	Heavy system requirements	Complex, extensive dependencies, or requires specialized configuration
Example	Microsoft Office, Adobe Acrobat, antivirus	Corporate developed applications.	Call Centre, Sales applications	CAD/CAM, developer suite	Healthcare management, business management applications
Preferred delivery model	Installed on desktop	Installed or streamed to desktop	Installed on Personal vDisk	Installed or streamed to desktop	Hosted RDS-based application

This is a starting strategy that must be modified to meet the specific implementation, based on decisions made during the [application assessment](#). For example, the following items will play a crucial role in determining the proper mix of delivery techniques for a given implementation.

Compatibility

Desktop virtualization typically requires significant changes to be made to an organization's application delivery strategy. For example, many organizations will take the opportunity to upgrade their desktop operating system and to simplify management by reducing the number of applications installed into the base image using techniques such as application streaming and seamless applications. These are significant changes that require comprehensive compatibility testing. Important compatibility requirements that may need to be verified include:

- **Desktop operating system** – If the application is to be streamed or installed into a hosted VDI desktop, the application must be compatible with the preferred desktop operating system.
- **Server operating system** – Some applications may be more appropriate for delivery via a hosted shared desktop or published application. In these situations, the compatibility of the application must be verified against the chosen server operating system, for example Microsoft Windows Server 2012 R2.
- **Application architecture** – It is important to understand whether the application includes 16-bit, 32-bit or 64-bit code so that an appropriate operating system can be selected. 16-bit code cannot be executed on a 64-bit operating system.
- **Interoperability** – Some applications may experience complications if they coexist on the same operating system. Possible causes include shared registry hives, dll files or INI files as well as incompatible dependencies. Application interoperability issues should be identified so that appropriate

remediation steps can be taken or an alternative delivery model selected.

- **Application streaming** – The use of application streaming helps to simplify image management by reducing the number of applications installed into the base image. However, not all applications are suitable for streaming because they may install device drivers, use COM+ or form part of the operating system. Microsoft App-V limitations will be covered in greater detail in the next section [Decision: Application Streaming](#).

Note: Citrix Application Streaming is not supported on operating systems running Windows Server 2012 and higher, or Windows 8 and higher. Microsoft App-V is the preferred technology for streaming applications. If Citrix Application Streaming is currently being used to stream applications in XenApp 6.x deployments, it will continue to be supported for those environments.

- **Virtual IP Loopback** – Virtual IP allows applications to bind to a unique IP address. Often an application request to bind to a port listening on the address 0.0.0.0. When an application does this and uses a static port, it cannot launch more than one instance of the application. Using the Virtual IP address enables more than one application to listen on the same port on the same computer because they are listening on different addresses. Enabling the virtual loopback policy will allow each session to have its own loopback address (127.0.0.1) in a Winsock call, the virtual loopback feature replaces 127.0.0.1 with 127.x.x.x where x.x.x is a representation of the session ID + 1. The virtual loopback feature is available in XenApp 7.6 and higher.

There are three main techniques that can be used to perform the application compatibility testing component of the application assessment:

- **Manual** – With manual application compatibility testing, each application is installed into the target environment and manually tested. Manual testing is very time consuming because it

requires a full application test per delivery model. In addition, it is difficult to test every aspect of the application and almost impossible to identify every application interoperability conflict. As a result, most of the compatibility issues will be identified by production users.

- **Pre-verified applications** – Most application vendors supply detailed information on which operating systems and delivery models their application supports. Please refer to the application vendor's documentation for more information. Applications compatibility can be verified by checking the [Windows Compatibility Center](#) site for Windows 7 and 8 or the [Windows Server Catalog](#) site for Windows Server 2008 R2 and Windows Server 2012. Microsoft also maintains a spreadsheet of applications that have been verified as compatible with Microsoft Windows 7 by the software vendor or the Windows 7 Logo Program. For more information, please refer to the Microsoft spreadsheet - [Windows Application Compatibility List for IT Professionals](#). The problem with pre-verified applications is that they are unlikely to provide compatibility information for every application identified during the inventory and there won't be any information available on applications that have been fully or partially developed in-house. At best, there will be limited information available on application interoperability issues.

Note: When developing in-house applications, ensure that they are tested using the Microsoft [Application Test Framework](#) for Windows XP and the [App Certification Kit](#) for Windows 7, 8 and 8.1.

- **Automated Tools** – [Citrix AppDNA](#) allows applications to be quickly and accurately analyzed for compatibility across all relevant desktop and server operating systems and delivery models including Citrix XenDesktop and Microsoft App-V. By leveraging AppDNA multiple applications can be tested on multiple platforms simultaneously, saving a significant amount of time and resources needed for compatibility testing.

Applications are imported into AppDNA and analyzed with a series of algorithms providing the user with detailed compatibility and interoperability information. When compatibility issues are identified, AppDNA provides an explanation of the issue, possible remediation actions and an estimate on the time required to resolve the issue.

The advantages and disadvantages of each approach are summarized in the following table:

Application Compatibility Testing Approaches

Approach	Time Required	Cost	Complexity	Interoperability	Accuracy
Manual	High	Low	High	Limited	Low
Pre-verified	Medium	Low	Medium	Limited	Medium
Automated Tool	Low	Medium	Low	Yes	High

Regardless of the approach used, the compatibility testing results should be captured in the compatibility section of the application assessment worksheet:

- **Prerequisites** – Many applications will depend on certain prerequisite to function correctly, for example, the Java Runtime Environment, .Net Framework or a database driver. All essential prerequisites should be captured during the application assessment so that an appropriate image design can be created.
- **Dependent apps** – Applications may need to interact with each other to provide the users with a seamless experience. For example, applications that present information in a PDF format require a suitable PDF viewer to be available. Capturing application dependencies will help to ensure that an appropriate image design can be created.
- **16-bit code** – The application assessment should determine whether the applications include any 16-bit code because they cannot be supported on a 64-bit operating system. There are

three classifications available for 16-bit code in the application assessment worksheet: yes, no and unknown.

- **Windows 8** – Specify whether the application passed compatibility testing for Microsoft Windows 8. There are six classifications available for Windows 8 in the application assessment worksheet: yes – no remediation, yes – low remediation, yes – medium remediation, yes – high remediation, no or unknown.
- **Windows 7** – Specify whether the application passed compatibility testing for Microsoft Windows 7. There are six classifications available for Windows 7 in the application assessment worksheet: yes – no remediation, yes – low remediation, yes – medium remediation, yes – high remediation, no or unknown.
- **Windows XP** – Specify whether the application passed compatibility testing for Microsoft Windows XP. There are six classifications available for Windows XP in the application assessment worksheet: yes – no remediation, yes – low remediation, yes – medium remediation, yes – high remediation, no or unknown.
- **Hosted shared desktop supported** – Specify whether the application passed compatibility testing for running on RDS-enabled server environments. There are six classifications available for hosted shared desktop supported in the application assessment worksheet: yes – no remediation, yes – low remediation, yes – medium remediation, yes – high remediation, no or unknown.
- **Application streaming** – Specify whether the application passed compatibility testing for application streaming. There are three classifications available for application streaming in the application assessment worksheet: App-V, no, or Unknown.

It is unlikely that all users require all applications. Certain applications may be required by departments or a small numbers of

users within / across departments. Architects need to determine the optimal way to deliver applications to users, considering the delivery methods available and the management requirements for various applications. If a large number of applications will be delivered to a subset of users, a separate image may be required. If a small number of users require an application, it can be delivered using an alternative delivery method such as installed on a personal vDisk or streamed.

Business Characteristics

- **IT experience** – If IT has experience or infrastructure already in place for a given application delivery model, then it may make sense to have preference for that model. For example, if an organization has experience using Microsoft App-V for streaming applications, then streamed application delivery may be preferential vs. hosted providing applications can be delivered using both technologies.
- **Management requirements** – Application delivery strategy may depend on who owns and is responsible for a given application. If an application is owned by the corporation, centralized IT can maintain the application and it can be published or streamed. However, if the application is owned and managed by a department and management responsibility cannot be given to centralized IT, then the application may have to be installed in a personal vDisk and managed by the department as a unique entity.
- **Update frequency** – Time and effort required to update an application will impact the selection of a delivery model. If an application must be updated frequently, then architects will want to choose a delivery model which minimizes the number of instances of the application that need to be updated, as well as the complexity of the update process. Hosted applications can be updated in a few locations and the process will be simple. Applications on a personal vDisk will need to be managed

on each individual desktop where it is installed. Streamed applications may have a single source, but a relatively complex update process, including re-packaging and testing. These factors will impact the delivery decision.

- **Licensing** – If an application is available to an organization at no cost, or in a site licensing model, then it can be distributed to users even if they may not use the application. While this may add some complexity to the image, it can simplify the delivery process, reduce the number of images required and the overall complexity of the solution. If application licensing is a concern, architects will need to use a delivery model that conforms to the vendor-licensing model and allows for policy-based control over access to the application.
- **Security** – To protect sensitive data from unauthorized access architects will have to deploy a delivery method that secures the application. A hosted shared desktop or VDI solution can prevent data from ever leaving the datacenter should a user's laptop be stolen. Installed and streamed applications delivered through Citrix StoreFront can be configured so that the user is only able to see the applications for which they have been authorized to use. Achieving a secure application delivery method can also involve a number of measures, including a solution such as a Citrix NetScaler Gateway which can provide granular application-level policy and action controls to secure access to applications and data while allowing users to work from anywhere.

Technical Characteristics

- **Resource intensive** – Applications that make heavy use of system resources are better suited for execution on the virtual desktop as it allows for more granular control over assigned resources and isolation between users is maximized. Consider a delivery method such as installed, streamed or installed on personal vDisk for these applications.

Note: The use of application streaming and personal vDisk

technologies will increase performance overhead.

- **Technically challenging** – If an application is difficult to setup, requires a specialized configuration, or has significant dependencies on other applications, it should be considered technically challenging. Delivering technically challenging applications by hosting on an RDS-enabled server can reduce the complexity of the base image, but consideration needs to be made for dependencies upon other applications which may already exist as part of the base image set (for example, Microsoft Office applications, Adobe Acrobat Reader and email) and the number of users accessing the application.

Experience from the Field

Communications – A large telephone company delivers the majority of its applications using the installed on desktop model, with images focused on user segmentation. Applications that are used by only a few users are installed on personal vDisk to allow customization. Applications that are used by all users, but need isolation from the pooled desktop image are delivered through Microsoft App-V.

Energy – An energy company installs applications on the base image for the majority of users and streams departmental applications as required.

Financial – A banking customer maintains and deploys multiple desktop images containing user group focused applications as required by various departments.

Decision: Application Streaming

In some circumstances it may be necessary to stream applications which require isolation due to compatibility issues with the operating system or other technical issues with the application. XenDesktop 7 supports Microsoft App-V 5.0 as the preferred technology for streaming applications to user devices.

Not every application is capable of being of being streamed.

Microsoft App-V does not support application streaming for the following type of applications.

App-V limitations

Limitation	Description
Applications that start services at boot time	App-V requires a logged in user to initiate the launch of an application.
Applications that require device drivers	App-V cannot virtualize drivers. It is possible to bypass the issue and install driver locally on the target computer. Some user mode device drivers can be virtualized.
Applications that are required by several applications for information or access	For example, a program that will initiate a command and launch another program. Normally both programs would be included in the same suite. However, if this application launches or initiates commands in several applications it may not be feasible to include all of the applications in the same suite.
Applications that are a part of the OS	Such as Internet Explorer
Applications that use COM+	Because COM+ is dynamic and happens at runtime, the App-V Sequencer is not able to capture this information.
COM DLL surrogate virtualization	For example, DLLs that run in DLLhost.exe.

Applications to be streamed will generally fall into one of the following classifications:

- **Simple** – Applications that are relatively small or require very little modification. Retail and off-the-shelf applications usually fall into this category.
- **Moderate** – These applications require some modification during sequencing in order to function properly, or these applications are relatively large. These applications may require changes to the registry, altering the DeploymentConfig or UserConfig file to launch with additional parameters or scripts, or there may be additional applications that need to be installed together as a suite to allow cross functionality.
- **Complex** – These applications have technically challenging installations that require a significant number of customizations to function in a virtual environment. Packages are normally larger than 1GB in size and take an extended period of time to sequence. These applications may also have hard-coded

functions and require manually editing batch and command files to point to resources in the virtual environment. When sequencing these applications it is important to have someone who inherently understands the full functionality of the application.

For more information on Microsoft App-V, please refer to Microsoft TechNet Library – [Microsoft Application Virtualization 5.0 Administrator's Guide](#).

Decision: 16-bit Legacy Application Delivery

There are three options available for delivering 16-bit applications in a XenDesktop environment:

- **Deploy a 32-bit desktop operating system** – A 32-bit desktop operating system is limited to 4GB of RAM, but that is sufficient to support many of the standard business applications in use today, as well as support the legacy 16-bit applications.
- **VM hosted app** – This is the preferred method for delivering legacy applications that cannot run in a 64-bit operating system, or run in a RDS-enabled environment. The 16-bit application is installed on a VM running a 32-bit Windows operating system and made available as a published application. This solution only allows one HDX connection to the VM at a time, so it is intended for applications accessed by few users.
- **XenApp 5 for Windows Server 2008 (x86)** – The 16-bit application is installed on a 32-bit XenApp 5 farm running in parallel with XenDesktop 7. This is a common configuration for organizations migrating away from older versions of XenApp but having to support legacy applications during the transitioning period.

Note: XenApp 5.0 will be the last version of XenApp that supports 32-bit Microsoft Server (Microsoft Server 2008). Windows 2008 R2 or Windows Server 2012 is required for XenDesktop 7 and are 64-bit only. Mainstream support for Microsoft Server 2008 ends on January 13th, 2015. Extended support ends on January 14th, 2020.

End of life (EoL) for XenApp 5 is also January 13th, 2015. Extended support ends on January 14th, 2020.

Images

When selecting an operating system for each user group, architects need to consider the needs of the users and their applications as well as the imaging solution selected. Considerations will include operating system version, architecture, delivery strategy, appropriate policies and how the users and images are aligned into different sets of delivery groups and machine catalogs.

Decision: Operating System

To select an appropriate operating system for each user group, architects must understand user and applications requirements as well as the capabilities of the chosen FlexCast model. The following table outlines which operating systems are recommended in XenDesktop 7 based on FlexCast model:

Differences between 32-bit vs. 64-bit

FlexCast Model	Windows XP	Windows 7	Windows 8	Windows Server 08 R2	Windows Server 2012
Hosted Shared	●	●	●	●	●
VDI: pooled-random	●				
VDI: pooled-static		●	●		
VDI: pooled w/ PvD		●	●		
VDI: dedicated	●	●	●		
VDI: Existing	●	●	●		
VDI: physical / remote PC	●	●	●		
VDI: Streamed	●	●	●		
VDI: streamed with PvD		●	●		
Streamed VHD	●	●	●	●	●
Local VM	●	●	●	●	●
On demand apps				●	●
VM local apps	●	●	●		

● Recommended ● Viable

Decision: Operating System Architecture

A key decision during the XenDesktop design is whether to use a 32-bit (x86) or 64-bit (x64) version of Windows XP, Windows 7 or Windows 8. Windows Server 2008 R2 and Windows Server 2012 are 64-bit only.

The primary benefit of a 64-bit operating system is that significantly more physical memory can be assigned to each desktop - 128GB for Windows XP, 192GB for Windows 7 Professional and 512GB for Windows 8. In contrast, 32-bit operating systems are restricted to 4GB of physical memory.

One disadvantage from choosing a 64-bit desktop operating system is that 64-bit drivers will be required. Finding 64-bit drivers can be

difficult, especially for older peripherals and software. However the main disadvantage is that 64-bit operating systems can't support 16-bit applications and most companies still have some 16-bit applications in use. Even 32-bit applications often include elements of 16-bit code.

Note: [Citrix AppDNA](#) can be used to verify whether applications use 16-bit code or not, in addition to a wealth of additional compatibility information.

If 16-bit applications are required, consider one of the following approaches:

1. Deploy a 32-bit operating system limited to 4GB of RAM for the majority of users. Provide power users with a 64-bit operating system so that they can be assigned more than 4GB of RAM.

Note: Windows 8 is available in both 32-bit and 64-bit versions.

2. Deploy a 64-bit operating system for everyone and use Microsoft Windows 2008 x86 with Citrix XenApp 5.0 to deliver 16-bit applications.

Note: XenApp 5.0 will be the last version of XenApp that supports 32-bit Microsoft Server (Microsoft Server 2008). Windows 2008 R2 or Windows Server 2012 is required for XenDesktop 7 and are 64-bit only. Mainstream support for Microsoft Server 2008 ends on January 13th, 2015. Extended support ends on January 14th, 2020. End of life (EoL) for XenApp 5 is also January 13th, 2015. Extended support ends on January 14th, 2020.

3. Deploy a 64-bit operating system and use VM Hosted Apps to deliver 16-bit applications from 32-bit desktop operating systems.
4. Deploy a 64-bit operating system and replace or re-engineer all applications to be 32-bit or 64-bit.

Decision: Computer Policies

Citrix policies provide the basis to configure and fine tune the

[Click here to provide feedback](#)

XenDesktop environment, allowing organizations to control connection, security and bandwidth settings based on various combinations of users, devices or connection types. Correctly defining an initial baseline policy and assigning additional policies based on security requirements and specific access scenarios is an important aspect when delivering an exceptional user experience.

When making policy decisions it is important to consider both Microsoft Windows and Citrix policies as each have an impact on user experience and environment optimization. For more information on Windows related policies, please refer to the Microsoft spreadsheet – [Group Policy Settings Reference for Windows and Windows Server](#).

Developing a Computer Policy solution for the overall architecture includes the same set of design decisions defined within the [user policy](#) section. These include:

- Preferred policy engine
- Policy integration
- Policy filtering
- Policy precedence
- Baseline policy

Note: Baseline policies are provided in the [Appendix for Microsoft Windows](#), and [Folder Redirection](#). A Citrix baseline policy is provided in the [Citrix Policy Reference](#) spreadsheet.

Decision: Machine Catalogs

Machine catalogs are a collection of virtual or physical machines managed as a single entity. Machine catalogs specify:

- The virtual or physical machines available to host applications or desktops or both
- The Active Directory computer accounts assigned to those virtual machines or computers

- The type of virtual machine being allocated (static or random)
- The provisioning method used to generate the virtual machine
- The operating system installed on the machine
- In some cases, the master image that is copied to create the virtual machines

As a catalog is simply a definition of the virtual desktop resources, a single catalog can span across multiple hosts or hypervisor pools and associated resources such as storage. If a catalog spans across multiple hosts, it is important to ensure that the hosts have access to the appropriate templates for cloning and imaging, depending upon the FlexCast model of the desktop being delivered. Architects will also need to consider the methods used to ensure that the base desktop image is updated across all hosts within the catalog, as required.

Generally, in order to simplify management of the environment, each catalog created should provide machines of the same type (e.g. Static or Random desktops). Although this is not a product-based restriction, constraining catalogs by machine type will allow for simplified management and an easier to understand catalog structure.

Decision: Delivery Groups

Delivery Groups are collections of machines that specify which user groups can access desktops or applications. To deliver applications and desktops, Active Directory users or user groups are assigned to delivery groups. Assigning delivery groups to users can be performed 1:1 or 1 to many, and delivery groups can span multiple catalogs. This process allows architects to better align their desktop allocations with unique user requirements. For example, if a developer user group requires both a corporate desktop for day-to-day operations, and a set of dedicated desktops for development and testing, these desktops can be assigned from a single delivery group. It is important that the following items are considered when

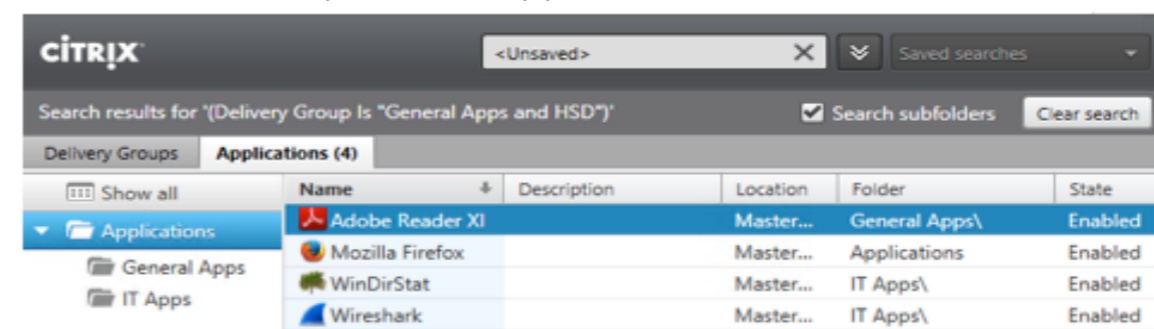
defining user groups:

- Machine catalogs can span multiple hypervisor pools
- Users can be assigned to multiple machine catalogs
- A machine catalog can be associated with one or more delivery groups
- Multiple machine catalogs can reference the same delivery group
- A machine cannot be used in more than one delivery group
- Delivery groups can be created from multiple machine catalogs with the same machine characteristics
- Mixed delivery groups cannot be created from machine catalogs with multiple machine types

Delivery groups also allow an organization to setup delegated administration allowing specific administrators to manage a group of virtual desktops. For example, a delivery group can be setup for executive use, with a defined set of delegated administrators providing VIP support.

Decision: Application Folders

Application folders is a XenApp 7.6 feature that enables administrators to organize applications into logical groups. By organizing applications into folders, administrators will find it easier to manage applications, especially in large environments that can have hundreds of published applications.



Name	Description	Location	Folder	State
Adobe Reader XI		Master...	General Apps\	Enabled
Mozilla Firefox		Master...	Applications	Enabled
WinDirStat		Master...	IT Apps\	Enabled
Wireshark		Master...	IT Apps\	Enabled

Note: In this context Application Folders, is referring to the administrator's view in Citrix Studio. This is not to be confused with folders that users may see when accessing applications in StoreFront or Web Interface.

By default, all applications are placed in a single folder called Applications in Citrix Studio. Folders can be created at the time Delivery Groups are created or they can be added later. Folders can also be created at the time applications are published.

In order to see application folders in Citrix Studio, administrators must be granted the “View Applications” permission. The Edit Application Properties permission is required to remove, rename, or delete a folder that contains applications. Administrators that have been delegated the Delivery Group Administrator rights are able to create and modify applications within application folders, but they cannot create, rename, or delete application folders.

Citrix Administrative Roles

Name	Type
<input checked="" type="radio"/> Delivery Group Administrator Can deliver applications, desktops, and machines; can also manage the...	Built In
<input type="radio"/> Full Administrator You cannot use the Full role with any Scope other than the All scope	Built In
<input type="radio"/> Help Desk Administrator Can view Delivery Groups, and manage the sessions and machines ass...	Built In
<input type="radio"/> Host Administrator Can manage host connections and their associated resource settings.	Built In
<input type="radio"/> Machine Catalog Administrator Can create and manage Machine Catalogs and provision machines.	Built In
<input type="radio"/> Read Only Administrator Can see all objects in specified scopes as well as global information, b...	Built In

There are several use cases for grouping applications into folders:

Use case	Examples
Applications grouped by department use	Finance Apps Sales Apps Marketing Apps
Applications grouped by region	Americas Apps EMEA Apps
Applications grouped by product version	MS Office 2013 Apps MS Office 2010 Apps
Hosting provider with multiple tenants	Tenant1 Apps Tenant2 Apps

For more information on how to configure Application Folders, please refer to eDocs – [Manage applications in a Delivery Group](#).

Decision: StoreFront Integration

StoreFront is used to facilitate the delivery of applications and desktops to end user devices. It is an integral component of XenDesktop that provides a single portal to access desktops and applications from any device. When creating or editing a delivery group in XenDesktop, StoreFront URLs can be pushed to the Citrix Receiver on the virtual desktop so that Citrix Receiver can connect to the store without requiring any user intervention. This is useful when the virtual desktops need access to XenDesktop server-hosted applications.

When creating or editing a Delivery Group in XenDesktop, there are two options available for integrating StoreFront:

- **Automatic** – StoreFront URLs are pre-configured in XenDesktop. This is the preferred method since it doesn't require the end user to enter the address of the StoreFront servers in the environment.
- **Manually** – The StoreFront URL is entered manually when Receiver is launched on the virtual desktop. Administrators may opt for this choice for advanced users, or those that require access to multiple StoreFront stores, like production and test & development environments.

Note: This feature is not available to application-only delivery groups.

Resource Allocation

Resource allocation determines the processor, memory and disk specification of the virtual machines. These decisions have a direct impact on the hardware and storage requirements calculated in the [hardware layer](#).

The key to successful resource allocation is to ensure that virtual desktops and applications can offer similar levels of performance to physical desktops. Otherwise, productivity and overall user satisfaction will be affected. Allocating resources to the virtual machines above their requirements however is inefficient and expensive for the business.

The resources allocated should be based on the workload characteristic of each user group, identified during the [user assessment](#) phase.

Decision: Virtual Processor (vCPU)

For desktop-based operating systems (XenDesktop), Citrix Consulting typically recommends two or more vCPUs per virtual machine so that multiple threads can be executed simultaneously. A single vCPU could be assigned for light workloads, however these desktops are more likely to experience session hangs. In addition, light workloads are more appropriate for server-based operating systems (XenApp) which offers higher levels of scalability.

For server-based operating systems (XenApp), Citrix Consulting typically recommends four vCPUs for Microsoft Windows Server 2008 R2 and eight vCPUs for Microsoft Server 2012 / 2012 R2. Internal scalability testing has shown that the number of users hosted on Microsoft Windows Server 2012 / 2012 R2 can be doubled when the number of processors is increased from four to

eight. The same testing showed that 2008 R2 does not provide the same linear scalability as 2012 / 2012 R2. Fewer, high density virtual machines are typically preferred for simplified management.

The following table provides guidance on the vCPUs that should be assigned based on workload and FlexCast model.

vCPU Assignment Guidelines

Workload	Host Shared/Application Servers ¹	Pooled/Assigned VDI ²	
		Configure For Density	Configure For User Experience
Light	Windows Server 2008 R2 = 4 Windows Server 2012/2012 R2 = 8	1	2
Medium		2	2-4
Heavy		2-4	4+

¹Hosted Shared desktops and application server vCPU assignments will also depend on the number of NUMA nodes in the physical processor. For more information see the [hardware sizing](#) section.

²Assigned VDI virtual machines typically require more vCPUs than Pooled VDI due to additional applications that users may install.

Decision: Virtual Memory (vRAM)

Assigning insufficient memory to the virtual machines will cause excessive paging. If the virtual machines are provisioned using Provisioning Services, insufficient memory will cause increased network traffic as less data can be cached locally. If the recommended “cache in RAM with overflow” feature in Provisioning Services is used, additional memory should be allocated to allow for a significant reduction in IOPS.

Most hypervisors support dynamic memory allocation to automatically provide additional memory to the virtual machines that require it by limiting those that do not. The way that dynamic memory is handled is hypervisor specific and is covered in the [hardware layer](#).

The following table provides guidance on the vRAM that should be assigned based on workload and FlexCast model. These are

recommended guidelines, however if a hardware model has already been selected, reverse sizing can be used to optimally allocate vRAM. For more details on reverse sizing please see the [hardware sizing](#) section.

Memory Assignment Guidelines (GB)

Workload	Host Shared/Application Servers ^{1,3}	Pooled/Assigned VDI ^{2,3,4}	
		Configure For Density	Configure For User Experience
Light	Windows Server 2008 R2 = 12 Windows Server 2012/2012 R2 = 24	1-2	2-3
Medium		2	3-4
Heavy		4	5+

¹Final vRAM allocation to XenApp machines will depend on the vCPU allocation and total RAM available in a process known as reverse sizing. For more on reverse memory sizing, see the [hardware sizing](#) section. As a rule of thumb, assign 3GB of vRAM for every 1 vCPU to your XenApp servers.

²Assigned VDI virtual machines typically require more vRAM than Pooled VDI due to additional applications that users may install.

³If using PVS with the recommended RAM Cache with overflow to disk option, additional memory allocation can reduce IOPS significantly (<1). Testing is ongoing, however benefits have been seen with an additional 512MB for XenDesktop and 2GB for XenApp virtual machines allocated to the RAM buffer.

⁴Memory allocation above 4GB requires a 64-bit version of Windows Desktop OS.

Decision: Storage Space

The amount of storage that each VM requires will vary based on the workload and the image type. If creating dedicated VDI machines without leveraging an image management solution, each VM will require enough storage for the entire OS and locally installed applications. Deploying machines through MCS or PVS can substantially reduce the storage requirements for each VM. Disk space requirements for the write cache and difference disk will depend on application usage and user behavior. However, the following table provides a starting point for estimating disk space requirements based on machine sized with vCPU and vRAM as per

the guidelines above:

Disk Space Requirements

Workload	Pooled VDI	Assigned VDI	Hosted Shared
	Windows 7/8	Windows 7/8 with Personal vdisk ³	2012/2008 R2
Provisioning Services - Write Cache^{1,2}			
Light	5 GB	15 GB	40 GB / 25 GB
Medium	7 GB	17 GB	
Heavy	10 GB	20 GB	
Machine Creation Services - Difference Disk^{4,5}			
Light	5 GB	15 GB	40 GB / 25 GB
Medium	7 GB	17 GB	
Heavy	10 GB	20 GB	

¹The write cache size is minimal following startup but will grow constantly during normal usage even when deleting files. The write cache will be cleared on reboot, unless write cache persistence has been configured.

²In some situations, it may be necessary to redirect data to the write cache disk so that it is preserved between reboots, for example Windows event logs and antivirus definition files.

³Includes PVS write cache or MCS difference disk and 10GB personal vDisk (default size).

⁴Storage that can leverage thin provisioning, such as NFS, is recommended for MCS differencing disks. Without thin provisioning, each virtual desktop will require a difference disk that is the same size as the master disk.

⁵A new difference disk is created each time a non-persistent MCS virtual desktop is started. Once the virtual desktop has successfully started, MCS will remove the old difference disk. The exact amount of additional disk space required depends on the number of concurrent boot cycles. Based on project experience, a good starting point is to allocate 10% of additional storage for non-persistent MCS virtual desktops.

Decision: IOPS

The following table provides guidance on the number of IOPS generated per virtual machine based on workload, operating system and FlexCast model. The IOPS estimates in the table below are an average of the steady state and are not a peak average that takes the boot process, logons and logoffs into account.

Note: The IOPS numbers in this table represent I/O operations within a virtual machine. The final IOPS number will vary based on the chosen RAID level and storage optimization technologies in the hardware layer.

¹Read/Write ratio may differ using Personal vDisk.

²Installed data refers to machines created outside of MCS or PVS and are not pooled desktops.

³The LoginVSI heavy workload does not generate a greater number of IOPS than the medium workload. The heavy IOPS workload is defined here based on experience from the field as well as LoginVSI testing. Certain workloads, such as developers, can require upwards of 150 IOPS and should be planned for using custom scalability testing.

⁴There is a significant variation in IOPS for heavy workloads. Therefore, testing should be performed to determine accurate numbers.

IOPS Requirements by Workload

Workload ¹	Pooled VDI		Assigned VDI (Personal vDisk) ¹		Hosted Shared	
	Windows 8	Windows 7	Windows 8	Windows 7	Windows 2012 -Per User	Windows 2008 R2 - Per User
Installed² (Steady State Read/Write Ratio 50/50)						
Light	7	7	n/a	n/a	5	3
Medium	13	13	n/a	n/a	9	6
Heavy ^{3,4}	26+	26+	n/a	n/a	17+	12+
Provisioning Services (Steady State Cache on HDD) (Read/Write Ratio 20/80)						
Light	5	5	4	4	3	2
Medium	10	10	10	10	6	4
Heavy ^{3,4}	20+	20+	20+	20+	12+	8+
Provisioning Services (Steady State Cache in RAM with overflow) (Read/Write Ratio 40/60)						
Light	1	1	TBD	TBD	1	1
Medium	1	1	TBD	TBD	1	1
Heavy ^{3,4}	1+	1+	TBD	TBD	1+	1+
Machine Creation Services (Steady State Read/Write Ratio 20/80)						
Light	7	7	5	5	5	3
Medium	13	13	12	12	9	6
Heavy ^{3,4}	26+	26+	26+	26+	17+	12+

Decision: Graphics (GPU)

A graphical processing unit (GPU) can be leveraged to improve processor scalability and user experience or enable the use of graphically intensive applications. During the desktop design it is important to decide how the GPU (if used) will be mapped to the virtual machines. There are three methods available.

- Pass-Through GPU – Each physical GPU is passed through to a single XenDesktop virtual machine (single user) or XenApp virtual machine (multiple users).
- Hardware Virtualized GPU – Using XenServer vGPU technology, an NVIDIA GRID is virtualized and shared between multiple machines. Each virtual machine has the full functionality of NVIDIA drivers and direct access to the GPU.
- Software Virtualized GPU – The GPU is managed by the hypervisor and intercepts requests made by the XenDesktop or XenApp virtual machines. The machines do not have direct access to the GPU making this the least preferred method. Software Virtualized GPU is inherently different from software rendered graphics processing within the CPU.

GPU Allocation Options

	Pass-Through GPU	Hardware Virtualized GPU	Software Virtualized GPU
XenServer			
XenDesktop	○	○	✗
XenApp	○	•	✗
Hyper-V			
XenDesktop	✗	✗	• ¹
XenApp	✗	✗	• ¹
ESX			
XenDesktop	○	✗	✗
XenApp	○	✗	✗

“○”: Recommended “•”: Viable “✗”: Not Supported

¹Remote FX requires an RDP connection

For sizing information recommended by NVIDIA and additional information on these technologies please refer to the [hardware sizing](#) section of this handbook.

Decision: Optimizations

Optimizing the virtual desktop image can help to reduce hardware requirements and improve overall user experience. However, image optimizations can introduce risk, complicate troubleshooting and require additional testing effort.

Citrix Consulting currently provide optimization guides for the following operating systems.

- [Microsoft Windows 7.x](#)
- [Microsoft Windows Server 2008 R2](#)
- [Microsoft Windows 8.x](#)

Note: There is no optimization guide for Microsoft Windows 2012 as testing showed limited scalability benefits over the default configuration.

Bandwidth Requirements

While bandwidth is abundant on most corporate LANs, bandwidth on the WAN can very quickly become the limiting resource and must be planned for accordingly. With the growing trend in mobile workstyles, even remote connections through broadband and cellular can be expected and need to be taken into account.

It is important to understand when and how bandwidth can become constrained in the overall architecture in order to compensate for it and deliver a quality experience on any network connection.

Decision: Bandwidth Minimizing Technologies

With the HDX protocol, bandwidth can be controlled using technologies such as Quality of Service (QoS), HDX RealTime, and WAN Optimization with Citrix's own CloudBridge solution. Although these technologies offer great benefit, they should be planned for accordingly and are highlighted briefly below.

- **Traffic Management** – If the network infrastructure is capable of supporting QoS, the XenDesktop traffic can be prioritized by channels to better utilize the available bandwidth. For example, traffic such as printing can be assigned lower priority to ensure that it does not degrade the performance of user sessions. For more on assigning network priorities, see Citrix eDocs – [Assign Priorities to Network Traffic](#). Furthermore, traffic for individual channels can be restricted with Citrix policies allowing for granular limits for how the bandwidth of a session is distributed.
- **HDX RealTime** – Microsoft Lync is a popular video conferencing solution, however high quality audio and video can strain the network in a XenDesktop deployment. HDX RealTime allows this to traffic to circumvent the WAN by flowing directly to the end point while using a virtualized version of Lync. To learn more on configuring this feature, see the Citrix eDocs – [System Requirements for HDX RealTime Optimization Pack 1.4](#).
- **CloudBridge** – Citrix CloudBridge is a WAN Optimization solution that offers TCP flow control, protocol optimization, and the optimization of video delivery. More information on Citrix CloudBridge can be found on the [CloudBridge Product Page](#).

Decision: HDX Encoding Method

The bandwidth consumption and user experience depend on various factors such as the operating system used, the application design, and screen resolution. In XenApp and XenDesktop the bandwidth will also be affected by the chosen HDX encoding

method. There are two main HDX encoding methods available in XenApp 7.x and three in XenDesktop 7.x.

- **Desktop Composition Redirection:** Based on the Aero Redirection feature introduced in XenDesktop 5.x, Desktop Composition Redirection (DCR) allows for the offloading of DirectX commands used by the Desktop Windows Manager to the user's Windows device to reduce the CPU load on the server. Since the Desktop Windows Manager is always on in Windows 8 and above, DCR is the default setting for capable end points. Requirements can be found on the [Citrix eDocs](#).
 - **Pros** – Great user experience, full Aero enabled desktop, low server CPU consumption (higher server scalability)
 - **Limitations** – Only available on DirectX capable Windows end points and requires high bandwidth. Not available on Server operating systems. Not available if HDX 3D Pro VDA is installed.
 - **Recommended Use Case** – Main office LAN or remote home user on broadband due to high bandwidth requirement.
 - **Policy** – ICA\Desktop UI\ “Desktop Composition Redirection” = Enabled
- **H.264 Enhanced SuperCodec** – Initially part of HDX 3D Pro, the H.264 encoder now in the enhanced HDX SuperCodec uses deep compression allowing for the delivery of server rendered video in low bandwidth connections. The codec runs completely on the server CPU and allows for a full Aero enabled desktop on any device.
 - **Pros** – Great user experience, full Aero enabled desktop on any device, low bandwidth consumption, greatly improved delivery of server rendered video.
 - **Limitations** – Increased CPU requirements on the VDA due to the H.264 processing. SuperCodec is optimized for Receiver clients that support H.264, clients without this support may

see increased bandwidth usage.

Minimum Receiver Versions: Windows 3.4+, Mac 11.8+, Linux 13+, HTML5 1.3+ Latest Receiver for iOS and Android.

Also, increased CPU requirements on the user device; therefore, this setting is typically not appropriate for older thin clients (slow CPU) unless they perform the H.264 decoding in separate silicon (not on the CPU).

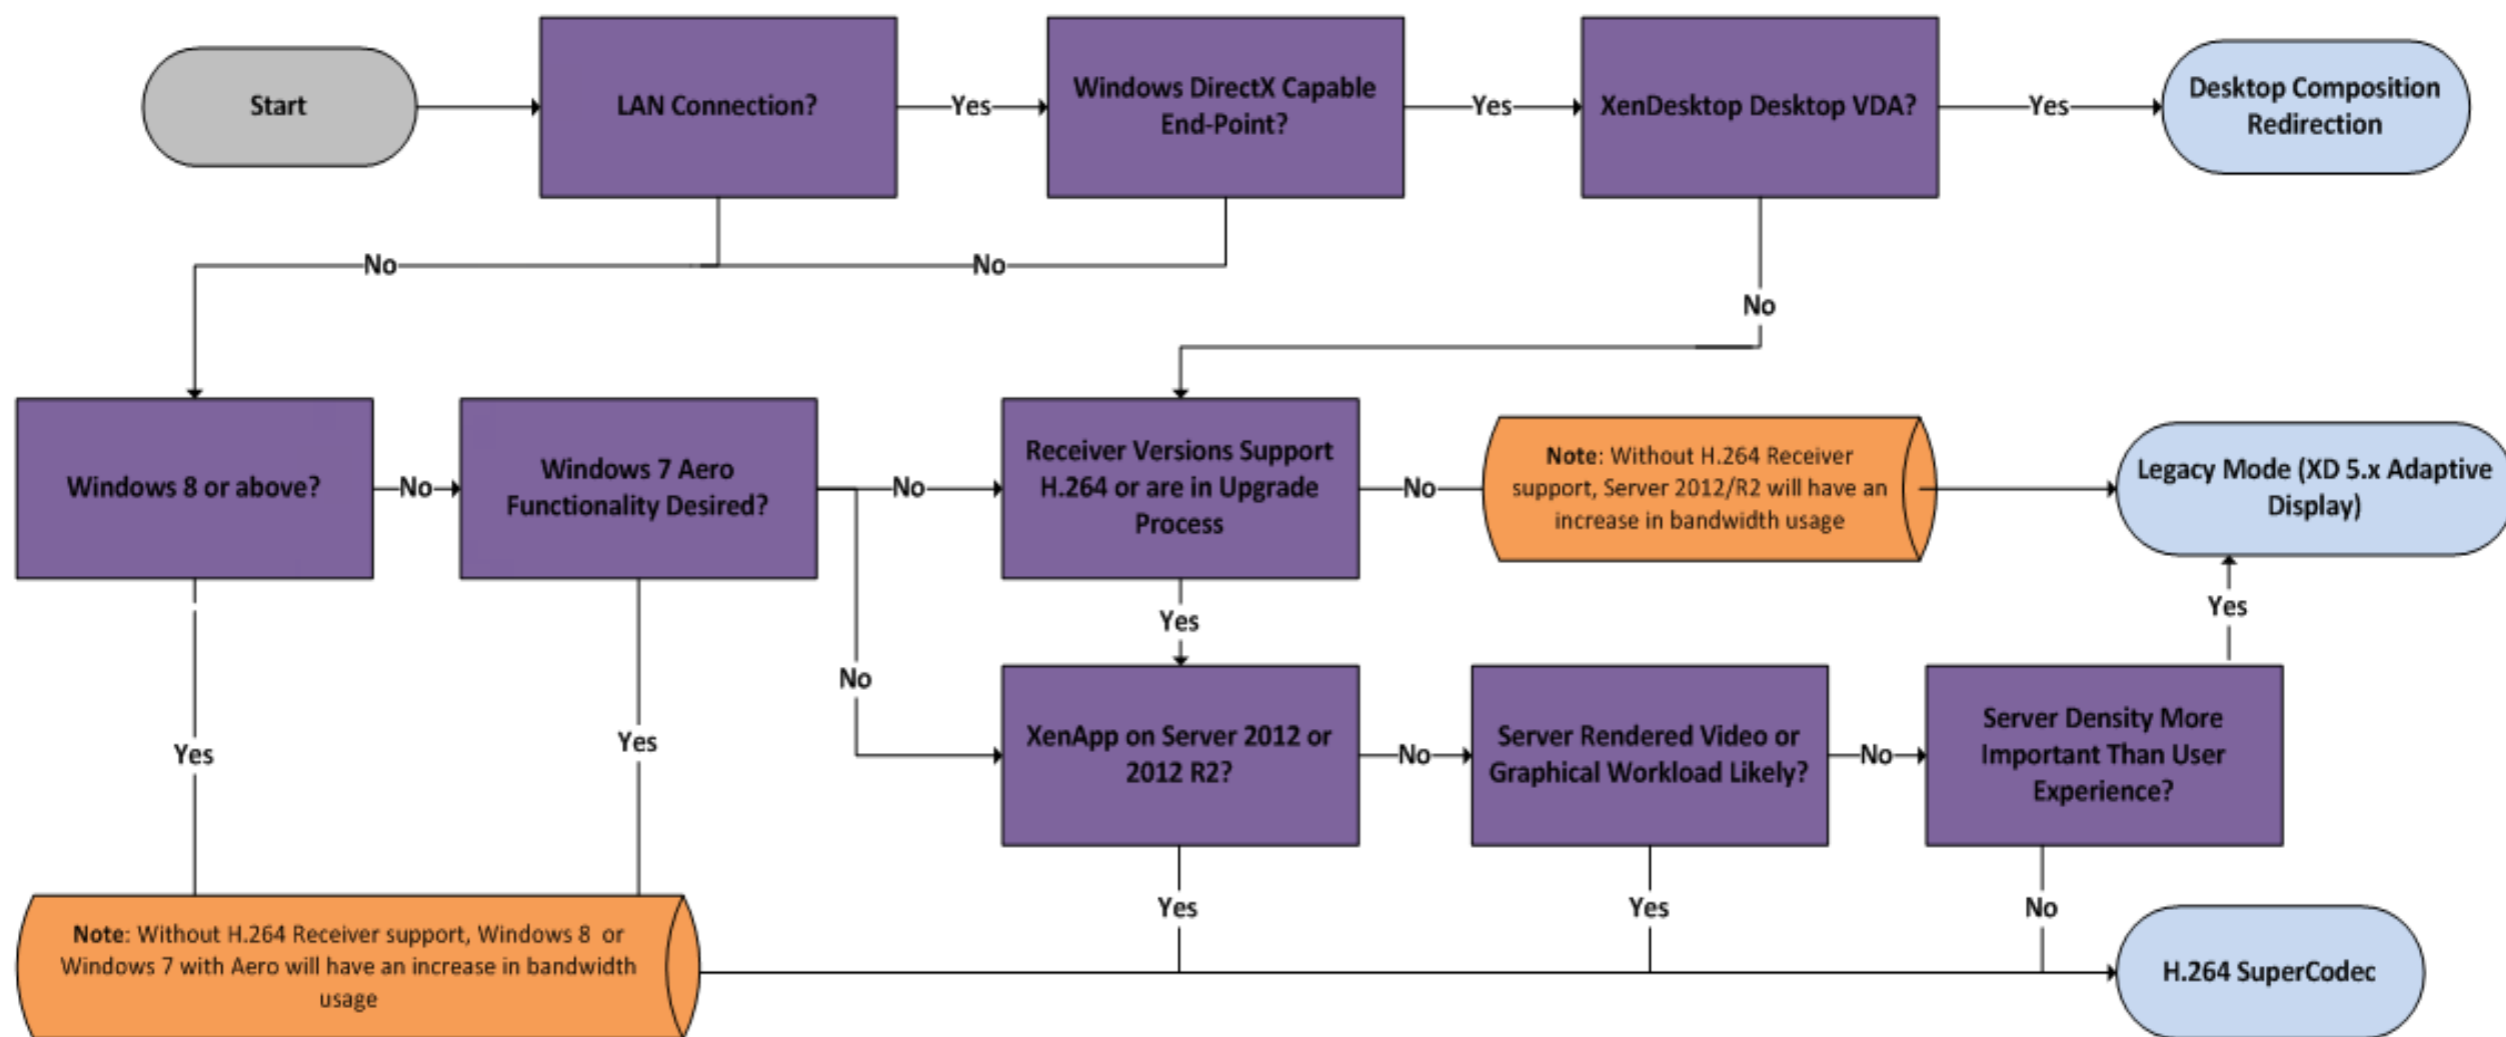
Note: Many thin client vendors create custom Receiver versions. Check with your vendor to ensure H.264 decoding is supported.

- **Recommended Use Case** – XenApp, WAN environments, mobile end points. Any endpoint which does not support DCR. Especially if accessing server rendered video or graphical content.
- **Policy** – To force using the H.264 codec for all supported end points use the following policy: ICA\Desktop UI\ “Desktop Composition Redirection” = Disabled
- **Legacy Mode (XenDesktop 5.X Adaptive Display)** – Legacy mode enables the XenDesktop 5.x settings allowing for what was known as Adaptive Display. It is only applied through Citrix Policy. Unavailable on Windows 8 and above due to Microsoft OS changes.
 - **Pros** – Long proven method which can deliver a rich user experience using a balance of CPU and bandwidth.
 - **Limitations** – Cannot deliver Aero desktops, not compatible with Windows 8 and above due to changes in the OS. Available on Server 2012 and above, but with limited benefits.
 - **Recommended Use Case** – WAN environments where end points cannot be updated to support H.264 and delivering Windows 7 or Server 2008 R2. If server scalability is more important than user experience with Windows 7 and Server 2008 R2.

- **Policy** – To enforce Legacy mode on supported operating systems, configure the following policy. ICA\Graphics\ “Legacy Graphics Mode” = Enabled. Refer to the graphics policies in [Citrix eDocs](#) for configuring legacy settings.

The following diagram provides guidance on selecting an optimal HDX Encoding method for user groups:

HDX Encoding Selection



Decision: Session Bandwidth

The bandwidth required for a connection to XenDesktop can vary greatly in the duration of a session. In a limited bandwidth situation, such as a WAN, XenDesktop can self-tune to use less bandwidth, but only up to a certain point. When bandwidth becomes too constrained, the user experience will degrade to the point of becoming unacceptable and some connections may even be dropped. The bandwidth required depends heavily on the user interaction with the desktop and the applications used. For example, multimedia and file copies will require significant bandwidth for short durations, while idle time may generate very little network traffic. These requirements can be broken down into the two categories of session bandwidths shown below.

- **Average Session Bandwidth** – The average bandwidth consumed by concurrent user sessions, this is the bandwidth of interest for general planning purposes especially in larger deployments.
- **Burst Session Bandwidth** – This is the bandwidth required to maintain the user experience during high consumption activities such as video. This bandwidth requirement is the most important to consider in small WAN deployments.

When determining the bandwidth required for a WAN connection, it is important to note which delivery method is used, discussed above. For the best user experience over a WAN connection, the H.264 enhanced SuperCodec is recommended. Once the delivery method has been chosen, it is highly recommended to test the production configuration and applications to determine bandwidth requirements.

If testing is not an available option, the following formula can be used as a rough estimate. This formula is based on both internal testing and customer analysis.

$$\text{Bandwidth (kbps)} = (200H) + (100D) + (1500X) + Z$$

Where:

H = Number of concurrent users requiring server rendered video

D = Number of concurrent users who are not watching video

X = Number of concurrent users who require 3D applications (Requires HDX 3D Pro)

Z = Additional 1000 to 2000 kbps minimum capacity to support peaks in smaller environments (<10 users)*

** The factor Z in the above formula accounts for the burst session requirement in smaller WAN networks. The burst capacity is most important to consider when deploying to a smaller WAN network. As the WAN deployment grows larger, the average bandwidth will become more accurate (to a point) and the burst less important. This formula is not meant to be a substitute for proper testing and monitoring of the environment, especially when custom applications are used. Note that this formula applies to desktops delivered using the H.264 codec described in the Encoding Method section.*

Decision: Latency

In a XenDesktop deployment, latency can have a large impact on the user experience and the perceived bandwidth that can be utilized. Although the HDX protocol can perform well at higher latencies, the user experience will begin to falter at a certain point. Server rendered video and other graphically intensive applications will be impacted most by the effects of latency and on extremely latent networks, the session may become unusable.

The HDX protocol can typically perform well up to latencies of approximately 300ms, but as mentioned this will depend on the applications being delivered and should be tested accordingly. For high network latencies, consider the global locations of desktops and applications being deployed.

Experience from the Field

A large manufacturing company deployed Hosted Shared Desktop from a central datacenter to multiple WAN sites worldwide. These desktops were used by multiple user groups, many with dual monitor requirements which require additional bandwidth. During a production pilot, where best practices policies and optimizations were employed and all end points support the H.264 Enhanced SuperCodec, the average concurrent bandwidth requirement was measured at 100kbps. The total WAN connection between the main datacenter and the remote site at which the bandwidth was monitored was 20Mbps which was more than capable of sustaining the burst requirements of the users and had an average WAN latency of 50ms.

Control Layer

The control layer is the fourth layer of the design methodology. Every major design decision made for all user groups in the upper three layers are used as a whole to help design the control components of the overall solution.

The design decisions for each user group are met by incorporating the correct control layer components, which includes access controllers, desktop controllers and infrastructure controllers. Determining capacity, configuration, topology and redundancy for the respective components creates a control environment capable of supporting the user requirements.

[Click here to provide feedback](#)

Infrastructure Controllers

Active Directory

Active Directory (AD) is required for authentication and authorization of users in a Citrix environment. The Kerberos implementation in Active Directory is used to guarantee the authenticity and confidentiality of communications with the Delivery Controllers as well as maintain time synchronization between the servers. It should be noted that Kerberos is dependent on Service Principle Names (SPNs) and DNS. SPNs are defined in AD and are used in the Kerberos authentication process.

Decision: Forest Design

Multi-forest domains allow for an environment to be separated by security boundaries within the corporate network. Examples of this can include separating by geographical location, asset isolation, or by corporate department, i.e. separating financing and human resources into separate forests or placing a newly acquired corporate merger into its own forest.

Multi-forest deployments, by default, do not have inter-domain trust relationships between the forests. An AD administrator can establish trust relationships between the multiple forests, allowing the users and computers from one forest to authenticate and access resources in another forest.

For forests that have inter-domain trusts, it is recommended that the appropriate settings be configured to allow the Delivery Controllers to communicate with both domains. When the appropriate trusts are not configured, multiple XenDesktop sites for each forest must be configured.

For more information, please refer to the following eDocs article – [Deploy in a multiple forest Active Directory environment](#).

Decision: Site Design

Proper AD site design involves ensuring that the domain controller is highly available to the Delivery Controllers and Virtual Desktop Agents (VDAs). This can be achieved by having a locally deployed domain controller or a domain controller accessible by multiple redundant WAN connections. Citrix recommends that each site has at least two domain controllers to provide high availability. Sites should be configured so that users authenticate against the most appropriate domain controller. This will usually be the physically closest domain controller. Generally, the further away a Delivery Controller is located from a domain controllers, the longer the authentication process will be. An administrator can enforce this action by configuring AD sites with subnet affinity.

For more information, please refer to the following Microsoft TechNet article – [Understanding Sites, Subnets, and Site Links](#).

Operation master roles provide a method to avoid AD update conflicts by specifying which servers perform certain AD updates. All domain controllers are treated equally in an environment. The first domain controller in the AD forest will have all the operation master roles by default. If the primary domain controller becomes overworked due to its management of all the roles, it is recommended to spread out the different roles to the other domain controllers.

For more information, please refer to the following Microsoft TechNet article – [Introduction to Flexible Single-Master Operation](#).

Decision: Organizational Unit Structure

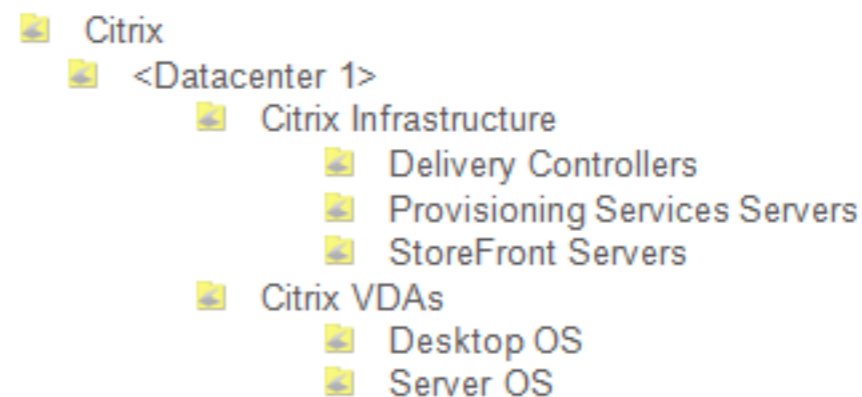
All Citrix infrastructure components for XenApp and XenDesktop as well as worker hosts should reside within their own dedicated organizational units (OUs); separating workers and controllers for management purposes. By having their own OUs, the objects inside will have greater flexibility with their management.

Organization Unit Structure

[Click here to provide feedback](#)

Citrix administrators can also be granted delegated control of the Citrix specific OUs. A basic Citrix OU Structure will have all Citrix servers placed in its own OU structure, separating infrastructure and VDA components. A sample Citrix OU structure can be seen below.

Example Citrix OU Structure



Decision: Naming Standards

A standard naming convention should be defined for all key user groups and infrastructure components, including servers and service accounts. This naming scheme should identify the role or function and location of the component or user group. This will improve the manageability and organization of the Citrix environment. Some examples of Citrix naming schemes are included in the following table.

Example Citrix Naming Schemes

Component	Naming Scheme	Example
Service Accounts	CTX_SVC_<Component>	CTX_SVC_PVS
Citrix Servers	CTX_<Component><Number>_<Datacenter>	CTX_PVS01_P
User Groups	CTX_<Group Name>_<Datacenter>	CTX_HelpDesk_P
Organizational Units (OU)	CTX_<Name of server/user role>	CTX_Admin

Note: The DTAP acronym stands for Development, Test, Acceptance, Production. It is used to identify a staging environment.

Decision: User Groups

Whenever possible, permissions should be assigned to user groups rather than individual users. Groups should be created for specific and unique roles so that permissions can be assigned to a large number of users simultaneously. Following the principles in Microsoft's Role-based Access Controls (RBAC) and Account, Global, Universal, Domain local, Permission (AGDLP), system administrators do not assign permissions directly to individual user accounts, eliminating the need to edit a large amount of resource permissions and user rights when creating, modifying, or deleting user accounts.

Permission application example:

- An application published to one group of 1,000 users requires the validation of only one object for all 1,000 users.
- The same application published to 1,000 individual user accounts requires the validation of all 1,000 objects.

For more information on Role-based Access Control, please refer to the following Microsoft article – [Role-based Access Control](#).

Decision: Service Accounts

Service accounts are special user accounts that an application or service can use to interact with the operating system. Administrators can create service accounts and manage them centrally in Active Directory. Service accounts are recommended to be used to avoid issues with permissions and account issues effecting administrators and users.

For more information, please see the Microsoft TechNet article – [Service Accounts Step-by-Step Guide](#).

Decision: Policy Control

Policy Inheritance

Group policies can be applied to users and computers at a site,

domain, or OU level. When GPOs are applied to a parent OU, the GPOs are inherited by the child container by default. The precedence of the inherited GPOs is determined by the order of processing of the GPOs. Blocking inheritance on a child container will prevent all GPOs from the parent container from being applied to the child containers. Since inherited policies may affect the usability and performance of the Citrix environment, it is recommend to document and test policies in a test environment before moving to production.

GPO Sizing

An administrator can implement a single monolithic GPO over the entire environment, or use several smaller, more specific GPOs to achieve the same goal. In general, Citrix recommends have as few GPOs as possible, merging smaller GPOs where possible. A single GPO can be linked to multiple OUs, reducing the amount of GPOs in the environment as well as promoting consistency between multiple environments. The more GPOs that need to be processed will have a negative effect on logon times. The table below shows the pros and cons for monolithic vs small GPOs.

Pros vs Cons for monolithic vs small GPOs

	Monolithic GPOs	Small GPOs
Delegation and Isolation	Monolithic GPOs will contain a lot of control settings for multiple areas of the environment. Delegation can only occur at the GPO level.	Each GPO will govern a single policy area. Delegation over this GPO can then be set to specific administrator roles.
Manageability and Complexity	Monolithic GPOs will be simpler to manage due to having all settings set within the GPO.	The more GPOs that are in the environment, the more difficult it will be to track down any delinquent policies.
Performance	The fewer the amount of GPOs that need to be processed, the faster the logon process will be for the end user.	Users will experience a degraded logon experience if there are many GPOs that need to be processed before logon can occur.

Block Inheritance

This policy setting will prevent the policy settings from higher-level OUs from being applied to a child OU. It should be noted that a higher-level OU can have their policies configured to have a No Override, preventing the Block Inheritance setting from being applied.

Loopback Policy

There are a number of user Citrix policies specific to user experience and security that should be applied to a user's Citrix session. Because user accounts can be located anywhere within a company's AD, it can become difficult to ensure that these user policies are applied. Applying Citrix policies at the domain level would affect every user logging into the environment, either through a Citrix connection or not. Applying Citrix policies to the OU which contain the Citrix XenDesktop/XenApp server or virtual desktops would not work either, as the users would have to be located within that OU. Enabling a loopback policy would allow user policy settings applied to a computer OU to get applied to users located in any OU.

For more information, please refer to the Microsoft Knowledgebase Article KB231287 – [Loopback processing of Group Policy](#).

Active Directory Policy Filtering

Policy Filtering can be used for instances when a policy needs to be applied to a small subset of users, such as Citrix Administrators. Enabling the Loopback Processing option will not work because it will apply user policy settings to everyone that logs into the system, instead of the desired group of users. AD Policy Filtering can be set up using the Security properties of a target property. For more information, please see the Microsoft TechNet article – [Security filtering using GPMC](#).

[Click here to provide feedback](#)

Database

The majority of Citrix products discussed within this document require a database. The following table outlines the usage on a per product basis:

Database usage

Product	Configuration Data	Runtime Data	Audit / Change Log Data	Monitoring Data
XenDesktop	•	•	•	•
Provisioning Services	•		Optional	
XenClient	•	•	•	

Leverage a database to store configuration related data helps to simplify administration by allowing configuration settings to be applied across multiple objects of the same type. Many of the Citrix products can utilize a new or existing database to log administrative changes performed within the infrastructure. In addition, XenDesktop leverages a shared database to store dynamic runtime information, for example connected users and available desktops. In this scenario, the database effectively becomes a message bus between controllers within the same site. Because of these dependencies the database design is a vital part of any Citrix architectural design.

Decision: Edition

There are five editions of Microsoft SQL Server 2012: Express, Web, Standard, Business Intelligence, and Enterprise. The table (shown below) outlines the features of each edition, which are most applicable to Citrix databases.

Based on the capabilities of the various SQL Server versions available, Citrix typically recommends using the Standard edition for hosting Citrix databases in production environments. The Standard edition provides an adequate amount of features to meet the needs of most environments. For more information on the databases supported with Citrix products please refer to the Citrix Database Support Matrix. Different versions of Citrix products support different versions of the SQL server; therefore it is important to check the support matrix to ensure the version of SQL server used is compatible with the Citrix product being deployed.

Decision: Database and Transaction Log Sizing

When sizing a SQL database, two aspects are important:

- **Database file** – Contains the data and objects such as tables, indexes, stored procedures and views stored in the database.
- **Transaction log file** – Contains a record of all transactions and database modifications made by each transaction. The transaction log is a critical component of the database and, if there is a system failure, the transaction log might be required to bring the database back to a consistent state. The usage of the transaction log varies depending on which database recovery model is used:
 - *Simple recovery* – No log backups required. Log space is automatically reclaimed, to keep space requirements small, essentially eliminating the need to manage the transaction log space. Changes to the database since the most recent backup are unprotected. In the event of a disaster, those changes must be redone.
 - *Full recovery* – Requires log backups. No work is lost due to a lost or damaged database data file. Data of any arbitrary point in time can be recovered (for example, prior to application or user error). Full recovery is required for database mirroring.
 - *Bulk-logged* – This model is an adjunct of the full recovery

SQL 2012 Edition Feature Comparison

Feature	Enterprise	Busines Intelligence	Standard	Web	Express
Scalability and Performance					
Compute capacity	OS maximum	4 Sockets or 16 cores	4 Sockets or 16 cores	4 Sockets or 16 cores	1 Sockets or 4 cores
Maximum memory utilized	OS maximum	128 GB	128 GB	64 GB	1GB
Maximum database size	524 PB	524 PB	524 PB	524 PB	10GB
High Availability					
AlwaysOn failover cluster instances	Yes (Node support: OS maximum)	Yes (2 nodes)	Yes (2 nodes)	-	-
AlwaysOn availability groups	Yes	-	-	-	-
Database mirroring	Yes	Yes (Safety Full Only)	Yes (Safety Full Only)	Witness Only	Witness Only

model that permits high-performance bulk copy operations. It is typically not used for Citrix databases.

For further information, please refer to the Microsoft Developer Network – [SQL Server Recovery Models](#).

In order to estimate storage requirements, it is important to understand the disk space consumption for common database entries. This section outlines the storage requirements on a per product basis and provides sizing calculations. For more information, please refer to Citrix article: CTX139508 – [XenDesktop 7.x Database Sizing](#).

XenDesktop General

XenApp 7.x and XenDesktop 7.x use three distinct databases:

- **Site Configuration database** – Contains static configuration and dynamic runtime data
- **Monitoring database** – Contains monitoring data which is accessible via Director
- **Configuration logging database** – Contains a record for each administrative change performed within the site (accessible via Studio)

Site Database

Since the database of a XenApp or XenDesktop site contains static configuration data and dynamic runtime data, the size of the database file depends not only on the physical size of the environment but also user patterns. The following factors all impact the size of the database file:

- The number of connected sessions
- The number of configured and registered VDAs
- The number of transactions occurring during logon
- VDA heartbeat transaction

The size of the Site Database is based on the number of VDAs and

[Click here to provide feedback](#)

active sessions. The following table shows the typical maximum database size Citrix observed when scale testing XenApp and XenDesktop with a sample number of users, applications, and desktop delivery methods.

XenDesktop Site DB sample size calculations

Users	Applications	Desktop Types	Expected Maximum Size (MB)
1,000	50	Hosted Shared	30
10,000	100	Hosted Shared	60
100,000	200	Hosted Shared	330
1,000	N/A	VDI	30
10,000	N/A	VDI	115
40,000	N/A	VDI	390

Note: This sizing information is a guide only. Actual database sizes may differ slightly by deployment due to how databases are maintained.

Determining the size of the transaction log for the Site database is difficult due to factors that can influence the log including:

- The SQL Database recovery model
- Launch rate at peak times
- The number of desktops being delivered

During XenDesktop scalability testing, Citrix observed the transaction log growth rate at 3.5 MB an hour when the system is idle, and a per user per day growth rate of ~32 KB. In a large environment, transaction log usage requires careful management and a regular backup, to prevent excessive growth. This can be achieved by means of scheduled jobs or maintenance plans.

Monitoring Database

Of the three databases, the Monitoring database is expected to be the largest since it contains historical information for the site. Its size is dependent on many factors, including:

- Number of Users
- Number of sessions and connections
- Number of VDI or HSD workers
- Retention period configuration – Platinum customers can keep data for over a year (default 90 days). Non-platinum customers can keep data for up to 7 days (default 7 days).
- Number of transaction per second. Monitoring service tends to execute updates in batches. It is rare to have the number of transactions per second go above 20.
- Background transaction caused by regular consolidation calls from the Monitoring service.
- Overnight processing carried out to remove data outside the configured retention period.

The following table shows the estimated size of the Monitoring database over a period of time under different scenarios. This data is an estimate based on data seen within scale testing XenApp and XenDesktop (assuming a 5 day working week).

Monitoring DB Size Estimations

Estimates with 1 connection and 1 session per user with a 5 day work week					
Users	Type	1 week (MB)	1 month (MB)	3 months (MB)	1 year (MB)
1,000	HSD	20	70	230	900
10,000	HSD	160	600	1,950	7,700
100,000	HSD	1,500	5,900	19,000	76,000
1,000	VDI	15	55	170	670
10,000	VDI	120	440	1,400	5,500
40,000	VDI	464	1,700	5,400	21,500
Estimates with 2 connections and 1 session per user with a 5 day work week					
Users	Type	1 week (MB)	1 month (MB)	3 months (MB)	1 year (MB)
1,000	HSD	30	100	330	1,300
10,000	HSD	240	925	3,000	12,000
100,000	HSD	2,400	9,200	30,000	119,000
1,000	VDI	25	85	280	1,100
10,000	VDI	200	750	2,500	9,800
40,000	VDI	800	3,000	9,700	38,600

Note: The 100,000 HSD tests are based on a test environment consisting of:

- 2 Delivery Controllers
- 43 HSD VDA workers
- 3 SQL servers, configured with databases held within one Always On Availability Group.

For more information, please see the Citrix Support article: CTX139508 – [XenDesktop 7.x Database Sizing](#).

The size of the transaction log for the Monitoring Database is very hard to estimate, but XenApp and XenDesktop scalability testing showed a growth rate of about 30.5 MB an hour when the system is idle, and a per user per day growth rate of ~9 KB.

Configuration Logging Database

The Configuration Logging Database is typically the smallest of the three databases. Its size and the size of the related transaction log depends on the daily administrative activities initiated from Studio, Director or PowerShell scripts, therefore its size is difficult to estimate. The more configuration changes are performed, the larger the database will grow. Some factors that can affect the size of the database include:

- The number of actions performed in Studio, Director and PowerShell.
- Minimal transactions which occur on the database when no configuration changes are taking place.
- The transaction rate during updates. Updates are batched whenever possible.
- Data manually removed from the database. Data within the Configuration Logging Database is not subject to any retention policy, therefore it is not removed unless done so manually by an administrator.

- Activities that have an impact on sessions or users, for example, session logoff and reset.
- The mechanism used for deploying desktops.

In XenApp environments not using MCS, the database size tends to fall between 30 and 40MB. For MCS environments, database size can easily exceed 200MB due to the logging of all VM build data.

Temporary Database

In addition to the Site, Monitoring, and Configuration Logging databases, a system-wide temporary database (tempdb) is provided by SQL Server. This temporary database is used to store Read-Committed Snapshot Isolation data. XenApp 7.x and XenDesktop 7.x uses this SQL Server feature to reduce lock contention on the XenApp and XenDesktop databases. Citrix recommends that all XenApp 7.x and XenDesktop 7.x databases use Read-Committed Snapshot Isolation. For more information please see CTX137161 – [How to Enable Read-Committed Snapshot in XenDesktop](#).

The size of the tempdb database will depend on the number of active transactions, but in general it is not expected to grow more than a few MBs. The performance of the tempdb database does not impact the performance of XenApp and XenDesktop brokering, as any transactions that generate new data require tempdb space. XenApp and XenDesktop tend to have short-lived transactions, which help keep the size of the tempdb small.

The tempdb is also used when queries generate large intermediate result sets. Guidance and sizing the tempdb can be found on the Microsoft TechNet article – [Optimizing tempdb Performance](#).

Provisioning Services

The Provisioning Services farm database contains static configuration and configuration logging (audit trail) data. The record size requirements can be used to help size the database:

Provisioning Services Farm DB sample size calculations

Configuration Item	DB Space Required (KB)	Example	
		# of Items	Total (KB)
Base farm configuration	112	-	112
User group w/ farm access	50	10	250
Site	4	5	20
Device collection	10	50	500
Farm view	4	10	40
Farm view to device relationship	5	1	5,000
Site View	4	5	20
Site view to device relationship	5	1	5,000
Device	2	5,000	10,000
Device bootstrap	10	-	-
Device to disk relationship	35	1	175,000
Device printer relationship	1	-	-
Device personality data	1	-	-
Device status (when booted)	1	5,000	5,000
Device custom property	2	-	-
vDisk	1	20	20
vDisk version	3	5	300
Disk locator	10	1	200
Disk locator custom property	2	-	-
Server	5	10	50
Server IP	2	1	20
Server status (when booted)	1	20	20
Server custom property	2	-	-
vDisk store	8	5	40
vDisk store to server relationship	4	1	40
Connection to XenServer (VirtualHostingPool)	4	-	-
vDisk update task	10	10	100
Administrative change (auditing enabled)	1	10,000	10,000
Total overall			211,732KB (~212MB)

During the PVS farm setup a database with an initial file size of 20MB is created. Due to the nature of the data in the PVS farm database the transaction log is not expected to grow very quickly, unless a large amount of configuration is performed.

In contrast to XenApp, which also offers the ability to track administrative changes, the related information is not written to a dedicated database but directly to the Provisioning Services farm

database. In order to limit the size of the Provisioning Services database it is recommended to archive the audit trail data on a regular schedule.

XenClient

The XenClient database contains both static configuration and active computer data. The static data includes endpoint computers, policies, virtual machine images and their assignments. The active computer data includes last check-in time and disk usage. This data is only refreshed and does not progressively build up. However, the XenClient event log data will progressively build up as the environment grows. This data includes all events that occur such as image creation, image updates, remote help desk logons and scheduled tasks. Initially, the XenClient database will be less than 10MB in size including the transaction logs (full) although it will grow as additional computers, images and policies are added.

Decision: Database Location

By default, the Configuration Logging and Monitoring databases are located within the Site Configuration database. Citrix recommends changing the location of these secondary databases as soon as the configuration of the site has been completed, in order to simplify sizing, maintenance and monitoring. All three databases can be hosted on the same server or on different servers. An ideal configuration would be to host the Monitoring database on a different server from the Site Configuration and Configuration Logging databases since it records more data, changes occur more frequently and the data is not considered to be as critical as the other databases. For more information, please refer to Citrix eDocs – [Change secondary database locations](#).

Note: The location of the Configuration Logging database cannot be changed when mandatory logging is enabled.

Decision: High-Availability

The following table highlights the impact to XenApp, XenDesktop,

[Click here to provide feedback](#)

XenClient and Provisioning Services when there is a database outage:

Impact of a database outage

Component	Impact of Database Outage
Site configuration database	<p>Users will be unable to connect or reconnect to a virtual desktop.</p> <p><i>Note: Connection leasing in XenApp and XenDesktop 7.6 allows users with Hosted Shared Desktops, On-Demand Applications, and Assigned VDI Desktops to reconnect to their most recently used applications and desktops even when the site database is unavailable.</i></p> <p>Administrators are unable to use Studio or Director. Any users with existing connections will be unaffected</p>
XenDesktop monitoring database	<p>Director will not display any historical data and Studio cannot be started. Brokering of incoming user requests and existing user sessions will not be affected.</p>
XenDesktop configuration logging database	<p>If allow changes when the database is disconnected has been enabled within XenApp and XenDesktop logging preferences, an outage of the configuration logging database will have no impact (other than configuration changes not being logged). Otherwise, administrators will be unable to make any changes to the XenApp and XenDesktop site configuration. Users are not impacted.</p>
Provisioning Services farm database	<p>When offline database support is enabled and the database becomes unavailable, the stream process uses a local copy of the database to retrieve information about the provisioning server and the target devices supported by the server. This allows provisioning servers and the target devices to remain operational. However, when the database is offline, the console and the management functions listed below become unavailable:</p> <ul style="list-style-type: none"> • AutoAdd target devices • vDisk creation and updates • Active Directory password changes • Stream process startup • Image update service • PowerShell and MCLI based management <p>If offline database support has not been enabled, all management functions become unavailable and the boot and failover of target devices will fail.</p>
XenClient database	<p>Administrators will be unable to access the Synchronizer webpage to manage the environment, new computers will be unable to register and existing computers will be unable to receive image updates or policies. Existing users will still be able to function normally with access to their virtual machines being uninterrupted.</p>

Note: Please review HA options for 3rd party databases (for example, App-V, SCVMM or vCenter) with the respective software vendor.

In addition to the built-in Citrix database redundancy options, Microsoft SQL Server, as well as the underlying hypervisor (in virtual environments), offer a number of high availability features. These enable administrators to ensure single server outages will have a minimal impact (if any) on the Citrix infrastructure. The following the SQL / hypervisor high availability features are available:

- **VM-level HA** – This high availability option is available for virtual SQL servers only, which need to be marked for High Availability at the hypervisor layer. In case of an unexpected shutdown of the virtual machine or the underlying hypervisor host, the hypervisor will try to restart the VM immediately on a different host. While VM-level HA can minimize downtimes in power-outage scenarios, it cannot protect from operating system level corruption. This solution is less expensive than mirroring or clustering because it uses a built-in hypervisor feature and requires shared storage. However, the automatic failover process is slower, as it can take time detect an outage and start the virtual SQL server on another host. This may interrupt the service to users.
- **Mirroring** – Database mirroring increases database availability with almost instantaneous failover. Database mirroring can be used to maintain a single standby or mirror database, for a corresponding principal or production database. Database mirroring runs with either synchronous operation in high-safety mode, or asynchronous operation in high- performance mode. In high-safety mode with automatic failover (recommended for XenDesktop) a third server instance, known as a witness, is required, which enables the mirror server to act as a hot standby server. Failover from the principal database to the mirror database happens automatically and is typically completed

within a few seconds. It is a good practice to enable VM-level HA (or a similar automatic restart functionality) for at least the witness to ensure SQL service availability in case of a multi-server outage.







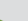





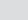
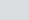
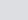
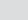



Note: Microsoft is planning to remove mirroring as a high availability option in a future release of SQL Server and is discouraging its use in new network development. Please refer to the Microsoft article – [Database Mirroring \(SQL Server\)](#) for more information.

- **AlwaysOn Failover Cluster Instances** – Failover clustering provides high-availability support for an entire instance of Microsoft SQL Server. A failover cluster is a combination of two or more nodes, or servers, using shared storage. A Microsoft SQL Server AlwaysOn Failover Cluster Instance, introduced in SQL Server 2012, appears on the network as a single computer, but has functionality that provides failover from one node to another if the current node becomes unavailable. The transition from one node to the other node is seamless for the clients connected to the cluster. AlwaysOn Failover Cluster Instances require a Windows Server Failover Clustering (WSFC) resource group. The number of nodes supported in the WSFC resource group will depend on the SQL Server edition. (Please refer to the table in the [Decision: Edition](#) earlier in this chapter.) For more information please refer to MSDN – [AlwaysOn Failover Cluster Instances \(SQL Server\)](#).
- **AlwaysOn Availability Groups** – AlwaysOn Availability Groups is an enterprise-level high-availability and disaster recovery solution introduced in Microsoft SQL Server 2012, which enables administrators to maximize availability for one or more user databases. AlwaysOn Availability Groups require that the Microsoft SQL Server instances reside on Windows Server failover clustering (WSFC) nodes. Similar to AlwaysOn Failover Cluster Instances, a single virtual IP / network name is exposed to the database users. In contrast to AlwaysOn Failover Cluster

Instances, shared storage is not required since AlwaysOn Availability Groups uses replicas copied locally to each node in the group. Both synchronous and asynchronous replication to one or more secondary servers is supported. As opposed to mirroring or clustering secondary servers can be actively used for processing incoming read-only requests, backups or integrity checks. This feature can be used to offload user resource enumeration requests to a secondary SQL server in XenApp or XenDesktop environments to essentially scale-out a SQL server infrastructure. Since the data on active secondary servers can lag multiple seconds behind the primary server, the read-only routing feature cannot be used for other XenApp or XenDesktop database requests at this point in time. For more information, please refer to MSDN – [AlwaysOn Availability Groups \(SQL Server\)](#).

The following table outlines the recommended high availability features for Citrix databases.

Recommended SQL high availability options

Component	VM-Level HA	Mirroring	Failover Clustering	AlwaysOn Availability Groups
XenDesktop Site Database				 ²
XenDesktop configuration Logging Database		 ¹		 ²
XenDesktop Monitoring Database				 ²
Provisioning Services Farm Database				 ³
XenClient Database				

 Recommended  Viable  Not Supported  Recommended for Test Environments Only

¹ If “Allow changes when the database is disconnected” has been enabled no redundancy is required, otherwise mirroring is recommended

² Mirroring is preferred due to the lower SQL license requirements

³ PVS 7.6 and above supports AlwaysOn for SQL Server 2012 and 2014

Decision: Connection Leasing

Connection leasing is a new XenApp and XenDesktop 7.6 feature that allows Hosted Shared, On-Demand Apps and Assigned VDI users to connect and reconnect to their most recently used applications and desktops even, when the site database is unavailable. Connection Leasing is not available for users with a Pooled VDI desktop.

The lease information along with the application, desktop, icon, and worker information is stored on the controller’s local disk and synchronized between controllers in the site. If the site database becomes unavailable, the controllers enter a “leased connection mode” and replay cached operations from an XML file on the local disk to connect or reconnect users to a recently used application or desktop.

Administrators familiar with the local host cache in XenApp 6.5 and earlier should understand the similarities and differences with connection leasing because it can have an impact on the design and scalability of the XenApp and XenDesktop 7.6 solution. In XenApp 6.5 and earlier, the IMA service is responsible for synchronizing the local host cache with the data store. In XenApp and XenDesktop 7.6, the FMA service caches the brokering operations (leases) to an XML file containing the address of the VDA, application path, and other details required for the session to launch. The FMA also caches dynamic information such as user sessions, VDA registrations, and load. These files are uploaded to the SQL database and synchronized between all controllers in the site. The controllers will download the files on a regular basis so that any controller in the site can connect a user to their session.

Each controller needs additional disk space for the cached lease files. At a minimum, 4KB is required for each lease file. Each resource entry in the enumeration lease will take anywhere from 200 bytes to a few KBs depending on the number of entries and resources published. Citrix testing has shown that 200,000 leased

connections for server hosted applications and desktops required approximately 3GB of disk space. 40,000 leased connections for assigned desktops required approximately 156MB of disk space.

By default, connection leases have an expiration period of two weeks. Applications and desktops must have been launched within the two last weeks to still be accessible when the database is unavailable. The expiration period is configurable using PowerShell cmdlets or editing the registry and can be set from 0 minutes to several years. Setting the expiration period too short will prevent users from connecting to their virtual desktops and applications in the event of an outage. Setting the expiration period too long will increase storage requirements on the controllers.

By default, connection leasing affects the entire site, however, leases can be revoked for specific users, which prevents them from accessing any applications or desktops when the site database is unavailable.

For more information on connection leasing considerations and configuration, please refer to eDocs – [Connection leasing](#).

Decision: Microsoft SQL Server Sizing

The SQL server must be sized correctly to ensure the performance and stability of an environment. Since every Citrix product uses SQL server in a different way, no generic all-encompassing sizing recommendations can be provided. Instead, per-product SQL server sizing recommendations are provided below.

XenApp and XenDesktop

XenApp and XenDesktop Brokers use the database as a message bus for broker communications, storing configuration data and storing monitoring and configuration log data. The databases are constantly in use and the performance impact on the SQL server can be considered as high.

Based on results from Citrix internal scalability testing the following SQL server specification for a server hosting all XenDesktop

databases are recommended:

- 2 Cores / 4 GB RAM for environments up to 5,000 users
- 4 Cores / 8 GB RAM for environments up to 15,000 users
- 8 Cores / 16 GB RAM for environments with 15,000+ users

The database files and transaction logs should be hosted on separate hard disk subsystems in order to cope with a high number of transactions. For example, registering 20,000 virtual desktops during a 15 minute boot storm causes ~500 transactions / second and 20,000 users logging on during a 30 minute logon storm causes ~800 transactions / second on the XenDesktop Site Database

Provisioning Services

In addition to static configuration data provisioning servers store runtime and auditing information in the database. Depending on the boot and management pattern, the performance impact of the database can be considered as low to medium.

Based on this categorization, a SQL server specification of 4 Cores and 4 GB RAM is recommended as a good starting point. The SQL server should be carefully monitored during the testing and pilot phase in order to determine the optimal configuration of the SQL server.

XenClient

Similar to provisioning servers, XenClient synchronizers store static configuration data in addition to runtime and auditing / event information in the database. The performance impact of the database can be considered as low to medium, depending on the client synchronization and management pattern.

Based on this categorization, a SQL server specification of 4 Cores and 4 GB RAM is recommended as a good starting point. The SQL server should be carefully monitored during the testing and pilot phase in order to determine the optimal configuration of the SQL

server.

Note: Please review SQL server sizing for 3rd party databases (for example, App-V, SCVMM, vCenter) with the respective software vendor.

Decision: SQL Service Accounts

Database access service accounts should be created for every Citrix product, as outlined within the Active Directory chapter. For reference purposes, the minimum SQL server and database access rights required for initial setup, maintenance and runtime are detailed below:

- **XenApp and XenDesktop**
 - **Setup and maintenance** – The administrator responsible for initial database creation, adding controllers, removing controllers or applying database schema updates requires dbcreator and securityadmin server rights as well as the db_owner database right.
 - **Runtime** – During initial database setup, XenApp and XenDesktop configure specific database security roles, which limit the services to read, write, and execute and should not be modified.
- **Provisioning Services**
 - **Setup and maintenance** – For initial database creation, the administrator who is performing this task requires dbcreator and securityadmin server rights.
 - **Runtime** – During normal runtime, the Provisioning Services service account requires read (db_datareader), write (db_datawriter) and execute permissions only. Further information can be found in eDocs – [Installing and Configuring Provisioning Services](#).
- **XenClient**
 - **Setup, maintenance and runtime** – The service account used

for connecting to the XenClient database requires db_owner rights for all phases. Further information can be found in CTX138715 – [XenClient Synchronizer Installation Guide Version 5.1](#).

Decision: Database Creation

The XenApp or XenDesktop site database can be created automatically using the Site creation wizard or manually using SQL database scripts. If the database is created automatically, then the account used must have SQL server permissions as specified in the previous section [Decision: SQL Service Accounts](#). If the account does not have these permissions then the Site creation wizard will prompt for the SQL server user credentials.

If it is decided upon to create the database manually, then the Site creation wizard will generate two SQL scripts that can be imported into the database. One script sets up the database and the other is used in a mirrored environment. After the scripts are generated, they are run on the SQL server to create the site database.

Manual database creation is necessary in situations where the Citrix administrator does not have the required SQL server permissions to create the site database. Some organizations will have a separate team of SQL database administrators responsible for the creation and management of production SQL databases, in which case the scripts should be given to them to run on the SQL servers. Once the scripts are run and the database is created, the Citrix administrator can complete the site configuration. For more information on the database attributes to set when doing a manual configuration please refer the Citrix eDocs article – [Create a Site](#).

Citrix Licensing

Citrix offers customers the flexibility of multiple licensing models that align with common usage scenarios. The different licensing models vary based on the Citrix product used, but can include

per user/device and per concurrent user. Several Citrix products use the license server, while other products require a license to be installed on the product itself.

License locations

Product	License Location
XenDesktop	Citrix License Server
XenApp	Citrix License Server
Provisioning Services	Citrix License Server
XenServer	Citrix License Server
XenClient	On the product
NetScaler	On the product
NetScaler Gateway	On the product

For more information on XenDesktop 7.x licensing, please refer to CTX128013 – [XenDesktop Licensing](#).

For more information on Microsoft Licensing, please refer to the Microsoft document – [Licensing Microsoft's Virtual Infrastructure Technology](#).

Decision: License Type

To determine the number of licenses needed, customers must determine the licensing model and the total number of users or devices that will access the Citrix environment. Licenses can be checked out to a user or device, depending on the licensing model for each product.

Note: A Supplemental Grace Period was introduced with XenDesktop 7.6. If all licenses are consumed, grace licenses will be granted for a period of up to 15 days, allowing the administrator time to acquire additional licenses. The administrator will be alerted that licenses are over consumed in Director via the License Policy Engine.

User/Device Licenses

With user/device licensing, the license server can assign the

same license to a user or a device. When assigned to a user, the license allows access from an unlimited number of devices. When assigned to a device, the license allows access from the device by an unlimited number of users. The license server determines how to minimize license consumption based on the number of users and devices connected. For more information on this process, please refer to the Citrix eDocs page – [Types of licenses](#).

Concurrent Licenses

Licenses can also be concurrent, not tied to a specific user or device. As a user launches a product, a license is checked out to the specific computer or device. When the user logs off or disconnects from the session, the license is checked back in, becoming available for another user to consume.

Additional licenses can be consumed in the following scenarios:

- Multiple sessions at different computers will consume multiple licenses. When a user launches a product, a license is checked out until the user closes the session at that computer/device (the license is checked back in at this point). For example, if a user launches a session from one computer and then launches another from a different computer without closing the first session, two licenses are checked out.
- License servers do not communicate usage information with each other. Therefore, multiple licenses can be checked out when multiple license servers are used. This can be avoided by ensuring that all product servers in an environment are pointed to the same license server.
- Multiple licenses are consumed when a single device connects to multiple product servers configured with different editions. For example, if a user connects to an application published on a server using the Advance edition, and then uses the same client to connect to an application published on a different server running the Enterprise edition, two licenses are consumed.

For more information on licensing models and types, please refer to the following Citrix Product Pages – [Licensing Models & Types of Licenses](#).

Decision: Version

New license servers are backwards compatible and will work with older products and license files; however, new products often require the newest license server to check out licenses correctly. You can find the latest version from the [Citrix Downloads site](#).

Decision: Sizing

Internal scalability testing has shown that a single virtual license server with two cores and 2GB of RAM can issue approximately 170 licenses per second or 306,000 licenses per half hour. If necessary, the specification of the license server can be scaled out to support a higher number of license requests per second.

Note: Citrix recommends engaging in proper scalability testing to ensure that the license server is capable of meeting the demands of the environment.

Decision: High Availability

For a typical environment, a single license server is sufficient. Should the license server become unavailable, dependent Citrix products will enter a 30-day grace period, which provides more than enough time to resolve connectivity issues and/or restore or rebuild the license server.

Note: If the license server and the Citrix product do not communicate within 2 heartbeats (5-10 min), the Citrix product will enter a grace period and will allow connections for up to 30 days. Once communication with the license server is re-established, the license server will reconcile the temporary and actual licenses.

Note: A CNAME record in DNS is a convenient way to reference the license server. Using CNAMEs allows the license server name to be changed without updating the Citrix products.

[Click here to provide feedback](#)

If additional redundancy is required, Citrix supports the following high availability solutions for the license server.

- **Windows Clustering** - Cluster servers are groups of computers that work together in order to increase availability. Clustering allows the license server role to automatically failover in the event of a failure. For more information on clustering, please see the Citrix eDocs article - [Clustered License Servers](#).
- **Duplication of license server** - Create a VM level backup of the license server. This backup should not be stored on the same host as the license server. Instead, it should be stored in a safe location, such as a highly available storage solution, or backed up to tape or disk. The duplicate server is not active and will remain on standby until the need arises to restore the active license server. Should the license server be restored using this backup, any new licenses must be re-downloaded to the server.

For more information, please refer to Citrix eDocs - [Licensing Architecture Overview](#).

Note: For XenApp 6.5 and older environments, the MPS-WSXICA_MPS-WSXICA.ini file should be redirected to a file share, as described within CTX131202 – Provisioned XenApp servers stop accepting connections if they are restarted when the license when the license server is unavailable. This is not an issue for XenApp 7.x environments since the controller handles the checking in and out of licenses.

Each method allows an administrator to exchange a single license server for another without an interruption in service; assuming that the change occurs during the grace period and that the following limitations are considered.

- License files will reference the server specified during the allocation process. This means that the license files can only be used on a server with the same binding information (Hostname) as the server that was previously specified.

- Two Windows-based, domain joined license servers cannot share the same name and be active in the environment at the same time.
- Because license servers do not communicate with each other, any additional licenses must be placed on both the active and backup license server.

Decision: Optimization

License server performance can be optimized by tuning the number of “receive” and “processing” threads. If the thread count is set too low, requests will be queued until a thread becomes available. Conversely, if the thread count is set too high, the license server will become overloaded.

The optimal values are dependent on the server hardware, site configuration, and license request volume. Citrix recommends testing and evaluating different values to determine the proper configuration. Setting the maximum number of processing threads to 30 and the maximum number of receiving threads to 15 is a good starting point for large scale deployments.

This optimization will improve the Citrix License Server ‘s ability to provide licenses by increasing its ability to receive and process license requests.

For more information, please refer to the Citrix eDocs – [Improving Performance by Specifying Thread Use](#).

Resource Controllers

The resource controller sub-layer is responsible for providing the infrastructure components to support the resource layer requirements for each user group. Desktop controllers often include some combination of XenDesktop or XenApp Delivery Controllers, Provisioning Services servers and XenClient Synchronizers.

[Click here to provide feedback](#)

XenDesktop and XenApp Delivery Controller

A Citrix site groups desktops and applications together to form a single architectural and management entity. All persistent and dynamic data for the site, including site configuration, desktop assignments, and session state, is stored in a central site database.

Site-wide functions are spread equally across all Delivery Controllers within a site. While it is possible to assign certain functions to specific controllers, it is not recommended as XenDesktop is self-optimizing and manual configurations can interfere with the automatic load balancing and failover mechanisms of XenDesktop.

Decision: Number of Sites

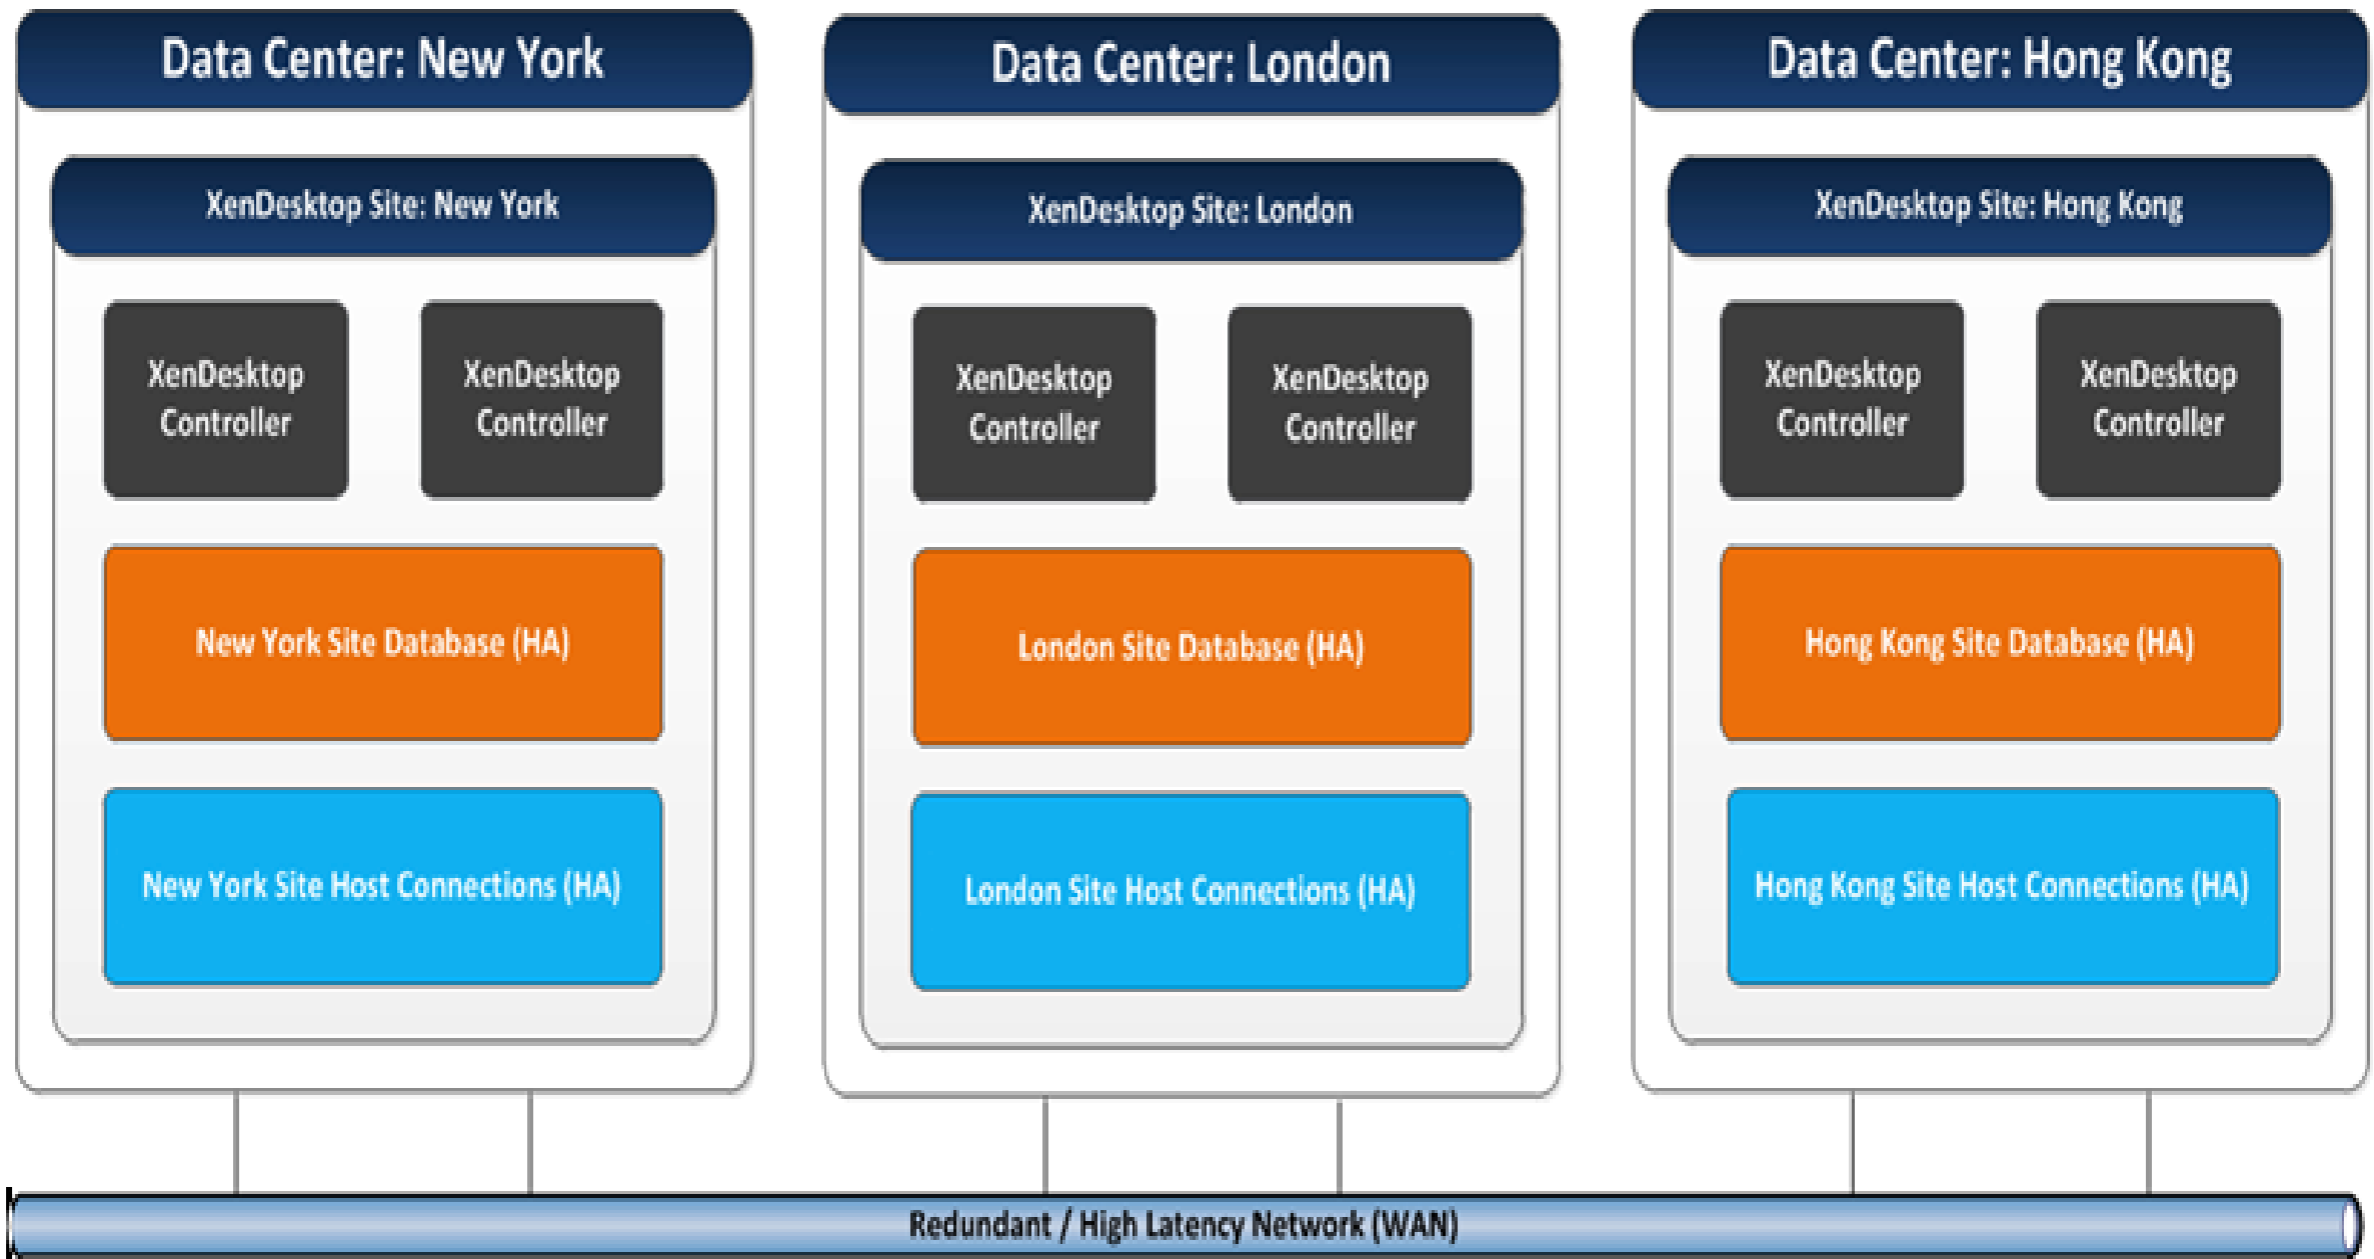
There are several factors that must be considered when determining the number of XenDesktop sites required.

- **Network** – XenDesktop Controllers constantly access the site database and regularly communicate with desktops and servers hosting applications. XenDesktop may exhibit unexpected behavior if significant levels of latency exist between the XenDesktop controllers, SQL database, desktop operating systems and server operating systems or when site communications are interrupted. Therefore, all site components (controllers, desktops, servers, virtualization hosts and database servers), should be connected by a redundant and high-speed network.

On the next two pages, two scenarios demonstrate how the network can affect XenDesktop site design.

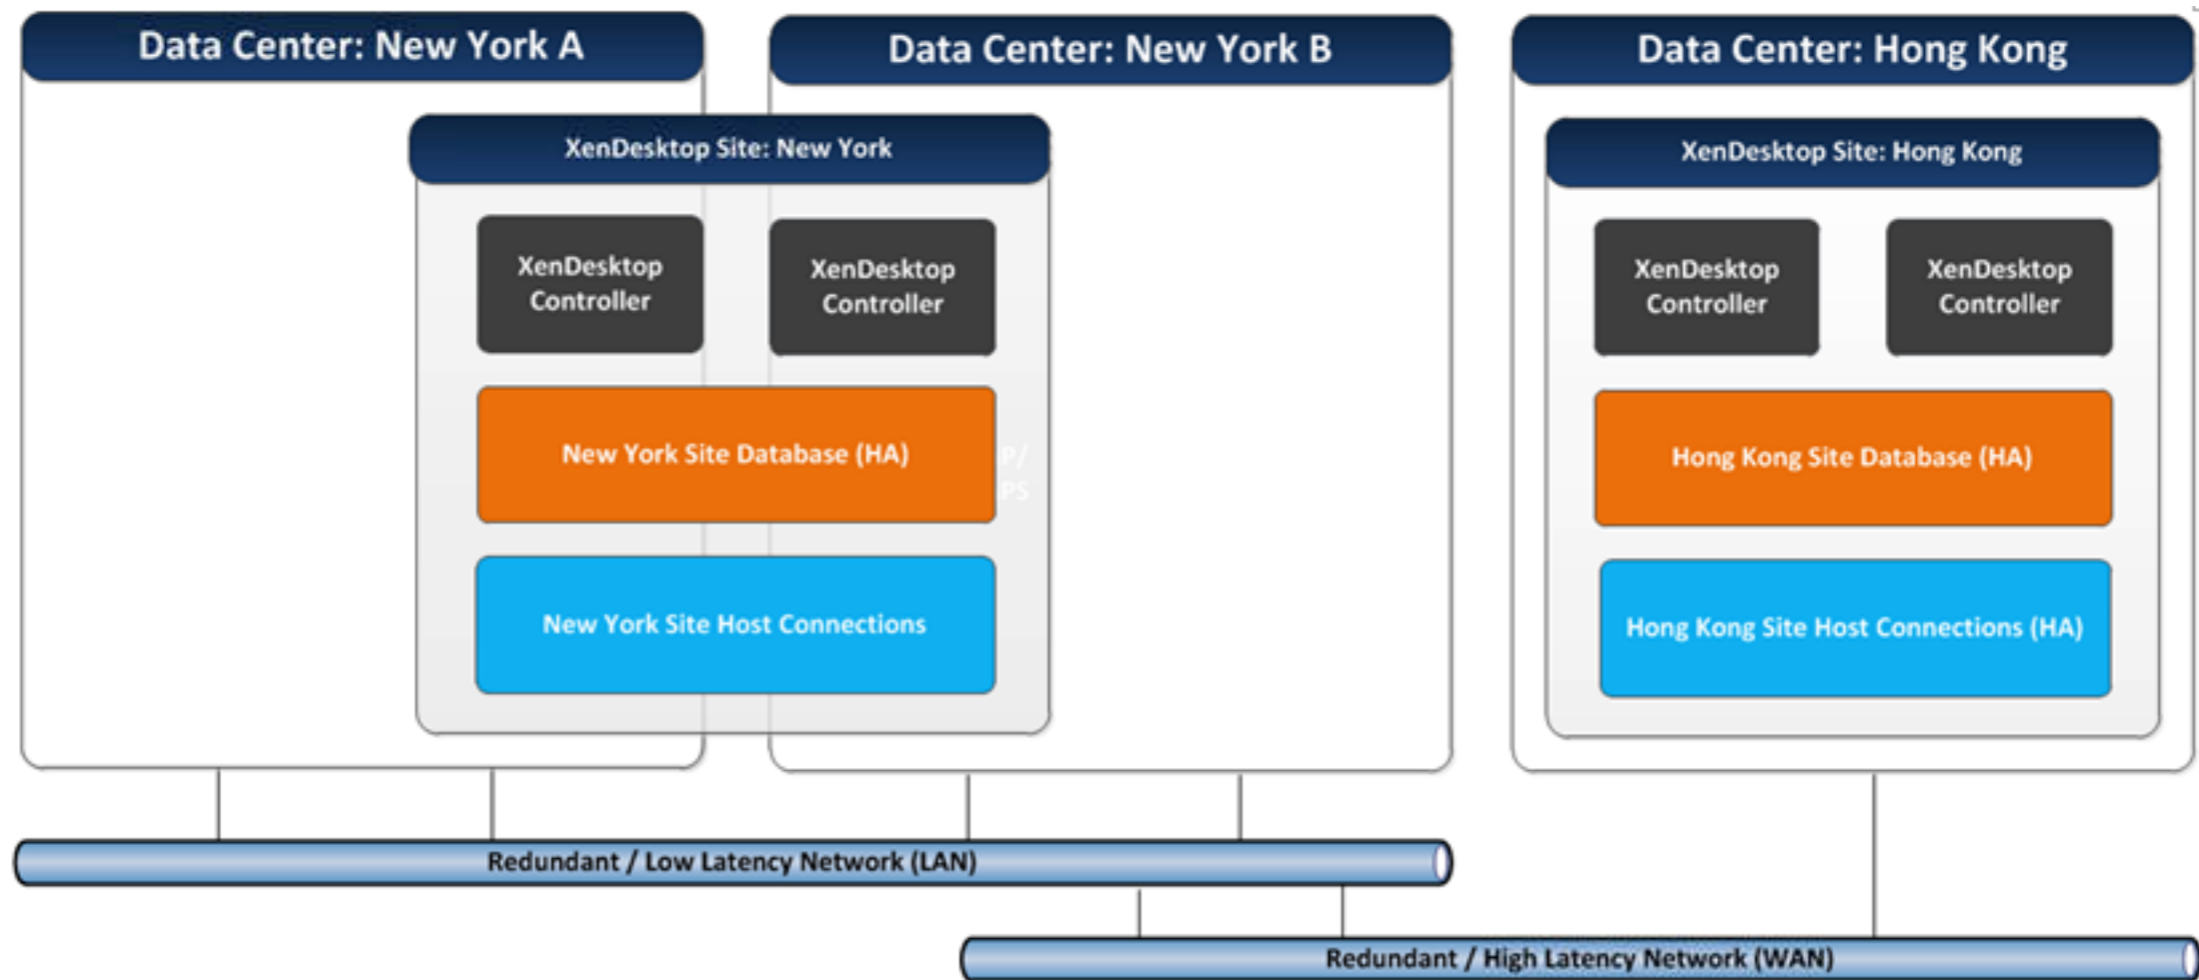
1. Three datacenters connected by a redundant but high-latency network - three sites:

Three XenDesktop Sites



2. Two datacenters connected by a fast, redundant network and a third datacenter connected by a high-latency, redundant network - two sites:

Two XenDesktop Sites



Note: Implementing a single XenDesktop site spanning multiple datacenters connected by fast, redundant links helps to reduce management overhead and supporting infrastructure. However, creating separate sites per datacenter is a safer alternative.

- **Risk Tolerance** – Multiple XenDesktop sites can be created to minimize the impact from a site-wide outage. For example, corruption of the XenDesktop site database could affect site-wide availability. For many organizations, the decreased risk from implementing multiple sites outweighs the additional management overhead and supporting infrastructure required.

Experience from the Field

Finance – A large financial institution hosts 10,000 desktops from a single datacenter. To reduce risk, it was decided that no site should exceed 5,000 desktops. Therefore, despite the desktops being connected by a fast and redundant network, two sites were created.

- **Security** – Although delegated administration is available, high-security organizations may require complete separation between environments to demonstrate compliance with specific service level agreements.

Experience from the Field

Retail – A retail organization required complete separation for employees responsible for managing financial data. To meet this requirement, two separate sites were created within the same datacenter – one for the financial employees and a second for all other employees.

In general, the number of XenDesktop sites should be kept to a minimum to reduce architectural complexity and administrative effort.

[Click here to provide feedback](#)

Decision: Delivery Controller Sizing

The Delivery Controllers authenticate users, enumerate resources, direct user launch requests and control desktop startups, shutdowns and registrations.

Delivery Controller scalability is based on CPU utilization. The more processor cores available, the more virtual desktops a controller can support. Each desktop startup, registration, enumeration and launch request impacts the controller's processor. As the storm increases in intensity, the CPU utilization of the controller will increase. If the CPU reaches a critical threshold, roughly 80%, the site will need to either scale up or scale out.

Adding additional CPU cores to a Delivery controller will lower the overall CPU utilization, thus allowing for greater numbers of desktops supported by a single controller. This is really only feasible when dealing with virtualized controllers as adding virtual CPUs is fairly easy and straightforward. The other alternative is to add another controller into the site configuration. The controller would have the same configuration as other controllers, and the load would be evenly distributed across all controllers, thus helping to reduce the overall load on each single controller.

Testing has shown that a single XenDesktop Controller, using the following configuration, can support more than 5,000 desktops.

XenDesktop Controller Specification for 5K Desktops

Component	Specification
Processor	4 vCPU
Memory	4GB RAM
Network	Bonded virtual NIC
Host Storage	40GB shared storage
Operating System	Windows Server 2012
XenDesktop	7

The following formula can be used to calculate the number of XenDesktop Controllers required:

$$\text{Number of Delivery Controllers} = \frac{\text{Number of Active Sessions per Site}}{5,000} + 1$$

Decision: High Availability

If the server hosting the Delivery Controller is unavailable, users will not be able to access their virtual desktops or published applications. Therefore at least two Delivery Controller servers (N+1 redundancy) should be deployed on different physical servers to prevent this component from becoming a single point of failure. If one controller fails, the others can manage connections and administer the site.

The locations of all Delivery Controllers are specified on the VDA, allowing it to automatically failover if communication with one Delivery Controller is unavailable. The VDA checks the following locations, in order, stopping at the first place it finds the Delivery Controller:

1. A persistent storage location maintained for the auto-update feature. This location contains controller information when auto-update is enabled and after the VDA successfully registers for the first time after installation. For its initial registration after installation, or when auto-update is disabled, the VDA checks the following locations.
2. Policy settings (Delivery Controllers, Delivery Controller SIDs)
3. The Delivery Controller information under the VDA ListofDDCs registry key. The VDA installer initially populates these values, based on the information specified when installing the VDA.
4. OU-based discovery. This is a legacy method maintained for backward compatibility.
5. The Personality.ini file created by Machine Creation Services. Citrix Consulting recommends utilizing the auto-update feature

[Click here to provide feedback](#)

(enabled by default). This feature will simplify management of the environment by keeping VDA's updated when adding and removing Delivery Controllers. For more information about methods for specifying Controllers, see Citrix eDocs – [Manage your Delivery Controller environment](#).

Decision: Host Connection Configuration

Host connections allow controllers within a site to access the hypervisor(s) hosting the virtual desktops or applications. Separate host connections will be required for each XenServer pool, Hyper-V SCVMM server or VMware vCenter server. For redundancy, XenServer host connections should have at least one HA server configured and SCVMM servers and vCenter servers should be highly available.

Host connections define the storage repositories and guest networks that will be used by the virtual machines. Both Provisioning Services and Machine Creation Services use these parameters during the desktop creation process to determine which guest network and storage location should be used. Each host connection can define multiple storage repositories and a single guest network. Therefore, separate host connections are necessary for each guest network required.

Note: Other host connection advanced options should only be configured under the guidance of Citrix support.

Decision: XML Service Encryption

In a typical session, the StoreFront server passes credentials to the Citrix XML Service on a XenDesktop Controller. The Citrix XML protocol uses clear text to exchange all data, with the exception of passwords, which are transmitted using obfuscation. If the traffic between the Storefront servers and the XenDesktop Controllers can be intercepted it will be vulnerable to the following attacks:

- Attackers can intercept the XML traffic and steal resource set information and tickets.

- Attackers with the ability to crack the obfuscation can obtain user credentials.
- Attackers can impersonate the XenDesktop Controller and intercept authentication requests.

For most organizations, the Citrix XML traffic will be isolated on a dedicated physical or virtual datacenter network making interception unlikely. However, for safety consider using SSL encryption to send StoreFront data over a secure HTTP connection.

Decision: Server OS Load Management

Default Load Management policies are applied to all Server OS Delivery Groups. The default settings specify the maximum number of sessions a server can host at 250 and do not consider CPU and Memory usage. Capping session count does not provide a true indication of load, which can lead to an overburdening of Server OS Delivery Groups resulting in a degradation of performance or an underutilization of Server OS Delivery Groups resulting in an inefficient usage of resources.

Citrix Consulting recommends creating unique “custom” Load Management policies for each Delivery Group based on performance and scalability testing. Different rules and thresholds can be applied to each Delivery Group depending on the different resource bottlenecks identified during testing. For more information on the available Load Management policy configurations refer to Citrix eDocs – [Load Management policy settings](#).

If adequate testing cannot be performed prior to production, Citrix Consulting recommends implementing the following “custom” Load Management policy which can be applied to all servers as a baseline:

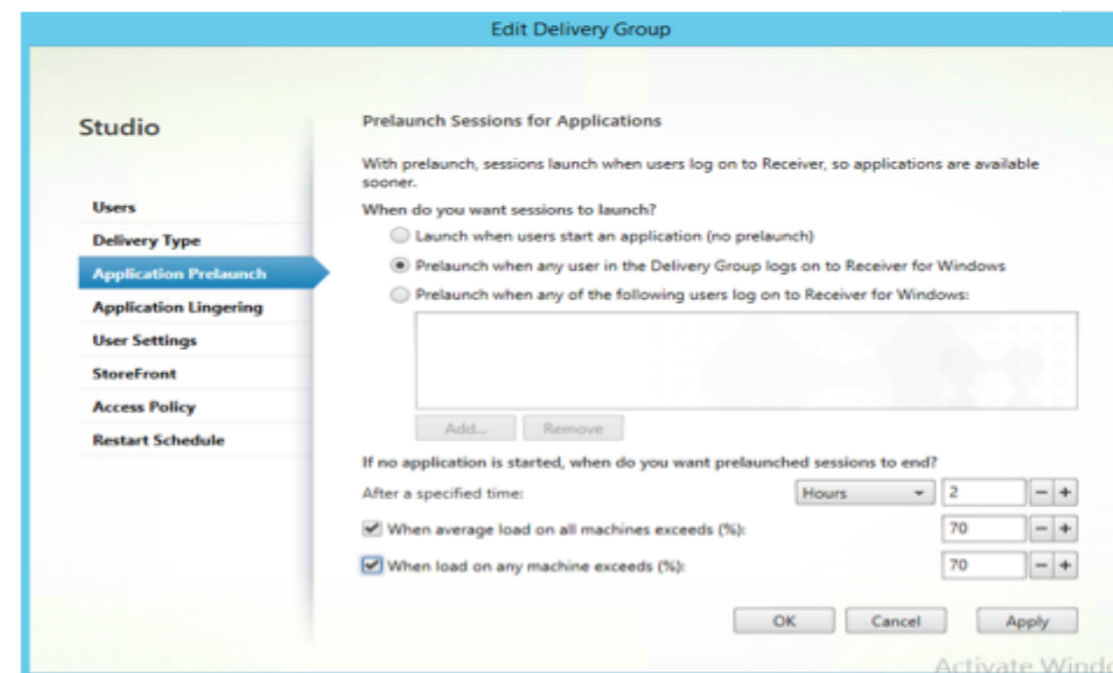
- CPU Usage - Full Load: 80%
- CPU usage excluded process priority – Below Normal or Low
- Memory Usage - Full Load: 80%

- Memory Usage base load – Report zero load (MBs): 786
- Maximum number of sessions – X

The “Maximum number of sessions” policy is included for capping purposes – this is considered a best practice for resiliency. Organizations can choose an initial value of 250 (denoted by “X” above). It is highly recommended that this value and others be customized based on the results from scalability testing.

Decision: Session Pre-Launch and Session Linger

Session prelaunch and session linger are XenApp 7.6 features designed to help users quickly access applications by starting sessions before they are requested (session prelaunch) and keeping user sessions active after a user closes all applications in a session (session linger). Session prelaunch will start a session and leave it open for a specified amount of time until the user connects to the session. Session linger will leave the session open for a specified amount of time in case the user decides to launch the application again. Session pre-launch and session linger are enabled in Studio by configuring Delivery Group settings.



The length of time an unused pre-launched application remains active, or a session is allowed to linger can be set based on:

- A specified time interval (1-99 days, 1-2376 hours, or 1-142,560 minutes). If set too short, a pre-launched application will end before it provides any benefit to the user. The time interval can also be set in PowerShell using the `New/SetBrokerSessionPreLaunch` cmdlet.
- When the average load on all machines within the Delivery Group exceeds a specified percentage (1-99%) the broker will select a session for termination across all VDAs in the group. By default the threshold is set to 70%.
- When the load on any machine within the Delivery Group exceeds a specified percentage (1-99%), a session on the VDA is selected for termination. By default the threshold is set to 70%.

When a threshold is exceeded, the sessions that have been in the pre-launched or lingering state longest are selected first. Sessions are ended one at a time until the load falls below the set threshold. While the threshold is exceeded, no new pre-launched sessions are started.

Note: Servers with VDAs that have not registered and servers in maintenance mode are considered fully loaded and cannot be used for session prelaunch and session lingering. An unplanned outage will cause prelaunch and lingering sessions to end automatically to free capacity.

Consider the following when planning a deployment using session prelaunch or session linger:

- Session Prelaunch will, on average, increase the amount of resources required. Sessions will be started on configured delivery groups regardless of whether they are necessary or not.
- Pre-launched and lingering sessions consume a Citrix license. Unused pre-launched and lingering sessions will disconnect after 15 minutes by default. This value can be adjusted using

PowerShell (`New/Set-BrokerSessionPreLaunch` cmdlet).

- Session Prelaunch is commonly used to improve the launch time of published applications integrated into a VDI desktop.

For additional considerations and information on how to configure Session Prelaunch and Session Linger, please refer to eDocs – [Sessions](#).

XenClient Synchronizer

Citrix XenClient is a client hypervisor that runs on bare metal. XenClient allows users to run multiple local virtual desktops simultaneously, side-by-side and in complete isolation. XenClient provides control over hard-to-manage laptops and provides users with access to their virtual desktop anywhere, anytime, even while disconnected from the network.

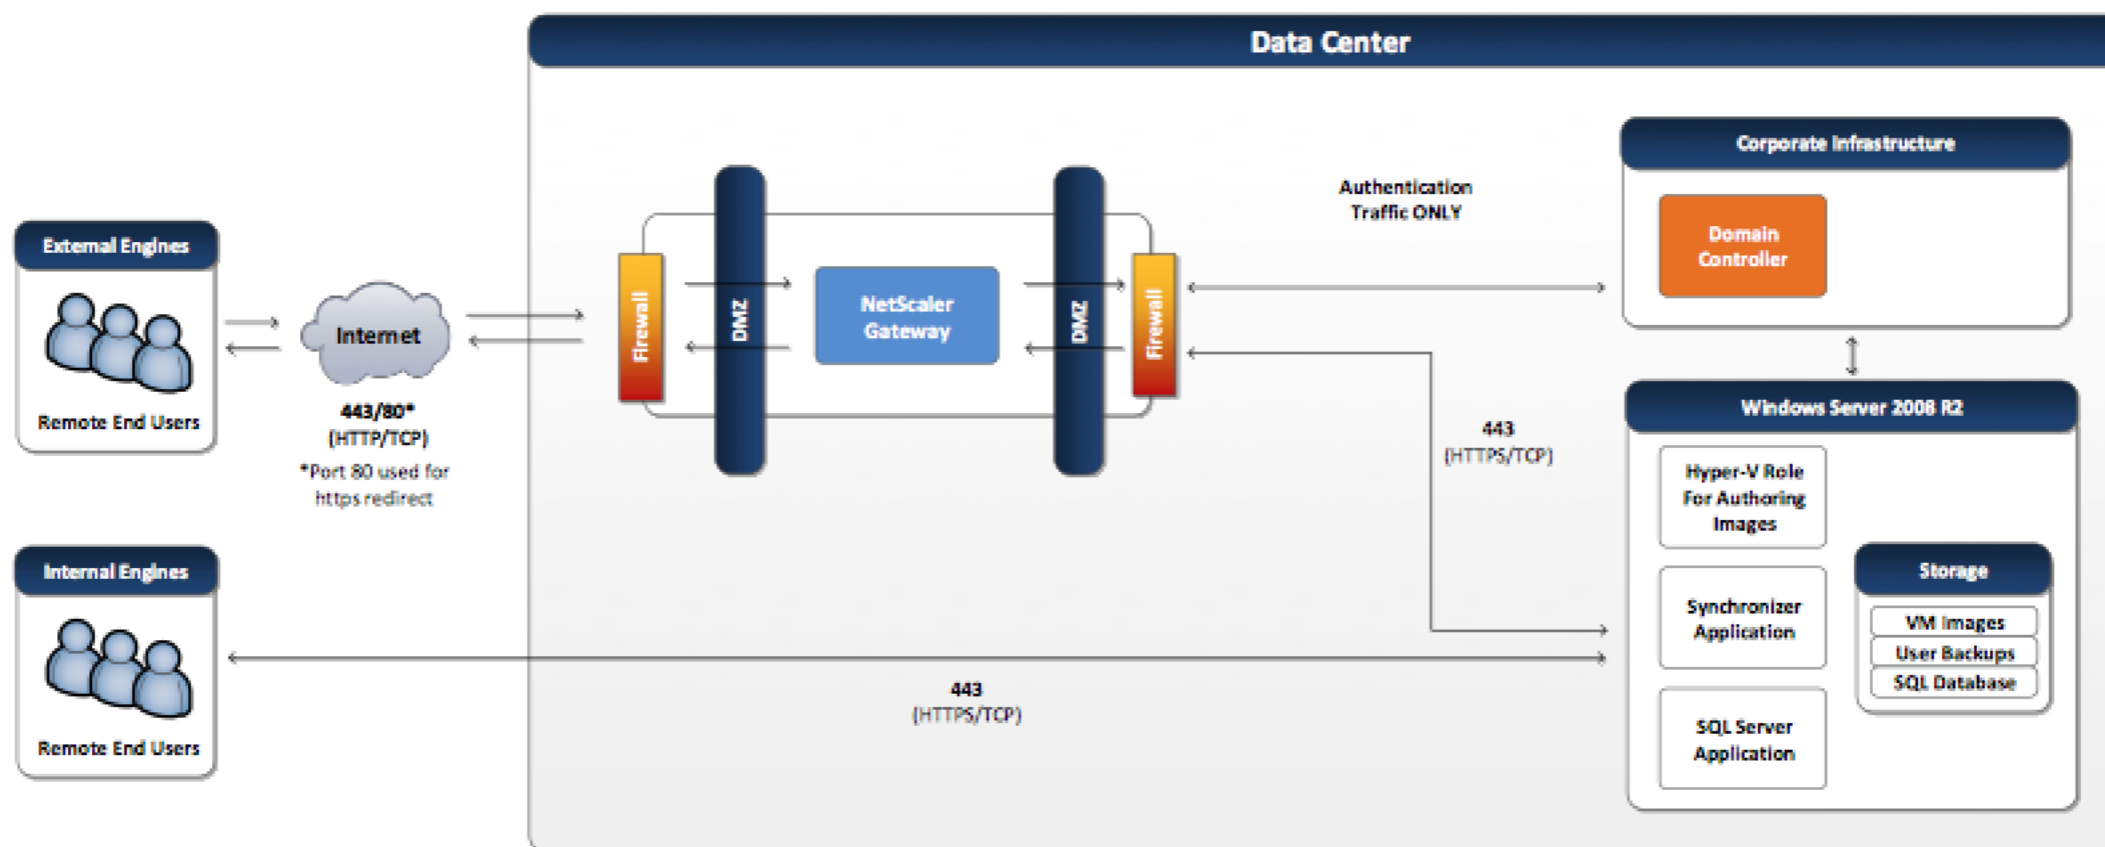
XenClient consists of two main components – synchronizers and engines. At the control layer, the synchronizer provides the delivery of images and policies operations. The engine refers to the XenClient hypervisor that is installed on a physical endpoint computer so that it can host virtual machine (VM) images.

Decision: Architecture

The XenClient Synchronizer can be hosted on a physical or virtual server. However, the Synchronizer is only compatible with Hyper-V when authoring VM images. Although Hyper-V must be used during the VM creation process, XenServer and VMware can also be used to host the synchronizer server. The decision on whether the synchronizer should be virtualized or not will depend on the existing capabilities of the environment:

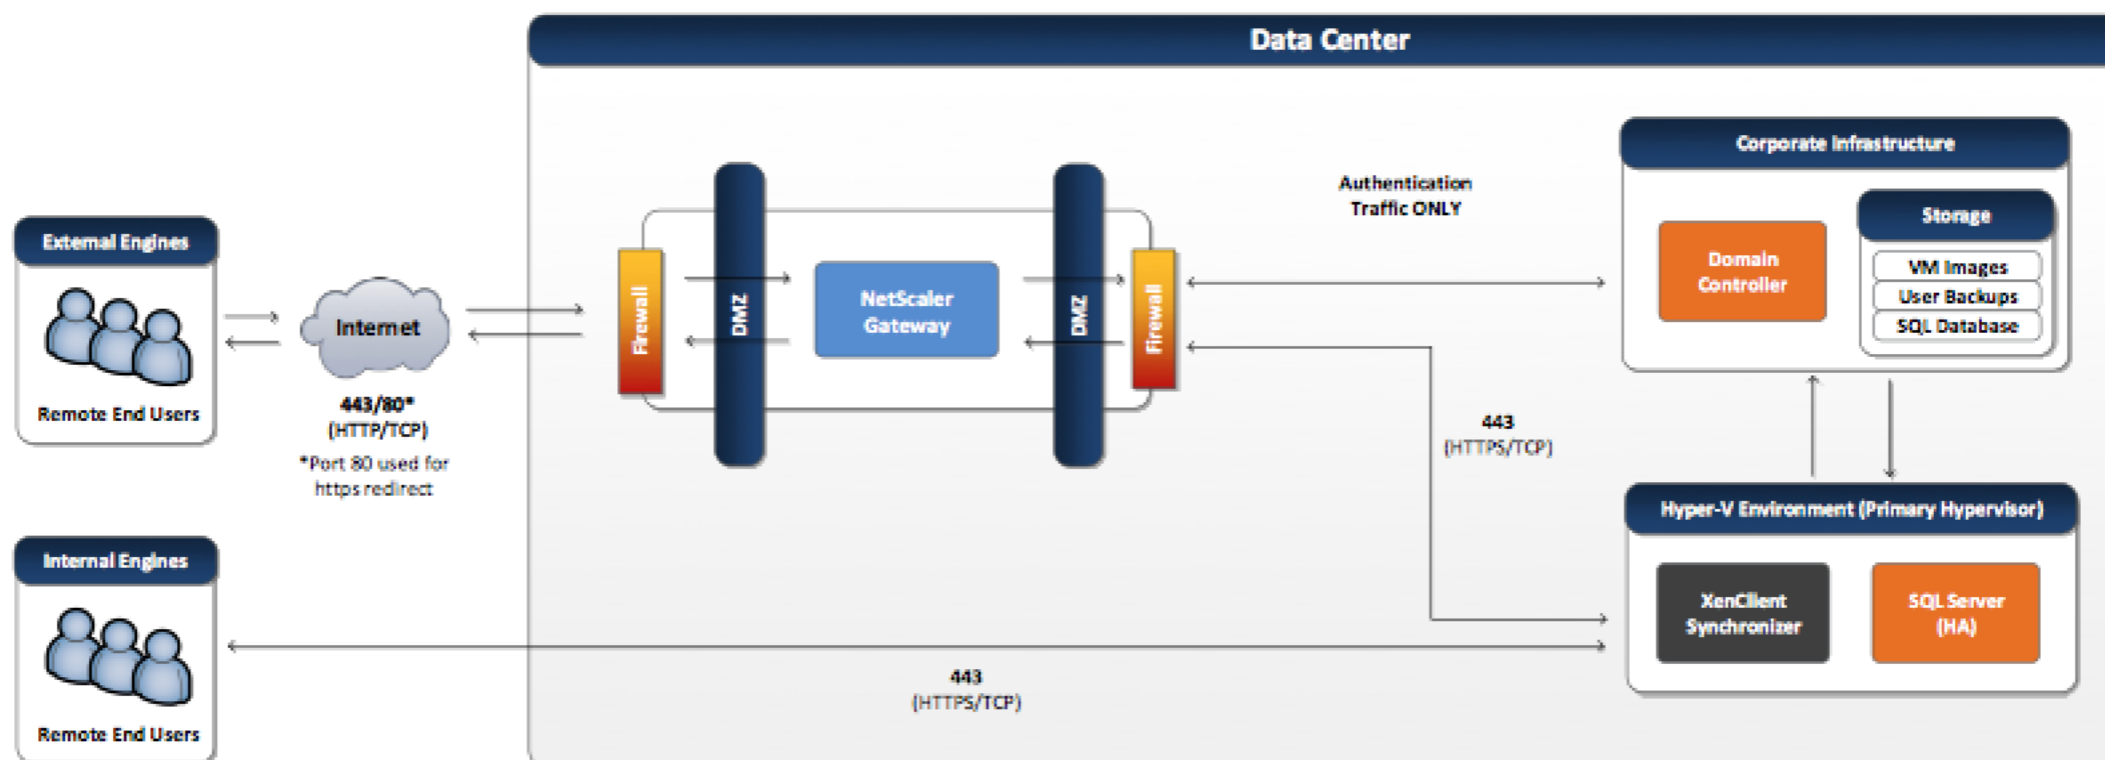
Physical server – If a Hyper-V deployment is not present, it is recommended that the synchronizer role be hosted on a dedicated physical server. In this scenario, the Hyper-V role is added to the same physical server that hosts the synchronizer role. Utilizing a dedicated physical server will ensure that other services do not impact the performance of the synchronizer role. In addition, time spent creating and updating images will improve since the data is transferred locally rather than over the network. Along with the Hyper-V role, the SQL database can also be installed on the same physical server if an existing SQL server is not available. A downside to using a physical server is that it does not offer the same high availability features as a virtual deployment with shared storage. In a virtual deployment, XenClient would be able to utilize the hypervisor's failover features if the host supporting the synchronizer failed.

Synchronizer on Physical Hardware



Virtual machine – If a Hyper-V server is already in use within the organization, it is recommended that the XenClient synchronizer utilize that instance. In this scenario, synchronizer would be installed as a VM on the primary hypervisor. This hypervisor can be Hyper-V, XenServer, or VMware but a Hyper-V server must be used for image authoring. Installing the synchronizer on a VM will help to reduce hardware requirements and allow the synchronizer server to benefit from existing high availability features.

Virtual Synchronizer



Decision: Processor Specification

The XenClient synchronizer can be processor intensive during certain processes. The following list details the major factors that will increase processor usage on the synchronizer as well as recommendations to reduce their impact:

- **Creating and updating images** – The image creation and update process utilizes Hyper-V to perform the necessary actions to the images. If Hyper-V is installed on the same physical server as the synchronizer, sufficient processor and memory should be pre-allocated to prevent contention with the synchronizer service.
- **Publishing images** – The publishing process, which takes approximately 20-40 minutes per image, is a very disk intensive process that causes some processor spikes. Publishing occurs in a serialized fashion, only allowing one VM at a time to be published. Therefore, images should be published during off peak hours. Since the image publishing process takes a relatively short amount of time, this is just a precaution.
- **Deploying images** – The deployment process is the most processor intensive activity due to the network traffic generated, ranging from 20 to 30GB per VM. Although the initial deployment of images will be intensive, subsequent updates will be much smaller and will be downloaded in the background. For this reason, it is important to allocate resources based on normal business activity rather than the initial deployment of VMs.

The synchronizer should be allocated three physical cores for every 1Gbps LAN connection. An additional physical core should be available for the synchronizer server process. At a minimum, it is recommended that two CPU cores be assigned to the synchronizer.

Decision: Memory Specification

The default recommendation is for 8GB of RAM to be assigned to

the synchronizer server. Additional memory will greatly increase performance by reducing disk I/O when distributing images to clients. If additional memory is available, Windows will attempt to cache images in the System Cache. To improve performance by offloading work from the storage subsystem, sufficient memory should be available to cache the compressed copy of the VM images. The compressed version of an image is approximately 50% smaller than the original image size.

Decision: Network Specification

The synchronizer performs bandwidth intensive operations when distributing VM images to endpoint computers. These operations will utilize all available bandwidth until the operation is complete.

The following items should be considered to overcome this potential bottleneck:

- **Capacity planning** – The type of VM deployed (Custom vs. Shared) has a direct impact on the future bandwidth load of the synchronizer. If Custom VMs are chosen, they will only be deployed a single time. Since incremental updates for Custom VMs cannot be deployed, these endpoints will only contact the synchronizer server for policy updates. In contrast, Shared VMs can be patched causing increased bandwidth since the incremental image updates are distributed to XenClient Engines.
- **Quality of service** – A Quality of Service (QoS) policy can be used to effectively manage bandwidth utilization between the XenClient engines and the synchronizer server. The QoS policy should be defined based on the destination IP of the synchronizer server. A port based policy cannot be used since the synchronizer listens on 443 which is also used for secure website traffic.
- **Bandwidth policies** – Policies can be created on the synchronizer to limit the bandwidth and number of connections that a synchronizer can utilize. This will help to ensure that the

synchronizer does not get overloaded with a large number of requests at the same time. While bandwidth policies are important, they do not replace QoS policies. Bandwidth policies enforce specified limits even when additional bandwidth is available.

- **Network rule** – XenClient policies are available which restrict the networks that can be used to download and upload images. This policy is established by specifying the DNS hostname for the networks that will allow downloads and uploads. Although this limits flexibility, it can help to relieve congestion on networks that are known to have limited bandwidth.

Decision: High Availability

The synchronizer, which is an integral component of the XenClient infrastructure, does not have any built-in high availability features. If the synchronizer server were to fail, users would still be able to operate normally although there would be some limitations. Users would not be able to receive images and policy updates. Administrators would be even more restricted, losing the ability to manage any of the endpoints. For the synchronizer, true high availability comes from a strong data backup strategy. The following items should be considered when designing a resilient synchronizer server:

- **Virtual machine HA** – If the synchronizer is deployed as a VM, the high availability features of the hypervisor can be utilized. In this scenario, the synchronizer is installed on a hypervisor with a shared storage solution. This allows the synchronizer VM to failover in the event of a host failure.
- **Data backup** – To guard against failures to shared storage, master VM images should be periodically backed up or archived to a separate storage repository. This will allow for the database to be easily recovered in the event of failure.

Decision: SQL Database

The synchronizer relies on Microsoft SQL, which stores the entire list of endpoints, images and policies assigned to each computer. If the database is unavailable, all user and administration functions will be unavailable. While the SQL database is very important, it does not consume a large amount of resources. The following recommendations should be considered when designing the SQL database for the synchronizer:

- **Server location** – While the SQL database can be installed locally on the synchronizer server, a highly available SQL environment is recommended. The SQL server should be allocated at least 2GB of RAM on the system hosting the instance and at least 500MB of reserved disk space.
- **Backups** – Automated backups of the XenClient database should be made daily. While downtime of the synchronizer can be tolerated, the loss of configuration data would cause irreparable damage.

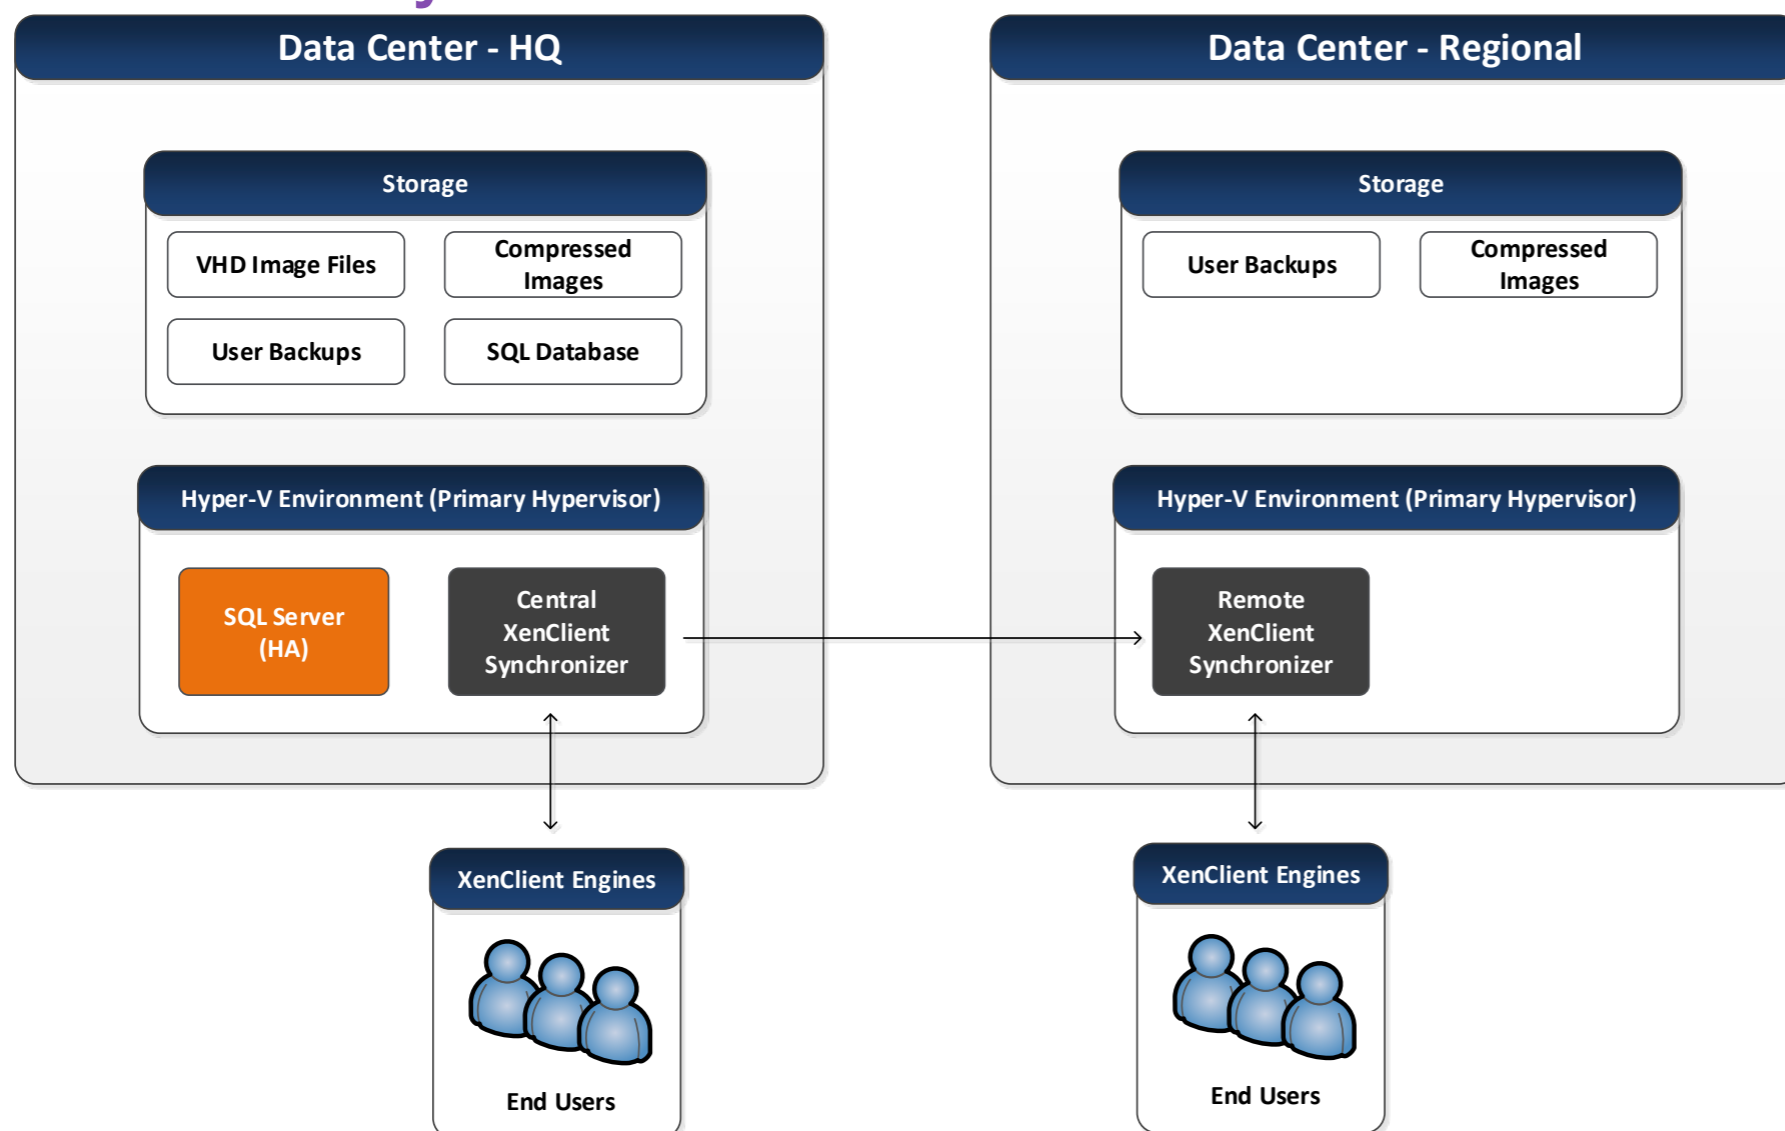
Decision: Remote Synchronizer Servers

When the XenClient user base is distributed across different physical locations, a decision must be made on whether additional central synchronizers or a remote synchronizer should be deployed. A remote synchronizer is a subordinate of the central synchronizer that can be placed at locations physically separated from the central synchronizer. Remote synchronizers allow users at these locations to retrieve and backup images locally instead of traversing the WAN. The pros and cons of each deployment option are discussed below.

Deploy additional central Synchronizers – Deploying additional central synchronizer servers will create a new XenClient environment with a new database. Since a remote synchronizer server requires database access, this may introduce reliability issues when reading or writing data across the WAN. A Central server would allow the SQL database to be hosted locally ensuring the reliability of the connection. The downside of deploying additional central servers is that they require separate databases and cannot be managed from the same console. These limitations increase management overhead since each central server must be managed independently.

Deploy a remote synchronizer – A remote synchronizer server acts as a subordinate of the central synchronizer and is managed from the same management console. The remote synchronizer shares the same SQL database as the central synchronizer allowing the same policies and configurations to be used. A remote synchronizer helps reduce bandwidth between the engine and central synchronizer server by intelligently caching content once a user has initially requested it. The downside to using a remote synchronizer is that it must be able to reach the central Directory environment over the network from where it is located. This may introduce some issues for high latency and low bandwidth sites. A diagram of the architecture is shown below.

Remote Synchronizer Architecture



If it is decided that a remote synchronizer should be deployed, the following items must be considered:

- **Storage** – Enterprise level redundant storage must also be available at the locations where the remote servers are located. This becomes even more important if user data is being backed up to the remote server. While VM image data can be re-downloaded from the central server, user data is only stored on the server that the users are registered with.
- **Management** – Locating servers at remote locations can introduce issues for the IT staff managing the remote hardware. The remote office where the synchronizer is placed may not have sufficient hosting resources or staff on hand that are able to resolve potential issues. This adds to the complexity of the environment and introduces additional risk areas.
- **Firewall configuration** – Since remote servers share the same Microsoft SQL database, the port used by the database server (1433) must be open on all firewalls between the remote server and the central server.
- **Active Directory** – If configured, each remote server must have connectivity to an Active Directory domain controller so that user passwords can be verified and computer objects created.

Experience from the Field

Financial – A mid-sized financial institution needs to provide multiple versions of their financial software running on different operating systems to hundreds of remote locations. These remote locations have very limited bandwidth so image updates from XenClient are delivered from a remote synchronizer located at each site. This allows for a one-to-many-deployment, since updates are deployed to a single server at each site and then distributed locally to each computer.

Personal Fitness – A large personal fitness company was in need of a solution to isolate general computing work from sensitive user data

such as credit card transitions. In order to become PCI compliant, these workloads must be separated and the physical space to locate multiple computers at each desk was unavailable. XenClient has been deployed to hundreds of fitness locations throughout the United States that do not house any server equipment. Since custom images have been deployed, small incremental image updates cannot be pushed out to computers. For this reason, the fitness company decided not to deploy remote synchronizer servers at each remote location. Their current package management system will be used to maintain the custom images.

Decision: Storage Repository Size and Location

The storage repository is the location where user data and VM images are stored. The repository size is directly dependent on the number of images, their size and number of users in the environment. Administrators can control the amount of data used to store the VM images since they create and update them. However, administrators have less control over user data that is backed up to the server. This data can grow at a fast rate depending on the amount of data modified and the number of backups retained.

User Data Backups

User data consists of all data stored on the U: drive of the XenClient endpoint, which includes the user profile and files stored in the My Documents folder. The table below provides an estimate on the storage space required for a 500-user environment. The user disk, which is 100GB, has a compression ratio of 90% by default, reducing the storage space to 10GB per user. The table is modeled on a 'moderate user' that changes approximately 7.5% of their data every time a backup occurs (Next page, top):

User Data Backups

	Initial	Backup One	Backup Two	Backup Three	Backup Four	Backup Five	Backup Six	Backup Seven
Storage Consumed on Synchronizer (Per User)								
Differential	0%	8%	15%	23%	30%	38%	45%	53%
Differential Size (GB)	10	0.8	1.5	2.3	3.0	3.8	4.5	5.3
Cumulative Size (GB)	10	10.8	12.3	14.5	17.5	21.3	25.8	31
Storage Consumed on Synchronizer (All Users)								
Total (GB)	5000	5375	6125	7250	8750	10625	12875	15500
Total (TB)	4.88	5.25	5.98	7.08	8.54	10.38	12.57	15.14

The general recommendation is to keep at least two backup copies of user data. This will allow for a moderate level of resiliency while significantly reducing storage requirements. It is recommended that user data be hosted on RAID-10 storage to offer fast reads/writes and redundancy.

Virtual Machine Images

XenClient can also require a large amount of storage for VM image data. Each VM is actually stored twice, once compressed for image deployment and a second time that is uncompressed for image updates. The compressed copy is a collection of Tar GZip files, which are downloaded to clients and remote servers. The compressed image typically requires 50% less storage than the uncompressed image. Additional storage space is also consumed each time a VM image is updated. For every update, additional differencing disk files will be created on the server. These files form a chain linked back to the original image that allow it to be reverted back to previous versions.

The following recommendations should be considered when designing storage space for virtual images:

- **Expansion** – VM images grow in size each time a new version of the image is created. Allocate sufficient space for each VM to grow approximately 50% of its original size.
- **Performance** – It is recommended that VM images be located

on RAID-5 storage due to its ability to offer redundancy and fast reads.

- **Reclaim space** – A rollup operation should be performed once old versions of a VM image become outdated. This process combines the differencing disk files of a VM eliminating previous revisions up to the version selected. It is recommended that a revision plan be established to help administrators decide how many versions should be kept before a rollup occurs. As a general recommendation, at least two revisions of each VM should be kept in case issues arise with the latest version.
- **Daily backup rollup** – The daily backup rollup process consolidates VM backups that have exceeded their retention limits. These limits are set through the backup policy assigned to the VM. Rolling up old backups reduces disk space usage on the server. It is recommended that this option be enabled and set to run during off peak hours.

Decision: External Access

To allow users to receive image and policy updates when they are outside the corporate network, the synchronizer must be accessible over the Internet. The following list provides some of the benefits from enabling external access to the synchronizer:

- **VM deployment** – When accessible over the Internet, the

synchronizer will allow mobile users located off the corporate network to receive VM images and policy updates. It is recommended that remote access be configured for all organizations with mobile users.

- **Lockout policy** – A lockout policy is set to enhance security by setting how long a XenClient computer can operate before it must contact the synchronizer server. If this limit is reached, a user will be locked out of their computer, although their data would still be intact. This policy exists to ensure that users receive the latest VM image and policy updates. It is recommended that this policy be used in high security environments to ensure that all Engines adhere to the latest image and policy updates deployed. If this policy is enabled for remote users, the synchronizer must be available externally.
- **Kill switch** – The kill switch option is available to provide a failsafe against stolen or compromised XenClient computers. Making the synchronizer available externally enables endpoint computers to be terminated when they connect to the Internet. The endpoint will receive the kill command once the policy update interval time has been reached.

Note: As the kill switch option requires that the Engine have Internet connectivity, it should be used in combination with the lockout policy.

If the synchronizer has been deployed externally, aliases should be configured so that it can be reached using the same hostname regardless of whether the Engine is internal or external. When deploying the synchronizer for external users, two options are available when exposing the synchronizer to external traffic:

- **Behind firewall** – The synchronizer server sits in the DMZ behind a firewall with only the appropriate port (443) accessible externally.
- **Behind Citrix NetScaler** – Citrix NetScaler provides additional security and scalability when external users connect to the

synchronizer. In this model, the NetScaler is located in the DMZ and the synchronizer is located inside the corporate network. NetScaler acts as a ‘man in the middle’ to ward off TCP/IP oriented exposures such as denial of service (DOS) attacks. There are two deployment modes that can be chosen when using a NetScaler:

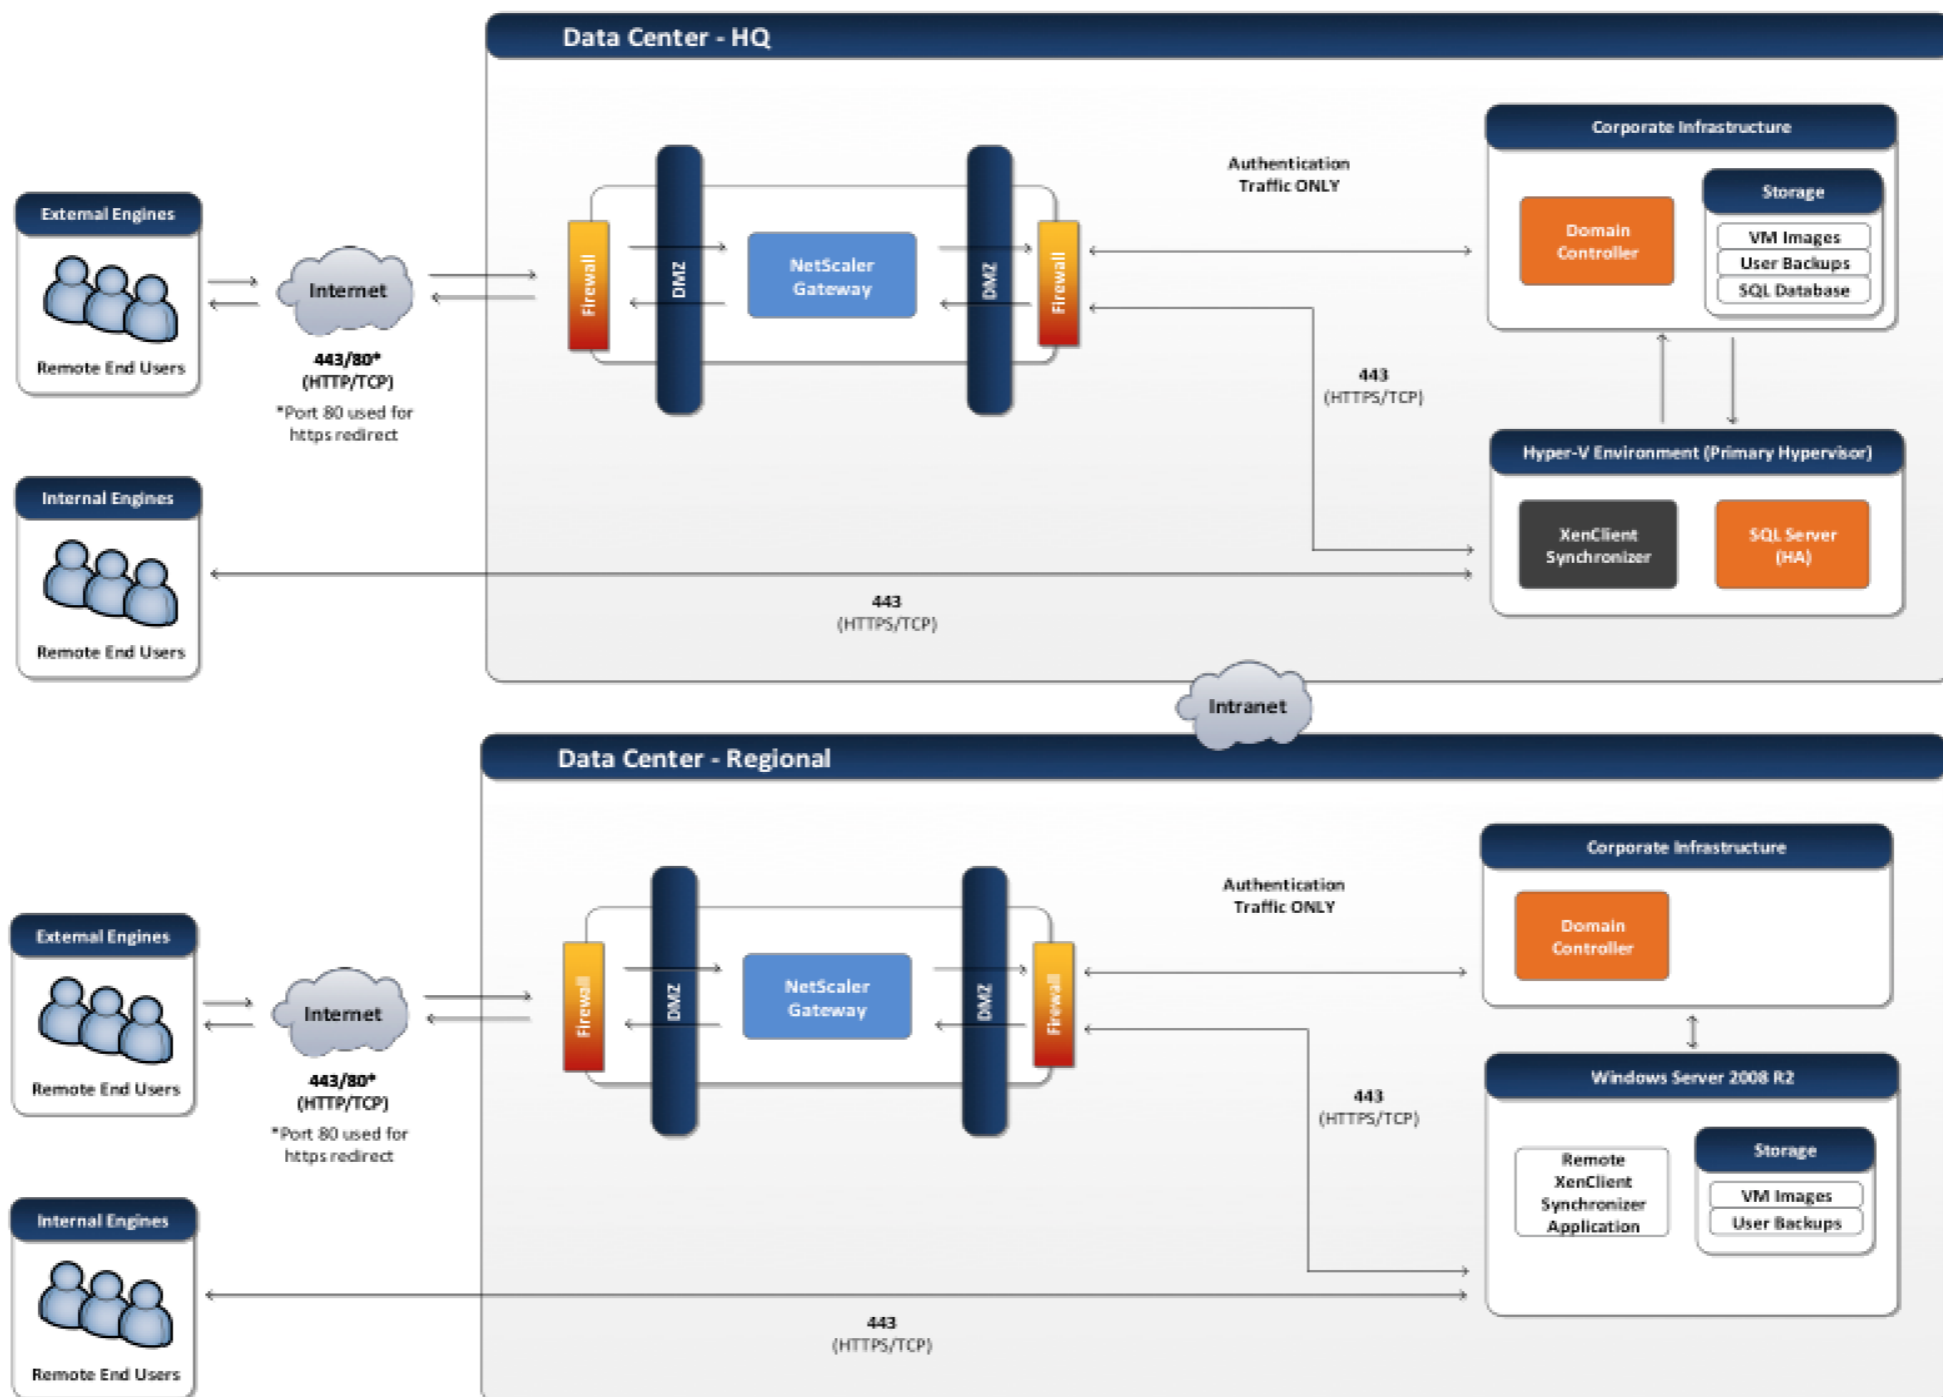
- *SSL offload* – Reduces the workload on the synchronizer allowing it to achieve improved single server scalability. The synchronizer certificates are placed on the NetScaler device so that it can offload the SSL processing from the synchronizer.
- *SSL bridging* – SSL traffic is passed directly from the NetScaler to the synchronizer. Although SSL Bridging reduced the scalability of the synchronizer, it does not require that the synchronizer certificates be placed on the NetScaler.

For added security and scalability, it is recommended that Internet facing synchronizers be placed behind a NetScaler with SSL Offload configured.

For example, a central and a remote synchronizer have been deployed behind redundant NetScaler devices. A remote synchronizer was deployed because the organization has two locations from which a large number of XenClient users are located. The remote synchronizer will allow users at ‘Datacenter – Regional’ to receive VM images from their local network, helping to save bandwidth. The central server has been virtualized while the remote server is a physical machine since no virtualization environment is available at ‘Datacenter – Regional’.

As the organization already utilizes NetScaler for load balancing web servers, using them to protect the XenClient environment is an added benefit. The NetScaler will help to guard against network level attacks such as denial of service (DOS) attacks. Additional features such as ACL lists, surge protection and rate limiting can also be implemented to help prevent the synchronizer server from being overloaded by malicious traffic.

Central (Virtual) and Remote Synchronizer (Physical) Servers



Decision: Active Directory Integration

The XenClient synchronizer can integrate with Active Directory to join images to the domain, authenticate users to their desktop and assign administrator privileges. The synchronizer server itself does not need to join the domain. However, a XenClient service account should be created with the correct privileges to create computer objects in the designated organizational unit. The following recommendations should be considered when designing the integration between the synchronizer and Active Directory:

- **Pre-join images to domain** – XenClient images can be deployed pre-joined to the domain. When this option is chosen, XenClient automatically creates the computer objects and generates unique host names for each VM. The host names can be customized automatically using the operating system type, user name and random number variables. While the host names can be tailored, only a single naming convention is specified for every VM in the environment. This constraint may not comply with organizational standards for naming to include location or user group in the hostname.
- **Administrator login** – It is recommended that administrator accounts to the synchronizer be assigned to active directory user accounts helping to streamline management.
- **User synchronization** – The synchronizer has the ability to schedule automatic updates to reflect any changes in Active Directory users and groups. It is recommended that this feature be enabled to ensure synchronizer has the correct user group assignments.
- **Active Directory SSL** – It is recommended that SSL be enabled when connecting to the Active Directory server for authentication. This will ensure that all user name and password data is encrypted.

Image Controllers

Provisioning Services

Citrix Provisioning Services (PVS) uses streaming technology to simplify the deployment of virtual and physical desktops. Computers are provisioned and re-provisioned in real-time from a single shared-disk image. In doing so, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is performed on the master image.

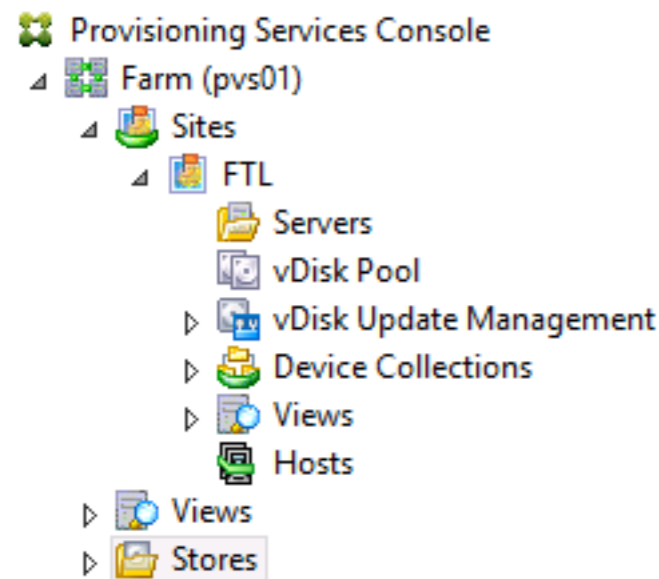
Decision: Farms

A Provisioning Services farm represents the top level of the Provisioning Services infrastructure. All provisioning servers in a farm share the same SQL database and Citrix license server. There are several factors that must be considered when determining the number of Provisioning Services farms required:

- **Network** – Provisioning servers are constantly communicating with the farm database to retrieve system configuration settings. Therefore, separate farms should be created for each physical location where target devices reside, unless they are connected to the database server by a fast and robust connection.
- **Administration** – Organizations may need to maintain the separation of administrative duties at a departmental, regional or countrywide basis. Additional Provisioning Services farms will add some complexity to the management of the environment. However, this overhead is typically limited to initial configuration, desktop creation and image updates.

Decision: Sites

Each Provisioning Services farm contains one or more sites. A site is a logical entity that contains provisioning servers, vDisk pools, target device collections, and hypervisor hosts. Multiple sites share the same database; however target devices can only failover to other provisioning servers within the same site.



Additional sites may need to be created in the following scenarios:

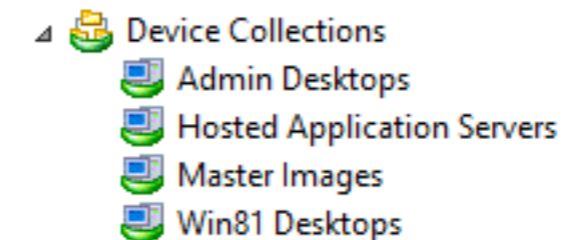
- **Networking** – Sites can be used to control the Provisioning Services streaming traffic. For example, a site can be created for a physical location, rack, or blade chassis to help ensure that streaming traffic stays local and does not become a bottleneck.
- **Organization** – Another practical reason for building multiple sites is due to organizational changes. For example, two companies may have recently merged through acquisition, but need to keep resources separate while integration takes place. Configuring the organization to use separate sites is one way to keep the businesses separate but managed centrally through the Provisioning Services console.

Only create additional sites if the business requirements warrant it.

A single site per farm is easier to manage and requires no additional configuration.

Decision: Device Collections

Device collections provide the ability to create and manage logical groups of target devices. Creating device collections simplifies device management by allowing actions to be performed at the collection level rather than the target device level.



Device collections can represent physical locations, subnet ranges, chassis or different departments within an organization. Collections can also be used to logically separate production target devices from test and maintenance ones.

Consider creating device collections based on vDisk assignment so that the status of all target devices assigned to a particular vDisk can be quickly identified.

Decision: High Availability

Provisioning Services is a critical component of the virtual desktop infrastructure. The following recommendations should be followed to eliminate single points of failure:

- **Database** – Provisioning Services provides two options for supporting a highly available database configuration:
 - Offline Database Support – Information can be found in the eDocs article – [Offline Database Support](#).
 - Database Mirroring – allows Provisioning Services to work from a mirrored copy of the SQL database in the event that the primary version becomes unavailable. For more

information, please refer to the Citrix eDocs article – [Database Mirroring](#).

- **Provisioning Server** – A minimum of two provisioning servers should always be implemented per site. Sufficient redundancy should be incorporated into the design so that a single server failure does not reduce the total number of target devices that can be supported per site.

The Provisioning Services boot file should be configured for high availability. Up to four Provisioning Servers may be listed in the boot file. Target devices will try to contact the servers in the order that they are listed. The server that responds may not necessarily be the server that will provide streaming services to the target device. If Load Balancing is enabled, the target device may be reassigned to another server in the site that is less loaded than the others.

- **vDisks and Storage** – For vDisk Stores hosted on Direct Attached Storage, replication should be used to synchronize the vDisks. If using Network Attached Storage, ensure that the vDisks are hosted on a highly available network share.
- **Networking** – The provisioning servers network adapters should be teamed, and Provisioning Services NIC failover enabled by specifying the NICs that the Provisioning Services drivers bind to. Provisioning Services could also be placed on a server with redundant network configuration. This would allow Provisioning Services to seamlessly failover to a secondary NIC if an issue occurs

Note: The target devices will only failover to NICs that are in the same subnet as the PXE boot NIC.

Trivial File Transfer Protocol (TFTP) is a communications protocol used for transferring configuration or boot files between machines. Provisioning services can use TFTP to deliver the bootstrap file to target devices. There are several options available to make the

TFTP service highly available. Some of the more commonly used options are:

- **DNS Round Robin** – A DNS entry is created for the TFTP service with multiple A records corresponding to the TFTP services running on the provisioning servers in the farm. This method is not recommended since the state of the TFTP service is not monitored. Clients could potentially be sent to a non-functioning server.
- **Hardware load balancer** – Use a hardware load balancer, such as Citrix NetScaler, to create virtual IPs that corresponds to the provisioning servers. The NetScaler can intelligently route traffic between the provisioning servers. In the event that one of the servers becomes unavailable, NetScaler will automatically stop routing TFTP requests to that server.
- **Multiple DHCP Option 66 entries** – This method is easy to implement but requires a DHCP service that supports entering multiple entries in option 66. Microsoft DHCP server allows one option 66 entry so this method would not be feasible in environments with Microsoft DHCP services. If using a non-Microsoft DHCP server or appliance, check with the manufacturer to verify that multiple option 66 entries is supported

For more information, please refer to the Citrix whitepaper – [High Availability for TFTP with Provisioning Services](#).

There are other options available that can achieve the same result without having to use TFTP:

- **Proxy DHCP** – Use the provisioning servers PXE service to provide the bootstrap information. If one of the servers is down, the next available server in the farm can provide the bootstrap information. This method requires the provisioning servers to be on the same broadcast domain as the target devices. If there are other PXE services running on the network (Altiris, SCCM, etc.) then multiple VLANs may be required to keep the PXE services

from interfering with each other.

- **Boot Device Manager** – Use the Boot Device Manager to create a bootstrap file that is either placed on the local hard drive, or used as a bootable ISO file. If the ISO file is used, configure the target devices to boot from the CD/DVD-ROM drive, and place the ISO file on a highly available shared network location or local storage of each target device. When either method is utilized, the TFTP service is not used at all.

High availability should always be incorporated into the Provisioning Services design. Although high availability may require additional resources and increased costs, it will provide a highly stable environment so that users experience minimal impact due to service outages.

Decision: Bootstrap Delivery

A target device initiates the boot process by first loading a bootstrap program which initializes the streaming session between the target device and the provisioning server. There are three methods in which the target device can receive the bootstrap program:

- **Using DHCP Options**

1. When the target device boots, the target device sends a broadcast for IP address and boot information. DHCP will process this request and provide an IP as well as scope option settings 66 (the name or IP address of the Provisioning Services TFTP server) and 67 (the name of the bootstrap file).

Note: If using a load balancer for the TFTP service then the address of the load balancer is entered in option 66.

2. Using TFTP, a request for the bootstrap file is sent from the target device to the provisioning server. The target device downloads the boot file from the provisioning server.
3. The target device boots the assigned vDisk image.

Note: Requires UDP/DHCP Helper to be configured when targets are not on the same subnet as the DHCP servers in order to receive PXE broadcasts.

- **Using PXE Broadcasts**

1. When a target device boots from the network, the target device sends a broadcast for an IP address and boot information. DHCP will process this request and provide an IP address. In addition, all provisioning servers that receive the broadcast will return boot server and boot file name information. The target device will merge the information received and start booting.
2. Using TFTP, a request for the bootstrap file is sent from the target device to the provisioning server which responded first. The target device downloads the boot file from the provisioning server.

Note: Make sure no other PXE services are in use on the same subnet, such as the Altiris PXE service, or conflicts may occur with Provisioning Services.

Note: Requires UDP/DHCP Helper to be configured when targets are not on the same subnet as the DHCP and PVS servers in order to receive PXE broadcasts.

- **Using Boot Device Manager** – The Boot Device Manager (BDM) creates a bootfile that target devices obtain through an ISO image mounted from a network share, a physical CD/DVD that is placed in the server, or the bootfile information is written to a hard drive partition local to the target devices.

A summary of the advantages and disadvantages for each delivery method is listed in the following table.

Bootstrap delivery options and advantages/disadvantages

Delivery Method	Advantages	Disadvantages
DHCP Options	Easy to implement	Requires changes to production DHCP service. DHCP service may only allow one options 66 entry. Requires UDP/DHCP helper for targets on different subnets.
PXE	Easy to implement	Can interfere with other running PXE services on the same subnet. Requires UDP/DHCP helper for targets on different subnets.
BDM	Does not require PXE or TFTP services	Extra effort required to boot physical target devices

Note: When configuring the bootstrap file, up to four provisioning servers may be listed. The order in which the provisioning servers appear in the list determines the order which the provisioning servers are accessed. If the first server does not respond, the next server in the list is contacted.

Decision: Write Cache Placement

The write cache uniquely identifies the target device by including the target device's MAC address and disk identifier. If data is written to the provisioning server vDisk in caching mode, the data is not written back to the base vDisk. Instead, it is written to a write cache file in one of the locations specified below:

- **Cache on Device Hard Drive** – The write cache is placed on the target device's hard drive. This option frees up resources on the provisioning server since it does not have to process write requests. While not as fast as RAM, this option provides a fast response time because the read/write to and from the write cache is local.
- **Cache on Device Hard Drive Persisted** – The option is the same as cache on device hard drive, except that the cache persists between reboots. This feature is experimental and only

supported on Windows 2008 R2, Windows 7 and later.

- **Cache in Device RAM** – The write cache is placed in the RAM of the target device. Although Cache in Device RAM offers high levels of performance it is expensive to implement and will negatively affect stability if RAM is exhausted.
- **Cache in Device RAM with overflow on hard disk** – The write cache is placed in the RAM of the target device. In the event that RAM is nearly exhausted, a portion of the hard disk is used as an overflow. This feature requires Windows 7 or Server 2008 R2 and above. It can significantly reduce the IOPS requirements, but may require additional RAM and disk capacity.
- **Cache on Provisioning Server Disk** – The write cache is stored on the provisioning server as a temporary file. All writes are handled by the provisioning server, which can increase disk IO and network traffic. This configuration is easy to setup and requires no additional resources, however performance of the target devices is affected since all read and write requests to the cache must cross the network. Placing the write cache on the provisioning server is typically only recommended for test environments.

Note: The write cache files on the provisioning server should be encrypted in high security environments

- **Cache on Provisioning Server Disk Persisted** – This option is the same as cache on the provisioning server disk, except changes are saved between reboots. The benefit of this option is that target devices are able to save changes made to the vDisk image and have them persist after a reboot. Any changes made to the vDisk will force the cache file to be marked invalid. For example, if the vDisk is set to Private Image Mode, all associated cache files are marked invalid. Any of the following changes listed below will also cause the cache file to be marked invalid.

- Placing a vDisk in maintenance mode
- Mapping the drive from the console
- Changing the location of the write cache file
- Using automatic update

Note: The write cache files on the provisioning server should be encrypted in high security environments.

The table shown below compares the advantages and disadvantages of the different write cache options.

Experience from the Field

Gas Distribution – A gas distribution company is designing a Provisioning Services solution to support five hundred virtual desktops. The design has the following requirements: it must be able to support a wide variety of desktop users, it must offer fast performance, and any problems with one desktop should not impact the other desktops. To meet these requirements, the gas distribution company went with a design that places the write cache on the target device hard drive.

Write Cache Destination comparison

Write Cache Destination	Performance	Windows OS Supported	Cost	Changes Persist After Reboot	Best Suited For
Device Hard Drive	Med-High	All	Low	No	Production; any size
Device Hard Drive Persisted	Med-High	Windows 7 Windows 2008 R2	Low	Yes	Test or niche use cases.
Device RAM	High	All	High	No	Environments where performance is critical; any size
Device RAM with overflow disk	High	Windows 7 Windows 8 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2	Low-Med	No	Environments where performance and high availability is important.
Provisioning Server Disk	Low-Med	All	Low	No	Test & small environments
Provisioning Server Disk Persisted	Low-Med	All	Low	Yes	Test & small environments

Decision: vDisk Format

Provisioning Services supports the use of fixed-size or dynamic vDisks:

- **Fixed-size disk** – For vDisks in private mode, fixed-size prevents disk fragmentation of the disk hosting the vDisk, and offers improved write performance over dynamic disks.
- **Dynamic disk** – Dynamic disks require less storage space than fixed-size disks, but offer significantly lower write performance. Although vDisks in Shared mode do not perform writes to the vDisk, the time required to complete vDisk merge operations will increase with dynamic disks. This is not a common occurrence as more environments choose to create new vDisks when updating.

Since most reads will be to the System Cache in RAM, there is no significant change in performance when utilizing fixed-size or dynamic disks. In addition, dynamic disks require significantly less storage space. Therefore dynamic disks are recommended.

Decision: Estimating Store Size

A vDisk consists of a VHD base image file, a properties file (.pvp), and may also contain a chain of referenced VHD differencing disks (.avhd). Every time a vDisk is updated using the Provisioning Services versioning method, a new differencing disk file is created.

When estimating the size of the vDisk store, you will need to consider:

- **Total size of vDisks** – Add together the size of the planned vDisk images that will be delivered.
- **Maximum # of versions in vDisk chain** – vDisk versioning simplifies vDisk update and management tasks, providing a more flexible and robust approach to managing vDisks. Differencing disks are created to capture the changes made to the base disk image, leaving the original base disk unchanged. Each differencing disk that is associated with a base disk represents a different version.

Note: When too many differencing disks are chained from the base image disk, Provisioning Services performance begins to degrade. Therefore it is recommended that the number of differencing disks should not exceed more than five to seven versions deep.
- **vDisk version change %** – The size of the differencing disks will vary according to the changes made to the vDisk. The more significant the change is to the vDisk, the larger the differencing disk that will be created. If it is anticipated that new applications will be added to the vDisk then plan for differencing disks that could be large in size. The following table shows the affect that a small, medium, and large change has on the size of a vDisk differencing disk. A 20GB Windows 8 vDisk was used in this example.

Differencing disk growth by application size

vDisk Change	Application Install size	Approximate Differencing Disk size
Small (Install 1 application – Firefox)	20 MB	1.1 GB
Medium (Install 2 applications – Firefox and Adobe Reader)	56 MB	1.6 GB
Large (Install MS Office 2010 Pro x86)	2.2 GB	4 GB

The following formula can be used as a guide when estimating the size of the vDisk store:

$$vDisk\ Store\ Size = Total\ Size\ of\ vDisks + (Total\ Size\ of\ vDisks * vDisk\ Version\ Change\ %) * Maximum\ \#\ of\ Versions\ in\ vDisk\ Chain$$

Consider the following example. You plan to deploy three vDisk images:

- Windows 8 (x64) image = 40GB
- Windows 7 (x64) image = 40GB
- Windows 7 (x32) image = 35GB

Each vDisk will be limited to five differencing disks. You anticipate that the differencing disk will be 20% of the master vDisk image. The estimated size required for the Provisioning Services store will be:

$$vDisk\ Store\ Size = 115\ GB + (115\ GB * 20\%) * 5 = 230\ GB$$

Decision: vDisk Replication

vDisks hosted on a SAN, local, or Direct Attached Storage must be replicated between vDisk stores whenever a vDisk is created or changed. Provisioning Services supports the replication of vDisks from stores that are local to the provisioning server as well as replication across multiple sites that use shared storage. The replication of vDisks can be performed manually or automatically:

- **Manual** – Manual replication is simple, but can be time consuming, depending on the number of vDisks and vDisk stores. If an error occurs during the replication process, administrators can catch them straight away and take the appropriate steps to resolve them. The risk of manual replication is vDisk inconsistency across the provisioning servers which will result in load balancing and failover to not work properly. For example, if a vDisk is replicated across three servers and then one of the vDisks is updated, that vDisk is no longer identical and will not be considered if a server failover occurs. Even, if the same update is made to the other two vDisks, the timestamps on each will differ, and therefore the vDisks are no longer identical.
- **Automated** – For large environments, automated replication is faster than the manual method due to the number of vDisks and vDisk Stores required. Some automated tools, such as [Microsoft DFS-R](#), support bandwidth throttling and Cross File Remote Differential Compression (CF-RDC), which use heuristics to determine whether destination files are similar to the file being replicated. If so, CF-RDC will use blocks from these files to minimize the amount of data transferred over the network. The risk of automated replication is that administrator do not typically monitor replication events in real-time and do not respond quickly when errors occur, unless the automation tool has an alerting feature. Some tools can be configured to automatically restart the copy process in the event of a failure. For example, [Robocopy](#) supports “resume copying” in the event that the network connection is interrupted.

For medium and large projects, use a tool to automate vDisk replication. Select a tool that is capable of resuming from network interruptions, copying file attributes and preserving the original timestamp.

Note: Load balancing and high availability will not work unless the vDisks have identical timestamps.

[Click here to provide feedback](#)

Decision: Virtual or Physical Servers

Citrix Provisioning Services can be installed on virtual or physical servers:

- **Virtual** – Offers rapid server provisioning, snapshots for quick recovery or rollback scenarios and the ability to adjust server resources on the fly. Virtual provisioning servers allow target devices to be distributed across more servers helping to reduce the impact from server failure. Virtualization makes more efficient use of the system resources.
- **Physical** – Offers higher levels of scalability per server than virtual servers. Physical provisioning servers mitigate the risks associated with virtual machines competing for underlying hypervisor resources.

In general, virtual provisioning servers are preferred when sufficient processor, memory, disk and networking resources can be made available and guaranteed to be available.

Note: For high availability, ensure that virtual Provisioning Servers are distributed across multiple virtualization hosts. Distributing the virtual servers across multiple hosts will eliminate a single point of failure and not bring down the entire Provisioning Services farm in the event of a host failure.

Decision: Provisioning Server Memory

The Windows operating system hosting Provisioning Services partially caches the vDisks in memory (system cache) reducing the number of reads required from storage. Reading from storage is significantly slower than reading from memory. Therefore, Provisioning Servers should be allocated sufficient memory to maximize the benefit from this caching process. A 64-bit operating system should be used for Provisioning Services in order to maximize the amount of memory available.

The following formula can be used to determine the optimal amount

of memory that should be allocated to a provisioning server:

System Cache =

*512MB + (# of Active vDisks * Average Data Read From vDisk)*

Total Server RAM = Committed Bytes Under Load + System Cache

Note: A good starting point for determining the minimum amount of server memory is to allocate 2GB of RAM per active desktop vDisk and 10GB of RAM per active server vDisk. For most Provisioning Services implementations, the amount of RAM allocated to each provisioning server will fall between 8 and 32GB.

Decision: Provisioning Server Processor

The minimum requirements for the processor in a physical provisioning server:

- Intel or AMD x86 or x64 compatible processor
- Processor speed must run at a minimum of 2GHz. A 3GHz or better processor is preferred.

For virtual provisioning servers:

- Small environments (up to approximately 500 virtual machines) 2 vCPUs are recommended.
- Larger environments 4 vCPUs are recommended.

Decision: Scale Up or Out

As the farm grows, administrators will need to decide whether to add more resources to the provisioning servers or to add more provisioning servers to the farm.

There are a number of environmental factors that need to be considered when determining whether the Provisioning Servers should be scaled up or scaled out:

- **Redundancy** – Spreading user load across additional less-powerful servers helps reduce the number of users affected from

a single provisioning server failure. If the business is unable to accept the loss of a single high-specification server, consider scaling out.

- **Failover times** – The more target devices connected to a single provisioning server, the longer it will take for them to failover in the event that the server fails. Consider scaling out to reduce the time required for target devices to failover to another server.
- **Data center capacity** – The data center may have limited space, power and/or cooling available. In this situation, consider scaling up.
- **Hardware costs** – Initially, it may be more cost effective to scale up. However, there will be a point where scaling out actually becomes more cost effective. A cost analysis should be performed to make that determination.
- **Hosting costs** – There may be hosting and/or maintenance costs based on the number of physical servers used. If so, consider scaling up to reduce the long-term cost of these overheads.

Decision: Bandwidth Requirements

It is imperative that sufficient bandwidth is available for Provisioning Services to prevent network bottlenecks from affecting virtual desktop performance. Network utilization is most significant when target devices are booting. The following table shows the approximate amount of data that Provisioning Services requires to boot different operating systems:

Approximate boot data usage by OS

Operating System	Avg Boot Data Usage (MB)
Windows 8 x86	178
Windows 8 x64	227
Windows 7 x86	166
Windows 7 x64	210
Windows 2012	225
Windows 2012 R2	232
Windows 2008 R2	251
Windows Vista x86	190
Windows Vista x64	240

By default, Provisioning Services can boot 500 devices simultaneously. Any additional booting devices are paused until the number of current booting devices falls below the 500 limit. This number can be adjusted in the Provisioning Services console. In normal day to day operations it is unlikely that this many devices will need to be started concurrently, however, events such as a restart after a maintenance window or recovery from a power outage may cause a large number of desktops to boot simultaneously and temporarily saturate the network. Lowering the number of devices that can boot simultaneously can help to prevent this from occurring.

For example, 500 Windows 8 x86 target devices requiring 178MB requires:

$$500 \text{ Devices} * 178 \text{ MB} = 89 \text{ GB}$$

Determining how much time will be required to boot the target devices can be estimated using the following formula:

$$\text{Seconds to Boot} = \frac{(\text{Number of Targets} * \text{MB Usage})}{\text{Network Throughput}}$$

Booting 500 devices on a 1 Gbps network will require:

$$\text{If 1 Gbps Ethernet} = 125 \text{ MB/S}$$

Then,

$$\text{Seconds to Boot} = \frac{(500 \text{ Devices} * .178 \text{ GB})}{125 \text{ MB/S}}$$

$$\text{Seconds to Boot} = 712 \text{ Seconds} = 12 \text{ Minutes}$$

Booting 500 devices on a 10 Gbps network will require:

$$\text{If 10 Gbps Ethernet} = 1250 \text{ MB/S}$$

Then,

$$\text{Seconds to Boot} = \frac{(500 \text{ Devices} * .178 \text{ GB})}{1250 \text{ MB/S}}$$

$$\text{Seconds to Boot} = 71 \text{ Seconds} = 1.2 \text{ Minutes}$$

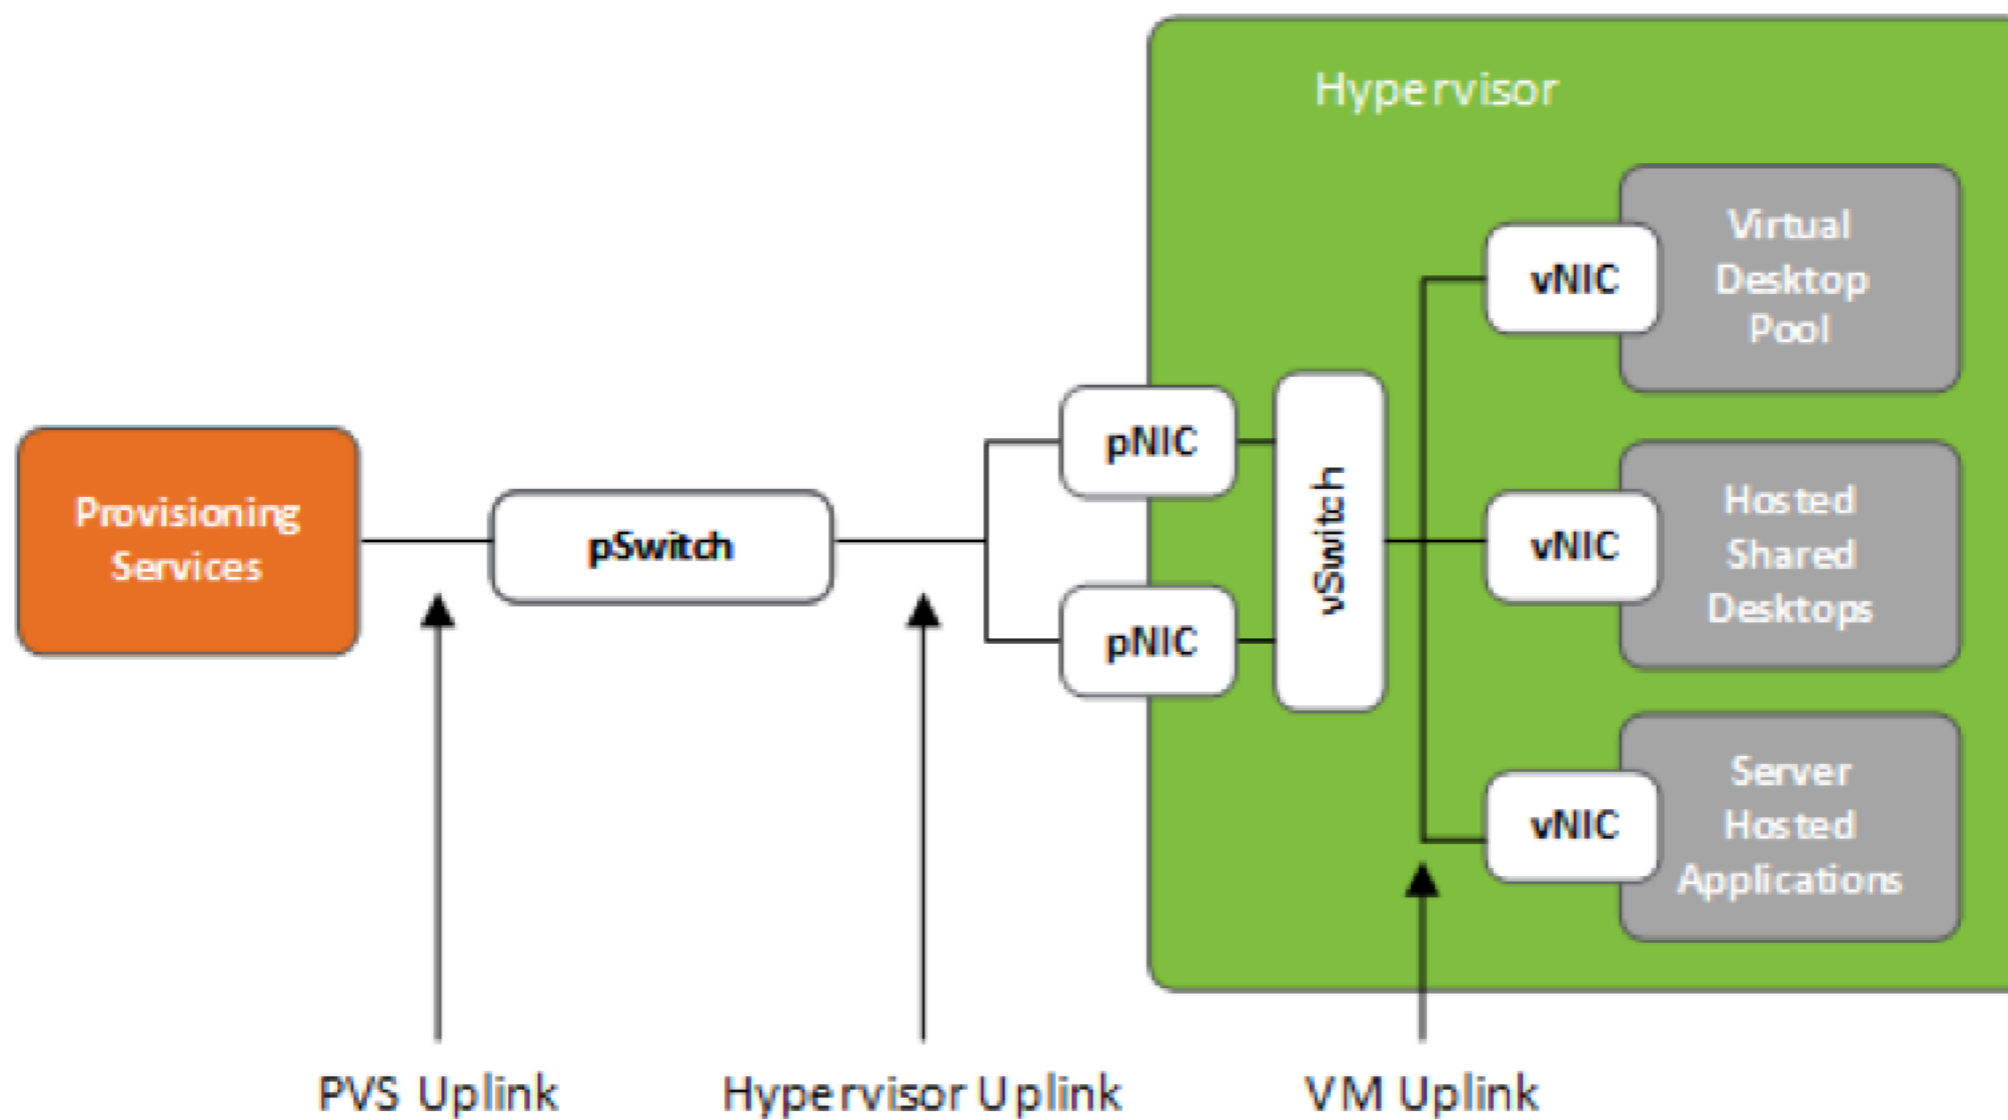
A 10Gbps network is recommended for use with Provisioning Services. If a 10Gbps network is not available, consider link aggregation to provide additional bandwidth to the provisioning servers, or a dedicated physical streaming network.

Tip: Firewalls can add latency and create bandwidth bottlenecks in Provisioning Services environments. If the use of firewalls cannot be avoided, refer to the Citrix whitepaper [CTX101810 – Communication Ports Used By Citrix Technologies](#), for the list of ports that should be enabled for full functionality.

Decision: Network Configuration

As mentioned before it is essential that the network is sized correctly to prevent network bottlenecks causing high disk access times and directly affecting virtual desktop performance. The following diagram outlines a common Provisioning Services network infrastructure:

Sample PVS Network Configuration



The following network configuration is recommended for the network sections outlined within the diagram:

- **PVS Uplink** – All disk access from the target devices will be transferred via the PVS network uplink. This means hundreds or even thousands of devices will use this network connection. Therefore it is vital that this connection is redundant and can failover without any downtime. Furthermore Citrix recommends a minimum bandwidth of 1Gbps per 500 target devices. For virtual provisioning servers a respective QoS quota or a dedicated physical network uplink should be configured to ensure best performance.
- **Hypervisor Uplink** – Used by all PVS target devices hosted on a particular hypervisor host. Therefore redundancy with transparent failover is strongly recommended. Unless the target devices run a very I/O intensive workload or perform I/O intensive tasks (e.g. booting) simultaneously, a bandwidth of 1Gbps is sufficient for this uplink.
- **VM Uplink** – All network traffic for a virtual machine, including PVS streaming traffic, will traverse this virtual network connection. Unless the workload is extremely I/O intensive a bandwidth of 100 Mbps is sufficient to handle even peak loads during I/O intensive tasks, such as booting from vDisk. For example a Windows 2012 R2 Server will read approximately 232MB during a period of 90 seconds from the vDisk until the Windows Logon Screen is shown. During this period an average data rate of 20.5 Mbps with peaks up to 90 Mbps can be observed.

The following switch settings are recommended for Provisioning Services:

- **Disable Spanning Tree or Enable Portfast** – In a switching environment the Spanning Tree Protocol (STP) places ports into a blocked state while it transmits Bridged Protocol Data Units (BPDUs) and listens to ensure the BPDUs are not in a loopback

configuration. The port is not placed in a forwarding state until the network converges, which depending on the size of the network, may incur enough time to cause Pre-boot Execution Environment (PXE) timeouts. To eliminate this issue, disable STP on edge-ports connected to clients or enable PortFast.

- **Storm Control** - Storm Control is a feature available on Cisco switches that allows a threshold to be set whereby, multicast, broadcast, or unicast traffic may be suppressed. Its purpose is to prevent malicious or erroneous senders from flooding a LAN and affecting network performance. PVS Servers may send a large amount of traffic by design that falls within a storm control threshold, therefore the feature should be configured accordingly.
- **Broadcast Helper** – The broadcast helper is required to direct broadcasts from clients to servers that would otherwise not be routed. In a PVS environment it is necessary to forward PXE boot requests when clients are not on the same subnet as the servers. If possible the recommended network design is to have PVS servers residing on the same subnet as the target devices. This mitigates the risk of any service degradation due to other networking infrastructure components.

Note: For more information on PVS networking best practices please refer to [Best Practices for Configuring Provisioning Services Server on a Network](#).

Decision: Network Interfaces

Teaming multiple network interfaces with link aggregation can provide greater throughput for PVS and Hypervisor uplinks, increasing network performance and helping to alleviate potential bottlenecks. Teaming NICs also prevents them from becoming a single point of failure. The table below lists provisioning server interface options, from most to least capacity:

Provisioning Service NIC Speeds

Interface	Bandwidth	Description
Chassis 10 Gbps+	80 Gbps+	PVS servers, targets, and infrastructure servers hosted within the same blade chassis. Offers the maximum potential bandwidth across a common backplane.
Bonded/Teamed 10 Gbps	Up to 80 Gbps	10 Gbps networks are common in many datacenters. A single network interface however creates a single point of failure and is not recommended.
Single 10 Gbps	10 Gbps	10 Gbps networks are common in many datacenters. A single network interface however creates a single point of failure and is not recommended.
Bonded/Teamed 1 Gbps	Up to 8 Gbps	When a 10 Gbps interface is not available Bonding/Teaming is an ideal option to improve server bandwidth by using multiple interfaces without the added complexity of multi-homing.
Single 1 Gbps	1 Gbps	Functions for small environments, but 1 Gbps targets quickly overrun 1 Gb servers. A single network interface creates a single point of failure and is not recommended.

Note: For Provisioning Services implementations on low bandwidth networks (1Gbps or slower), performance may be improved by isolating streaming traffic from other network traffic on the LAN.

Note: Microsoft does not support NIC teaming with Hyper-V on Windows Server 2008 R2; however, third party solutions are available. Microsoft does support NIC teaming with Hyper-V on Windows Server 2012/2012 R2. All support queries regarding teaming with Hyper-V should be directed to the NIC OEM.

The following network interface features should be taken into consideration when selecting a network interface for Provisioning Services:

- **TCP Offloading** – Offloading I/O tasks to the network interface reduces CPU usage and improves overall system performance, however, PVS Streaming Services can be negatively impacted when Large Send Offload is enabled due to the extra work placed on the network adapter. Many network adapters will have Large Send Offload and TCP checksum offload enabled by default.

Note: If Large Send Offload is enabled and the switch that the

traffic is passing through does not support the frame size sent by the Large Send Offload engine, the switch will drop the frame causing data retransmission. When retransmitting, the operating system will segment the frames instead of the network adapter, which can lead to severe performance degradation.

- **Receive Side Scaling (RSS)** – Receive side scaling enables packets received from a network adapter to be balanced across multiple CPUs which allows incoming TCP connections to be load balanced, preventing bottlenecks from occurring to a single CPU. In Windows Server 2008 R2 and Windows Server 2012/2012 R2, RSS is enabled by default.

Decision: Subnet Affinity

The Provisioning Services Subnet Affinity is a load balancing algorithm that helps to ensure target devices are connected to the most appropriate provisioning server. When configuring subnet affinity the following options are available:

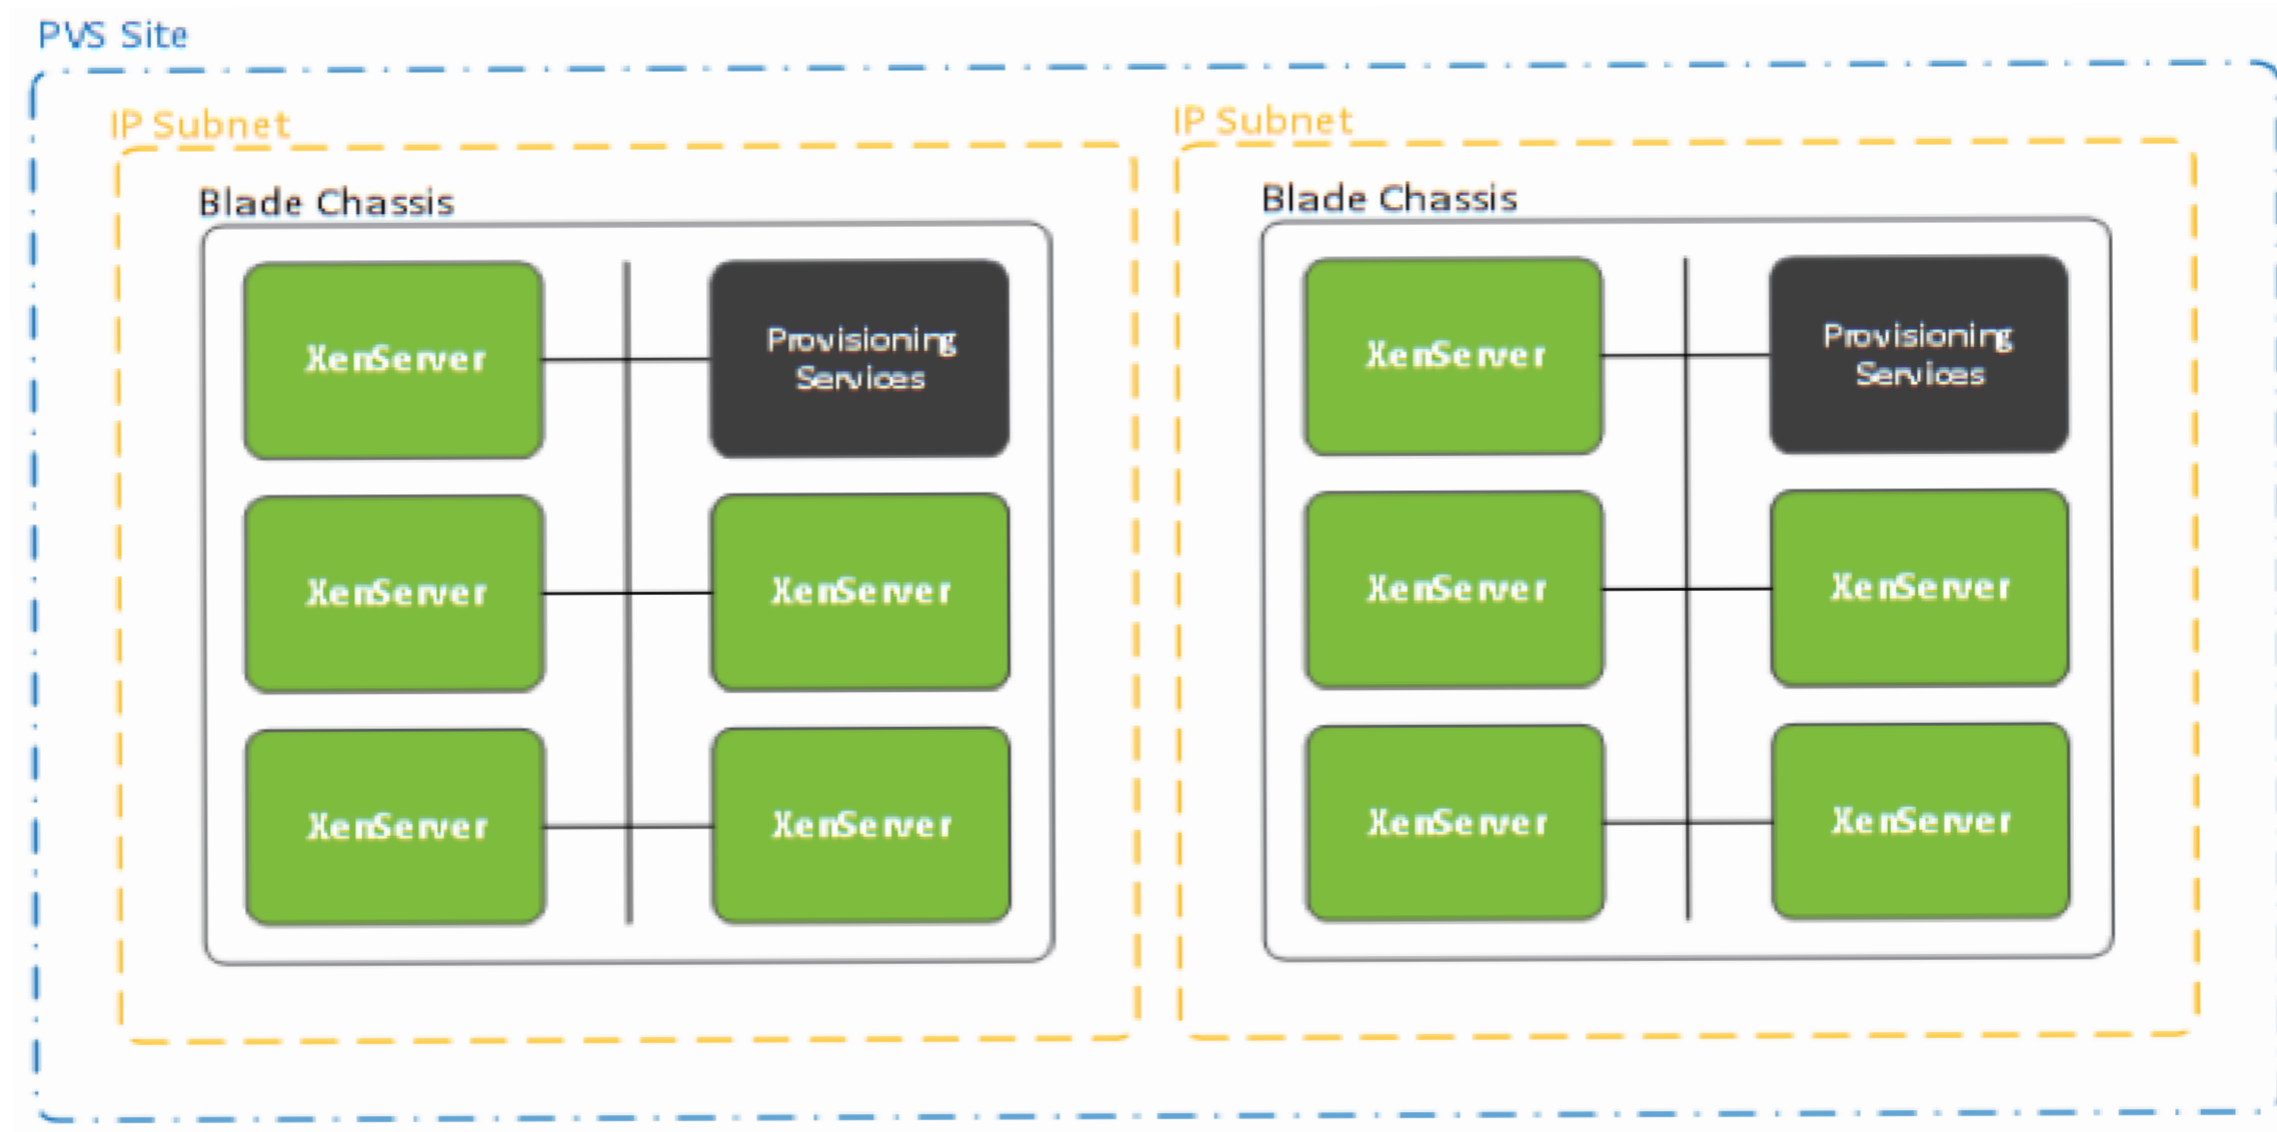
- **None** – Ignore subnets; uses the least busy server.
- **Best Effort** – Uses the least busy server/NIC combination from within the same subnet. If no server/NIC combination is available within the subnet, select the least busy server from outside the subnet. If more than one server is available within the selected subnet, perform load balancing between those servers. This is the default setting.
- **Fixed** – Use the least busy server/NIC combination from within the same subnet. Perform load balancing between servers within that subnet. If no server/NIC combination exists in the same subnet, do not boot target devices assigned to this vDisk.

The following examples show common network configurations for physical provisioning servers. Similar configurations can be implemented for virtual provisioning servers without compromising on performance or functionality.

Blade Design

The provisioning servers and the target devices that they support reside within the same chassis. In most cases, the chassis will have a dedicated 10Gbps switch shared among all blade servers within the chassis.

PVS Blade Enclosure Design

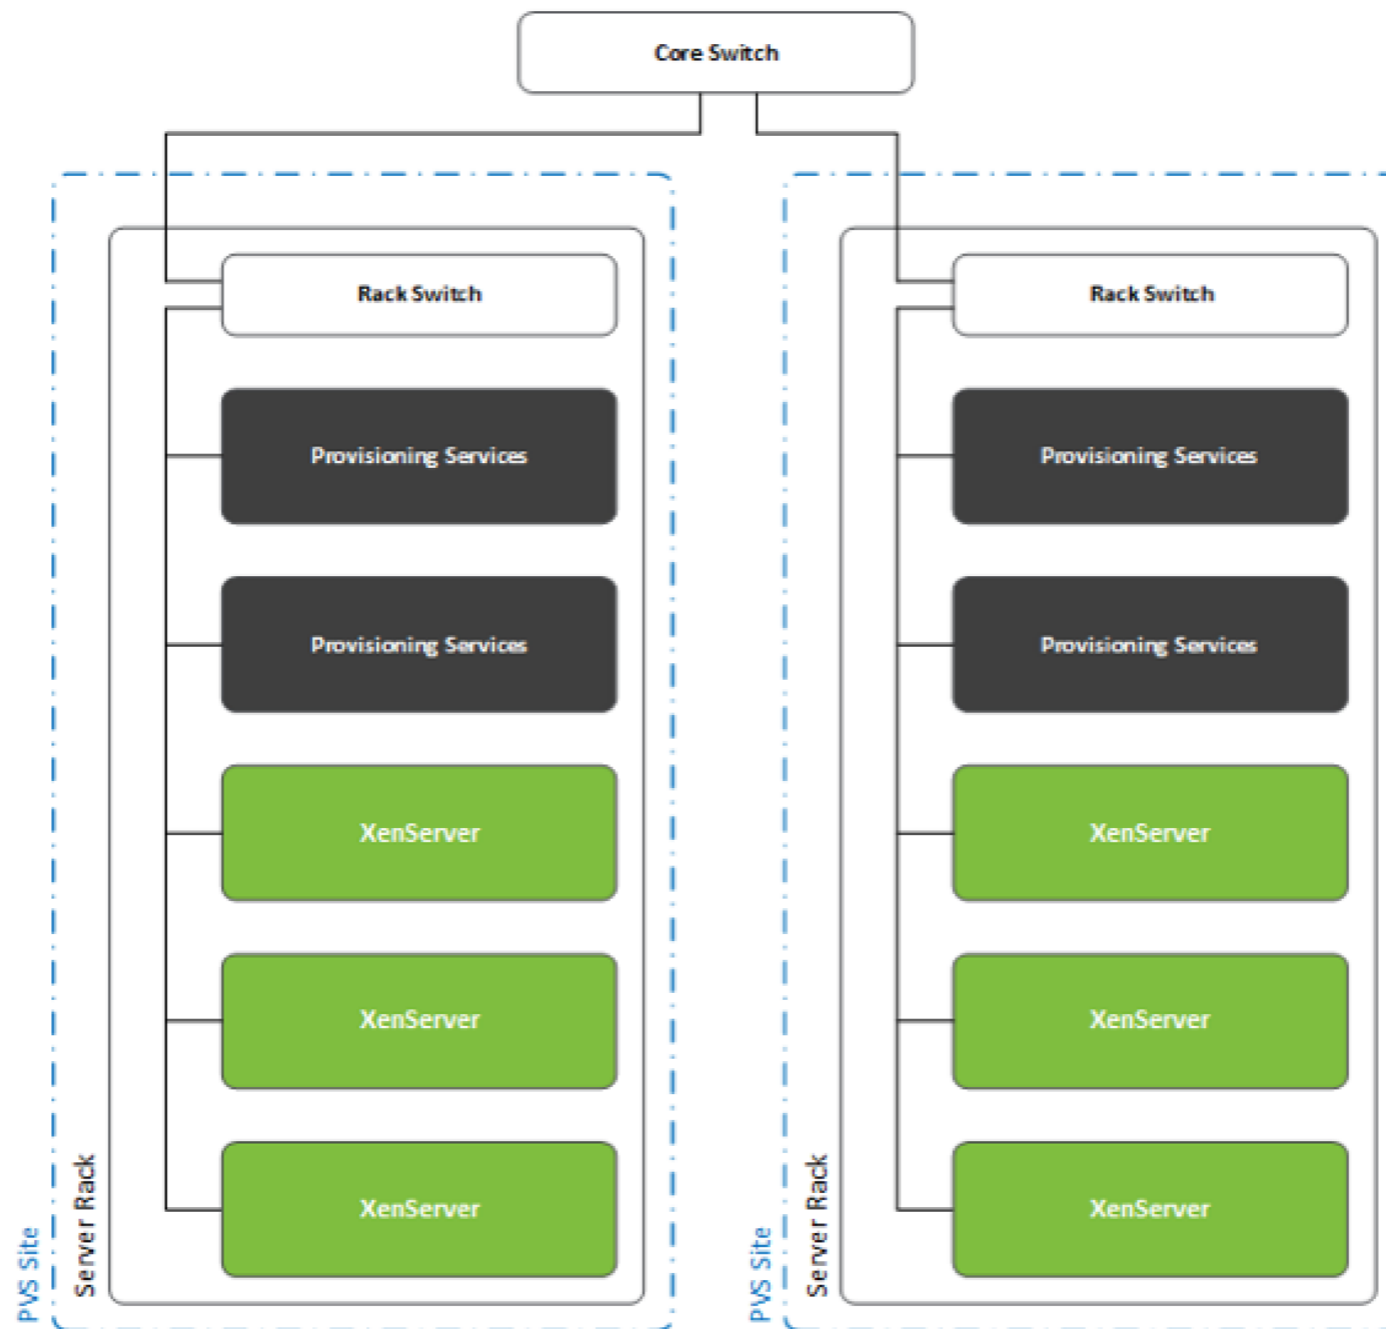


The “Best Effort” subnet affinity option is used to keep Provisioning Services traffic within the same chassis. Should the provisioning server become unavailable, the targets will failover to the second provisioning server in the second chassis, but same Provisioning Services site.

Rack Design

The second example is based on a rack design that uses rack switches to keep the provisioning traffic within the rack.

PVS Rack Design



As opposed to the blade chassis design, the subnet affinity feature is not used. Instead a Provisioning Services site with two provisioning servers will be configured per server rack. This will ensure that the target devices are streamed from provisioning servers within the same rack.

Experience from the Field

Manufacturing – A manufacturing company is designing a Provisioning Services solution to support five thousand virtual desktops. The company has concerns that Provisioning Services streaming traffic will create a bottleneck on the network affecting other applications. The company chose to build the environment on blade servers so that provisioning traffic is contained within the blade enclosure and will not impact other traffic on the network.

Decision: Streams per Server

The number of streams that a Provisioning Services server can run concurrently can be determined by the following formula:

$$\text{Max Number of Streams} = \# \text{ of Ports} * \# \text{ of Threads/Port}$$

By default the Streaming Service is configured with 20 ports, and 8 threads per port. Therefore, by default, a provisioning server can support 160 concurrent targets. If the environment needs to support more than 160 concurrent targets, the number of ports, and threads per port can be adjusted in the Provisioning Services console. Citrix testing has shown that best performance is attained when the threads per port is not greater than the number of cores available on the provisioning server. If the provisioning server does not have sufficient cores, the server will show a higher CPU utilization, and target devices waiting for requests to be processed will have a higher read latency.

The following table depicts an example of the number of possible streams available by various processor types when adjusting the number of PVS ports and PVS threads/port:

Example of determining max number of PVS streams

Processor Type	# of Processors	Threads	PVS Ports	PVS	Max Number of Streams
Quad-Core CPU with Hyperthreading	2	16	20	16	320
Quad-Core CPU	2	8	48	8	384
Dual-Core CPU	4	8	60	8	480
Single-Core CPU with Hyperthreading	4	8	60	8	320

Decision: Auditing

Provisioning Services provides an auditing tool that records configuration actions on components within the Provisioning Services farm. Auditing provides administrators with the ability to troubleshoot and monitor recent changes that might impact system performance and behavior. The Provisioning Services administrator privileges determine the audit information that can be viewed and the menu options that are visible.

By default the auditing feature is disabled. When enabled, the audit trail information is written to the Provisioning Services database along with general configuration data. The data written to the database is not archived automatically and will continue to grow indefinitely until archived manually.

Objects that can be audited include tasks performed on the farm, site, provisioning servers, collections, target devices, store, and vDisks. Only tasks, which are performed through the admin console, PowerShell, SOAP server, or the command line interface are audited. Since auditing is capturing user initiated actions it will have no impact on the performance of streaming services. Tasks that are not performed by these methods are not audited, such as booting target devices.

For organizations with high security concerns, enable auditing for tracking purposes. Ensure that archiving is performed on the

database at regular intervals in order to prevent indefinite growth.

Decision: Antivirus

By default, most antivirus products scan all files and processes, which will have a significant impact on performance. For optimal performance, antivirus software should be optimized for a Provisioning Services environment. For more information, please refer to CTX124185 – [Provisioning Services Antivirus Best Practices](#).

Antivirus software can cause file-locking issues on provisioning servers by contending with files being accessed by Provisioning Services. The vDisk Store and write cache should be excluded from any antivirus scans in order to prevent file contention issues.

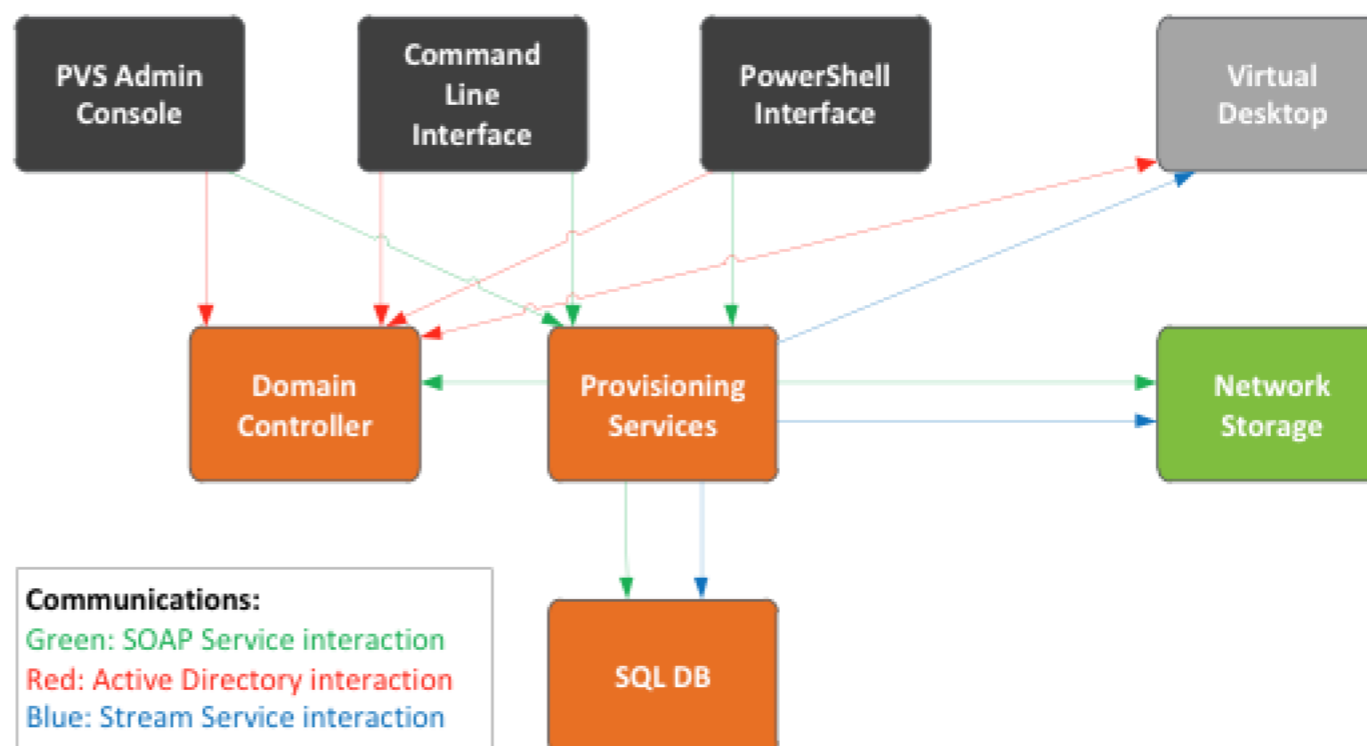
When a virtual disk is running in standard mode and needs to be restarted, it downloads all of the previously loaded virus definitions. This can cause performance degradation when restarting several target devices at a time, often causing network congestion while

the operation persists. In extreme cases, the target device and provisioning server can become sluggish and consume more resources than necessary. If the antivirus software supports it, definition files should be redirected to the write cache drive so that they are preserved between reboots.

Decision: Service Accounts

The Stream and SOAP services interact with many different components within the Provisioning Services architecture, as shown in the diagram below.

Provisioning Services Architecture



The Stream and SOAP services can be run in one of two different contexts:

- **Specified user account** – High security organizations should use a specified user account that is closely controlled and monitored. If the password of this account expires, Provisioning Services will not function until the password has been changed or reset.

Note: If the Provisioning Services store will be placed on a Windows file share, the Stream and SOAP services must be configured to run using a specified user account, which must be granted full permission to all files in the store. If the Stream and SOAP services are not configured as user accounts, Provisioning Services will be unable to access vDisk files on the remote server.

- **Local system account** – Running the Stream and SOAP services under the context of a local account such as the Network Service account is easy to configure, and requires no additional management.

Hardware Layer

The hardware layer is responsible for the physical devices required to support the entire solution including servers, processors, memory and storage devices. This layer is broken into three groups focused on providing the necessary resources for specific parts of the entire solution. One group of servers will support the XenApp (shared) components (if applicable). A second group of servers will support the XenDesktop (VDI) components (if applicable). A final group of servers will support the underlying infrastructure for the entire environment, which is identified as the Control Layer.

[Click here to provide feedback](#)

Hardware Sizing

This section covers hardware sizing for the virtual infrastructure servers, virtual desktops, and virtual application hosts. The sizing of these servers is typically done in two ways. The first and preferred way is to plan ahead and purchase hardware based on the workload requirements. The second way to size the hosts is to use existing hardware in the best configuration to support the different workload requirements. This section will discuss decisions related to both methods.

Decision: Workload Separation

When implementing a XenApp and XenDesktop deployment, the workloads for the infrastructure, XenDesktop, and XenApp workloads can be separated into dedicated resource clusters or mixed on the same physical hosts. Citrix recommends using resource clusters to separate the workloads, especially in an enterprise deployment. This allows better host sizing as each workload has unique requirements such as overcommit ratios and memory usage.

In smaller environments where resource clusters are cost prohibitive, the workloads may be mixed in a manner which still allows for a highly available environment. Citrix leading practice is to separate the workloads however mixed workloads is a cost based business decision. Mixed workloads will be covered in the following sections, however for a validated mixed workload design reference the [Mobilizing Windows Apps Design Guide](#).

Decision: Control Host Sizing

The required infrastructure for the control layer requires adequate resources in order to scale up and support the overall environment. Citrix recommends virtualizing the infrastructure required for XenApp and XenDesktop. The VM resource allocation for each control component should have been determined in the [control](#)

[layer](#). The infrastructure should be given adequate processing power to support the environment and therefore a 1 vCPU to 1 pCPU configuration is recommended. The underlying physical hosts for the infrastructure require less memory than the desktop and applications hosts as they will typically host fewer virtual machines as a result of no CPU overcommit.

For a highly available infrastructure, the access, desktop, and image controllers should be hosted on a minimum of two physical hosts. As the environment scales, the N+1 model should be used for additional controllers spread over physical hosts in the environment.

Desktop Host Sizing

The physical resources required for the VDI and SBC FlexCast models are based on workload characteristic of each user group which was first identified during the [user assessment](#) phase and then used to properly size the virtual machines earlier in the desktop [resource allocation](#). The decisions are broken down by physical resource below.

Decision: pCPU

The sizing of physical CPU cores differs for XenDesktop (VDI) workloads and XenApp (SBC) workloads due to the difference in CPU overcommit ratios and the number of virtual machines required to host user workloads.

The following table provides guidance on the number of virtual desktops that can be supported for light, medium and heavy workloads per physical core. Each desktop correlates to a single concurrent user.

Processor Requirements by XenDesktop Workload (VMs/ Core)

Workload ^{1,2}	Pooled VDI/Assigned VDI ^{3,4}	
	Windows 8	Windows 7 ⁵
Dual Socket		
Light	15	13
Medium	11	10
Heavy	6	5
Quad Socket⁶		
Light	13	11
Medium	10	8
Heavy	5	4

¹Processor architecture and speed have a direct impact on the number of users that can be supported. The estimates provided are based on a processor speed of 2.7 GHz and Intel Ivy Bridge processor architecture.

²The scalability estimates provided account for the performance benefits of hyper-threading. Hyper-threading has the potential to improve user density per VM (SBC) or VM density per host (SBC and VDI). The use of hyper-threading typically provides a performance boost of between 20-30%, but may also have an adverse effect on performance in rare cases. Therefore, it is recommended to monitor and compare the performance of workloads with and without hyper-threading.

³Personal vDisk (PVD) technology allows users to perform advanced customizations, while still providing administrators with the benefits of central image management. However, Personal vDisks incur a processor performance overhead, which varies according to the extent that they are used. Testing shows that user density can be reduced by 14% when Personal vDisks are used.

⁴Antivirus has a significant CPU overhead. The overhead can be safely estimated by reducing the user density by 25% although for more accurate estimates refer to your anti-virus vendor.

⁵With Aero disabled.

⁶The move from dual to quad sockets is not linear and has been accounted for by a 15% drop in user density on a per core basis.

To estimate the total number of physical cores required for the XenDesktop workload, use the following formula:

$$\text{Total Cores XenDesktop}^* = \frac{1}{1 - (V + M)} \sum_i \frac{a_i + P(b_i)}{W_i}$$

*This equation can be used based on either dual core machines or quad core

Σ represents the sum of all user group combinations i. (Light/Medium/Heavy on each OS)

a_i = Number of users per user group i

b_i = Number of users per PvD user group i

W_i = Number of users per core per user group i (from above table)

P = Impact of using PvD = 1.15

V = Antivirus impact (default = 0.25)

M = Monitoring tools impact (default = 0.15)

Sizing Example

A customer is deploying 200 non-persistent Windows 7 desktops to users which have been determined to be 'medium' users and 50 assigned Windows 8 desktops with PvD determined to be 'heavy'. The customer does not have any data on the impact of the AV and monitoring products they are using and therefore are using the default values. The estimated number of cores required if using dual socket hosts is:

$$\frac{1}{1 - (0.25 + 0.15)} * \left(\frac{200}{10} + \frac{1.15 * 50}{6} \right) = 50 \text{ Cores (Dual Socket)}$$

The number of users hosted per XenApp server will vary based on the number of vCPUs allocated to the virtual machines as outlined in the [resource allocation section](#). As previously mentioned, the vCPU allocation should be determined based on the number of NUMA (Non-Uniform Memory Access) nodes that the processor has. The vCPU assignment should be a factor of the NUMA nodes in order to avoid what is known as "NUMA thrashing". For example

if a physical host has a dual 12 core processor with 2 x 6 core NUMA nodes, the XenApp machines on that host should not be allocated more than 6 vCPUs. Allocating more than 6 vCPU can lead to the machine having to access memory across NUMA node. Although modern hypervisors will try and limit the effects of VMs placed across multiple NUMA nodes, it is a Citrix best practice to size appropriately taking NUMA into account. You can check with the processor vendor or within each hypervisor for the number of NUMA nodes.

To determine the number of XenApp machines required, the CPU overcommit ratio must be decided. Depending on the workload, an overcommit ratio of 1.5-2 has shown to result in the highest scalability although testing is required to determine the optimal number for each environment.

Once the virtual machine has been optimally sized, guidance on the expected number of users on those machines per workload per physical core can be referenced below.

Processor Requirements by XenApp Workload

Workload ^{1,2}	XenApp Hosted Shared Desktop ³	
	Windows 2012	Windows 2008 R2
Dual Socket		
Light	21	18
Medium	14	12
Heavy	7	6
Quad Socket⁴		
Light	18	15
Medium	12	10
Heavy	6	5

¹Processor architecture and speed have a direct impact on the number of users that can be supported. The estimates provided are based on a processor speed of 2.7 GHz and Intel Ivy Bridge processor architecture.

²The scalability estimates provided account for the performance benefits of hyper-threading. Hyper-threading has the potential to improve user density per VM (SBC) or VM density per

host (SBC and VDI). The use of hyper-threading typically provides a performance boost of between 20-30%, but may also have an adverse effect on performance in rare cases. Therefore, it is recommended to monitor and compare the performance of workloads with and without hyper-threading.

³Antivirus has a significant CPU overhead. The overhead can be safely estimated by reducing the user density by 25% although for more accurate estimates refer to your anti-virus vendor.

⁴The move from dual to quad sockets is not linear and has been accounted for by a 15% drop in user density on a per core basis.

To estimate the total number of physical cores required for the XenApp workload use the following formula:

$$\text{Total Cores XenApp}^* = \frac{1}{1 - (V + M)} \sum_i \frac{a_i}{W_i}$$

*This equation can be used based on either dual core machines or quad core

Σ represents the sum of all user group combinations i. (Light/Medium/Heavy on each OS)

a_i = Number of users per user group i

W_i = Number of users per core per user group i (from above table)

V = Antivirus impact (default = 0.25)

M = Monitoring tools impact (default = 0.15)

Sizing Example

A customer is deploying 500 Windows 2012R2 desktops to users which have been determined to be 'light'. The customer has spoken to their AV vendor and have estimated the scalability impact to be 20%. The customer does not use any monitoring within the environment and therefore the monitoring impact is 0. The estimated number of cores required if using quad socket hosts is:

$$\frac{1}{(1 - 0.20)} * \left(\frac{500}{18}\right) = 35 \text{ Cores (Quad Socket)}$$

[Click here to provide feedback](#)

Decision: pRAM

The recommended method for sizing memory to a physical VDI or SBC host is to size based on the total memory required to support the machines and the CPU capacity of the host. In order to calculate the total memory required for XenDesktop, simply multiply the number of users/desktops (1 user to 1 desktop) in a user group by the amount of memory allocated to that desktop. The sum of all of the user groups will be the total RAM required. This is shown in the formula below.

$$\text{Total RAM XenDesktop} = \sum_i a_i * R_i$$

a_i = Number of concurrent users per user groups (Light/Medium/Heavy on each OS)

R_i = vRAM assigned to each VDI machine as determined in the [resource allocation](#) section'

For XenApp the total memory is calculated by multiplying the number of XenApp machines by the amount of RAM that each will be allocated. These should be determined using the guidance in the resource allocation section and then recalculated based on any vCPU allocation changes to account for NUMA.

$$\text{Total RAM XenApp} = \sum_i X_i * R_i$$

X_i = Number of XenApp virtual machines per OS (2008R2 / 2012)

R_i = vRAM assigned to each XenApp machine (if the vCPU allocation has changed due to NUMA, a good rule of thumb is to assign 3GB of vRAM for each 1 vCPU on a XenApp VM)

Once the total required memory has been calculated to support the workloads the total memory per host will be calculated based on the number of hosts. The number of hosts should be determined based on CPU core requirements. To calculate the RAM required per host use the following formula for either XenApp or XenDesktop:

Total RAM per Host =

$$\text{Hypervisor RAM} + \frac{\text{Total RAM Calculated Above}}{\text{Number of Hosts Calculated by Core Requirement}}$$

Sizing Example

A customer is deploying 500 Windows 2012R2 desktops to users which have been determined to be 'light'. The customer has spoken to their AV vendor and have estimated the scalability impact to be 20%. The customer does not use any monitoring within the environment and therefore the monitoring impact is 0. The estimated number of cores required if using quad socket hosts is:

$$\frac{1}{(1 - 0.20)} * \left(\frac{300}{18}\right) = 21 \text{ Cores (Quad Socket)}$$

The customer decided to purchase a quad socket system with 24 total cores. The customer needs a single server to support the workload. Each XenApp machine will be assigned 6 vCPU due to the 6 core NUMA nodes on the server and 18GB of RAM. The customer will deploy 8 XenApp machines to support the users. The memory required is:

$$8 \text{ VMs} * 18 \text{ GB/VM} = 144\text{GB}$$

The memory per host required is $144\text{GB} + 8\text{GB} = 154\text{GB}$ for the hypervisor since only a single host is required. The final configuration would then be a host with 24 Cores and 192GB of RAM. Using the N+1 method for high availability, an additional host with the same specifications will be deployed.

Note: 192GB is a standard memory configuration (12 sticks x 16GB). Although now there is excess memory on the host, this can be allocated to the VMs which can be used for additional caching and increased RAM Cache if using PVS.

Decision: GPU

Hosts used to deliver graphical workloads require graphics processors to deliver a high end user experience. Specific hardware hosts and graphics cards are required to support high end graphics using HDX 3D Pro. An updated list of tested hardware is available in a [knowledge base article](#). Sizing of the desktop and application hosts of high end graphics users should be based on the GPU requirements ensuring that the host then has adequate CPU and memory resource to support the workload.

NVIDIA GRID cards can be leveraged with vGPU profiles to support multiple users. Sizing guidelines are provided from NVIDIA in the following table.

NVIDIA Sizing Guidelines

NVIDIA GRID Graphics Board	Virtual GPU Profile	Application Certifications	Graphics Memory	Max Displays Per User	Max Resolution Per Display	Max Resolution Per Graphics Board	Use Case
GRID K2	K260Q	✓	2,048 MB	4	2560x1600	4	Designer/Power User
	K240Q	✓	1,024 MB	2	2560x1600	8	Designer/Power User
	K220Q	✓	512 MB	2	2560x1600	16	Designer/Power User
	K200		256 MB	2	1900x1200	16	Knowledge Worker
GRID K1	K140Q	✓	1,024 MB	2	2560x1600	16	Power User
	K120Q	✓	512 MB	2	2560x1600	32	Power User
	K100		256 MB	2	1900x1200	32	Knowledge Worker

Decision: Number of Hosts

The number of physical hosts required will depend on several factors. In a small deployment with limited budget, the infrastructure and user workloads may reside on the same physical hosts. In an enterprise deployment, the infrastructure should reside on dedicated hosts separate from the user workloads.

Based on the physical host being used, the number of hosts will be dependent on the resource bottleneck. If sized properly, the resource usage should be maximized so that no single resource becomes the limiting factor. For example, a desktop host with a known number of desktops based on a CPU constraint should be given only the amount of memory needed to support the hypervisor, the desktops, and a small buffer to mitigate risk. Testing tools can be used to help identify any bottlenecks not only on the hosts with CPU, memory, and local storage, but with SAN storage and network bandwidth as well. Popular load test tools include [LoginVSI](#) and [HP LoadRunner](#).

Once bottlenecks are accounted for and the hosts are properly sized, it is important to have enough capacity for a hardware failure. Use the N+1 approach to ensure that no single hardware failure can adversely affect the production environment. The control layer should reside on a minimum of two hosts to ensure high availability. If additional control machines are required, they should be spread out among physical servers so that a single failure will not impact the environment.

You can use the following formula to estimate the number of hosts required for the user workloads. Note that it is based on the best practice of separating the XenApp and XenDesktop workloads due to the different recommended CPU overcommit ratios for each.

$$\text{Resource Hosts} = \left(\frac{\text{Cores Required XenDesktop}}{\text{Cores per Host Hardware}} + 1 \right)^* + \left(\frac{\text{Cores Required XenApp}}{\text{Cores per Host Hardware}} + 1 \right)^*$$

[Click here to provide feedback](#)

*Should be rounded up to ensure capacity (i.e. 5.2 + 3.3 should be rounded to 6+4=10)

From the Field: Hardware Sizing Example

A financial planning company has assessed their user base and come up with a user breakout. After reviewing the various user requirements, they determined that the users will be assigned the following FlexCast models:

User Group	Number of Users	FlexCast Model
Light Users	200	XenApp Hosted Shared Desktop on Server 2012 R2
Medium Users	400	XenDesktop Pooled VDI on Windows 7 SP1 x64
Heavy Users	50	XenDesktop Assigned VDI using Personal vDisk on Windows 7 SP1 x64

As per the corporate standard, all machines will have both anti-virus and monitoring tools running. The scalability impact of their anti-virus solution is estimated at 20%, while the monitoring tool impact has been estimated at 10%.

The company has decided to plan for dual socket machines for the better per core user density. Using the above estimates results in a total core estimate of:

$$\text{Total Cores XenDesktop} = \frac{1}{1 - (.20 + .10)} \left(\frac{400}{10} + \frac{50 * 1.15}{6} \right) = 71 \text{ Cores}$$

$$\text{Total Cores XenApp} = \frac{1}{1 - (.30)} \left(\frac{250}{21} \right) = 17 \text{ Cores}$$

In the resource allocation section, the vCPU assignment for each XenApp server was chosen as 8 following the guidance provided. Checking with the vendor for the hardware, there are two 8 core NUMA nodes on each host leaving an 8 vCPU assignment valid. The final XenApp server configuration will then be:

- 4 XenApp servers with 8 vCPUs and 24GB of vRAM

Also as per the resource allocation section, the VDI VMs were determined to have the following configurations:

- 400 XenDesktop VMs with 2 vCPUs and 3GB of vRAM
- 50 XenDesktop VMs with 4 vCPUs and 6GB of vRAM

So the total memory requirements are:

$$\text{Total RAM XenDesktop} = (400 * 3) + (50 * 6) = 1500 \text{ GB}$$

$$\text{Total RAM XenApp} = 4 * 24 = 96 \text{ GB}$$

If purchasing 16 core hosts, the total number of hosts required is:

$$\text{Hosts Required} = \left(\frac{71}{16} + 1\right) + \left(\frac{17}{16} + 1\right) = 6 + 3 = 9 \text{ Hosts}$$

In order to size the memory of each host, exclude the server reserved for high availability as there should be enough memory available in case of a single failure. Therefore the total memory per hosts are:

$$\text{XenDesktop Hosts RAM} = 8 \text{ GB} + \left(\frac{1500}{5}\right) = 308 \text{ GB per Host}$$

$$\text{XenApp Hosts RAM} = 8 \text{ GB} + \left(\frac{96}{1}\right) = 104 \text{ GB per Host}$$

The RAM should be rounded up to the nearest standard configuration for a server so in this case the final host configurations will be:

- 2 XenApp Hosts with 16 Cores and 128GB of RAM
- 6 XenDesktop Hosts with 16 Cores and 384GB of RAM

Note: An alternative configuration of 7 XenDesktop Hosts with 16 cores and 256 GB of RAM is possible with a potential to add more desktops in the future as CPU would not be the limiting factor. This might be done based on cost reasons and to better utilize memory, but can be determined by estimating the needs prior to procurement. It might be also worthwhile in this situation to further investigate the sizing if there is no plans for growth as the number of XenApp hosts is just over the threshold (17 cores vs 16 on host) based on the estimates.

[Click here to provide feedback](#)

Hypervisors

Hyper-V 2008 R2

This chapter provides design guidance for Citrix XenDesktop 7 deployments that leverage Microsoft Hyper-V 2008 R2.

Decision: Host Hardware

The hardware selected for the Hyper-V hosts will have a direct impact on the performance, scalability, stability and resilience of the XenDesktop solution. As such, the following key design decisions must be considered during the hardware selection process:

- **Hardware Compatibility:** For reasons of stability, it is imperative that the hardware selected is compatible with Hyper-V 2008 R2. The [Microsoft Windows Server Catalog](#) website maintains a list of servers and hardware components capable of supporting Hyper-V 2008 R2. Servers and hardware components with the “Certified for Windows Server 2008 R2” and “Certified for Windows Server 2008” logos pass the Microsoft standards for supporting Windows Server 2008 R2.
- **Processor Specification:** The Hyper-V design must ensure that the processor specification selected is capable of supporting Hyper-V:
 - 64-bit Intel VT or AMD-V processors with a clock speed of at least 1.4GHz, however 2GHz or faster is recommended.
 - Hardware Data Execution Protection (DEP), specifically Intel XD bit (execution disable bit) or AMD NX bit (no execute bit), must be enabled and available. For more information, please refer to Microsoft knowledge base article KB912923 – [How to determine that hardware DEP is available and configured on your computer.](#)
 - For improved performance and scalability, ensure that

Second Level Address Translation (SLAT) capable processors are selected. Intel calls this technology Extended Page Tables (EPT) and AMD calls it Nested Page Tables (NPT). The technology provides hardware assistance for memory virtualization, memory allocation, and address translation between the physical and virtual memory. The reduction in processor and memory overhead improves scalability allowing a single Hyper-V host to run more virtual machines.

Note: Live Migration allows administrators to move virtual machines between physical hosts. In order to use this feature, all physical host servers in a cluster must have processors from the same family. Live migration will not work from a host with an Intel based processor to a host with an AMD based processor. There is a circumstance in which live migration may fail even when processors are from the same family. For example, an application may be aware of CPU features it discovered on the host server that are no longer available after it has been migrated to a new host. In these situations, Hyper-V processor compatibility mode should be used to hide the processor features from the application. For more information, please refer to Microsoft TechNet article – [When to use Processor Compatibility Mode to Migrate Virtual Machines](#).

- **Memory Capacity:** The hardware selected for the Hyper-V hosts must have sufficient memory to support the parent partition as well as the guest virtual machines that they support:
 - Virtual desktops typically require between 768MB and 4GB of RAM depending on workload and operating system used. For more information, please refer to the [resource requirements](#) chapter.
 - Each hosted shared desktop typically requires between 200MB and 1GB of RAM per user, depending on their workload. For more information, please refer to the

[resource requirements](#) chapter.

Note: With the exception of Microsoft Windows Server 2008 R2 Standard edition, which is limited to 32GB, Hyper-V can support up to a total of 1TB of physical memory.

Decision: Host Scalability

During the design, it is necessary to estimate the number of physical host servers that will be required to support the XenDesktop implementation. Ideally, scalability testing should be performed prior to the hardware being ordered, however this is not always possible. In these situations, consider the following guidelines when determining single host scalability:

- The parent partition typically utilizes between 8 and 12% of the processor resources available leaving between 88 and 92% available for the child partitions.
- At a minimum, the parent partition requires 512MB of RAM, however 2GB or greater is recommended for VDI environments.
- Microsoft supports up to 12 virtual processors per logical processor for Hyper-V 2008 R2 hosts supporting Windows 7. A maximum of eight virtual processors per logical processor is supported for all other operating systems.

Maximum virtual processors for Windows

Physical Processors	Cores per Processor	Threads per Core (Hyper-Threading)	Max Virtual Processors – Win7	AlwaysOn Availability Groups
2	2	2	96	64
2	4	2	192	128
2	5	2	288	192
2	8	2	384	256
2	10	2	480	320
4	2	2	192	128
4	4	2	384	256
4	6	2	512 (Maximum)	384
4	8	2	512 (Maximum)	512 (Maximum)
4	10	2	512 (Maximum)	512 (Maximum)

It is important to realize that these are maximum values. Testing from Cisco shows that a Cisco UCS B230 M2 Blade Server (2 processor / 10 core / 2 threads) supports 145 Windows 7 desktops running the Login VSI medium workload, rather than the 480 maximum supported. Processor was identified as the primary bottleneck. For more information, please refer to the Cisco Whitepaper – [Citrix XenDesktop on FlexPod with Microsoft Private Cloud](#).

Decision: Scale Up or Out

There are a number of environmental factors that need to be considered when determining whether the Hyper-V 2008 R2 host hardware specification should be scaled up (reduced number of high-performance hosts) or scaled out (increased number of less powerful hosts), including:

- **Datacenter Capacity** – The datacenter may have limited space, power and cooling available. In this situation, consider scaling up.
- **Infrastructure and Maintenance Costs** – When determining the overall cost of the Hyper-V hosts, careful consideration should be taken when planning to include costs such as rack space, support contracts and the network infrastructure required.
- **Hosting Costs** – There may be hosting and/or maintenance costs based on the number of physical servers used. If so, consider “scaling up” to reduce the long-term costs of these overheads.
- **Redundancy** – Spreading user load across additional less-powerful servers helps reduce the number of users affected from hardware or software failure on a single host. If the business is unable to accept the loss of a single high-specification server, consider “scaling out”.

Decision: Hyper-V Edition

The design must determine which variation of Hyper-V will be used:

- **Hyper-V Server 2008 R2** – Standalone version of Hyper-V that does not include a Graphical User Interface (GUI). Administration is either performed remotely, or via the command line. Hyper-V Server does not require a Windows 2008 R2 license, however licenses are required for each guest virtual machine. For more information, please refer to the [Microsoft Hyper-V Server 2008 R2 website](#).
- **Windows Server 2008 R2** – Hyper-V is also available as an installable role for Windows Server 2008 R2 Standard, Enterprise and Datacenter editions, each of which includes different levels of “virtual server image use rights” which range from one Microsoft server operating system license to unlimited. For more information, please refer to the [Microsoft Windows Server 2008 R2 website](#).

The following table highlights the key differences between Hyper-V Server 2008 R2 and the different editions of Windows Server 2008 R2 when it comes to supporting XenDesktop:

Hyper-V Server 2008 R2 / Windows 2008 R2 Comparison

Capability	Hyper-V Server 2008 R2	Windows Server 2008 R2		
		Standard	Enterprise	Datacenter
Virtual Server Image Use Rights	0	1	4	Unlimited
Sockets	8	4	8	64
Memory	1TB	32TB	1TB	1TB
Cluster Shared Volumes	Yes	No	Yes	Yes

Hyper-V Server 2008 R2 has a decreased attack surface and does not require a Windows Server 2008 R2 license. Windows Server 2008 R2 Enterprise and Datacenter is recommended when familiarity with the command-line is limited or there is third-party software to be installed on the server that is not supported on the Hyper-V Server 2008 R2 version.

The virtual server image rights included with Windows Server 2008 R2 Datacenter make it an excellent choice for XenDesktop 7

servers hosting applications.

Note: The Standard edition of Windows Server 2008 R2 is rarely used due to the limited memory support. In addition, the Standard edition lacks Cluster Shared Volumes support, which is a recommended feature for hosting infrastructure servers and dedicated virtual desktops.

Decision: Network Connectivity

If unfamiliar with the concepts in this section, please refer to the Microsoft Whitepaper – [Understanding Networking with Hyper-V](#).

When integrating Hyper-V 2008 R2 with XenDesktop 7 it is important to consider the following key networking topics:

- **Physical Network Connections** – If sufficient infrastructure exists, performance may be improved by separating different types of network traffic across multiple physical Network Interface Cards (NICs), for example management, cluster, virtual machine, storage, provisioning and backup traffic can all be isolated from each other. A Hyper-V 2008 R2 host used to support a XenDesktop environment typically utilizes between four and eight physical network connections:
 - 2 x NICs (teamed) – Live Migration Network, Cluster Shared Volume Network, Virtual Machine Network
 - 1 x NIC – Management Network
 - 1 x NIC – Cluster Private Network

If standard network adapters, rather than HBAs, are used to access shared storage, it may be beneficial to separate out the storage traffic onto a dedicated network:

- 2 x NICs (teamed) – Storage Network

In addition, if the servers are network-constrained and Provisioning Services is utilized it may be beneficial to separate out the provisioning traffic:

- 2 x NICs (teamed) – Provisioning Network

For more information, please refer to CTX120955 – [How to Setup a Multi-Homed VM Separating Large Scale Traffic from ICA Traffic for XenDesktop](#).

Note: All hosts within a cluster must have the same number of network interfaces and the interfaces must be connected to the same networks.

- **NIC Teaming** – The implementation of NIC teaming provides increased resiliency by allowing up to two network cards to function as a single entity. In the event of one NIC failing, traffic is automatically routed to the second NIC. The Hyper-V 2008 R2 design should determine which network connections should be teamed for resiliency. Citrix recommends teaming all network connections apart from the Management and Cluster Networks because they can be failed over to alternative networks in the event of a failure.

Many servers today are supplied with NICs offering two or more ports. As such, it is important that any teams created consist of connections from two separate physical NICs so that a single card failure does not bring down the team.

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

Note: Microsoft does not support NIC teaming with Hyper-V 2008 R2; however, third party solutions are available. All support queries regarding teaming with Hyper-V 2008 R2 should be directed to the NIC OEM. In addition, some OEMs may not support TCP Chimney Offload and VMQ functionality with NIC teaming enabled. Please refer to the OEM's teaming documentation for more information. For more information, please refer to Microsoft Knowledgebase Article KB968703 – [Microsoft Support Policy for NIC Teaming with Hyper-V](#).

- **IP Addressing** – IP addresses need to be assigned to the Hyper-V 2008 R2 network 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 Hyper-V 2008 R2 hosts, ensure that DHCP reservations are created for the appropriate MAC addresses to prevent DNS resolution issues when IP addresses change.

The Hyper-V 2008 R2 network design should ensure that the Hyper-V traffic is routed via the appropriate virtual and physical networks. For example, shared storage traffic is routed via the parent partition and not directly from the virtual machine. Depending on the networking architecture used, static routes may need to be added to the parent partition's routing table to ensure that the correct interface is used.

- **VLANs** – Many network environments utilize VLAN technologies to reduce broadcast traffic, enhance security and to enable complex virtual network configurations that would otherwise not be possible. The Hyper-V 2008 R2 network design should determine which VLANs will be required. Typical VLANs used in a XenDesktop environment include:
 - Storage VLAN for storage traffic
 - Desktops and Applications VLAN for virtual desktops, hosted shared desktops, and published applications
 - Server management VLAN
 - Provisioning VLAN for streaming services

Separate physical networks are typically created for the Hyper-V Live Migration, storage, cluster and management networks.

Note: Hyper-V 2008 R2 supports the configuration and use of 802.1Q tagged VLANs. Each Hyper-V 2008 R2 host will need to have its physical NICs connected to specific VLAN trunk ports to allow for the correct routing of VLAN tagged traffic. For

more information, please refer to the TechNet Article – [Hyper-V: Configure VLANs and VLAN Tagging](#).

Decision: Virtual Network Adapter

Hyper-V 2008 R2 supports two types of network adapters:

- **Synthetic Adapter** – Virtual network adapter that uses a network driver installed via Hyper-V Integration Services to communicate with the hardware abstraction layer. Since the driver has to be loaded, the synthetic network adapter isn't available until the operating system boots.
- **Legacy Adapter** – Emulated network adapter based on the Intel 21140 PCI Fast Ethernet adapter. The legacy network adapter supports guest VMs that are not compatible with Hyper-V Integration Services and VMs that are configured to network boot.

The synthetic adapter shows better performance, requires reduced host CPU, and is installed by default when a VM is created. In XenDesktop 7 environments using Provisioning Services 7 the legacy network adapter must be used initially to boot the desktops and provisioned servers hosting applications. As such, two options must be considered during the Hyper-V 2008 R2 design:

1. Only use the legacy network adapter, which does not perform as well as the synthetic adapter.
2. Configure the Provisioning Services targets with dual network adapters: one synthetic and one legacy adapter. The legacy adapter is used to boot up the virtual machine initially. Once the virtual machine boots, Provisioning Services will detect the synthetic adapter and automatically begin to stream traffic over it. Both network adapters must reside on the same subnet.

Note: The automatic switchover from a legacy to synthetic network adapter can be disabled by editing the target's device registry setting:

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\BNISStack\Parameters]

"DisableHyperVLegacyNic"=dword:00000000

Note: If the dual NIC configuration is used, each virtual machine will require two IP addresses from the DHCP server.

The legacy network connection will be suitable for all but the most network intensive environments. Therefore, for simplicity, utilize the legacy network adapter unless an associated network bottleneck is identified.

Decision: Storage Connectivity

Storage has a major impact on the performance, scalability and availability of the XenDesktop implementation. As such, the storage design should focus on the following key areas:

- **Direct-attached Storage/Storage-attached network** – The Hyper-V 2008 R2 storage design must determine whether the virtual machines should be hosted on a Storage-attached Network (SAN), Direct-attached Storage (DAS) or a combination of the two. There are a number of factors that need to be considered, including:
 - Unlike DAS, SAN supports Live Migration, and Failover Clustering. Although these features are less critical when supporting pooled virtual desktops and server hosted applications, they are still very important for dedicated desktops and infrastructure servers. A SAN must be used if these features are required.
 - Another 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 are required to accommodate the number of IOPS generated.
 - Although using local storage may require additional disks and array controllers to be purchased per Hyper-V host,

the overall cost is likely to be significantly less than that of an enterprise storage solution.

- With a SAN there is a potential for a single point of failure. If the SAN experiences hardware failure or other type of interruption, all services running off the SAN may be impacted. If the SAN performance is degraded by a faulting service, all SAN dependent services may suffer as well.

In many cases, the optimal solution is to use SAN storage for the infrastructure servers and dedicated desktops so that they can benefit from Failover Clustering and Live Migration while using less expensive DAS storage for pooled desktops and server hosted applications.

Note: Network-attached storage (NAS) is not supported with Hyper-V 2008 R2. Please refer to the Microsoft TechNet article [KB2698995 regarding Microsoft support policy for Hyper-V environments utilizing network-attached storage](#).

Decision: Dynamic Memory

Citrix has conducted many scalability tests with Hyper-V 2008 R2 dynamic memory enabled and found it to be a very useful and recommended tool. However, there are a few areas to consider during planning:

- Dynamic memory does not oversubscribe the memory pool. When all of the physical memory has been assigned, desktops will not be able to request additional memory.
- The CPU costs to manage dynamic memory are real, but very reasonable (<1% additional CPU overhead).
- An increase in IOPS is normal and should be planned for. When a virtual machine requests additional memory, it will also start to utilize its page file in preparation for its request being denied by Hyper-V and the operating system running out of physical memory. Therefore, it should be expected that a virtual machine

will increase its I/O characteristics during this time, which can result in an increase in IOPS of up to 10%. This is a short-term increase, and one noticed more in mass scalability testing than in real-life situations. However, on systems where RAM is more plentiful than IOPS, dynamic memory settings should be conservative. For example if a Windows 7 virtual machine, with a subset of applications, has been tested and determined to require a minimum of 700MB to run then the dynamic memory default setting of 512MB should be increased to 700MB. By accepting the default setting there will be an increase in IOPS consumption in relation to paging occurring for the guest.

For more information, including recommended default values for startup RAM by operating system, please refer to the Microsoft TechNet article – [Hyper-V Dynamic Memory Configuration Guide](#).

Decision: SCVMM Architecture

When designing a VMM 2012 environment there are two design choices that should be considered.

- **Design 1** – A single VMM virtual machine exists within a host cluster, or a series of VMM servers each supporting no more than 2,000 virtual desktops each. These VMM virtual machines have no VMM application level HA, however they do have VM based high availability enabled with the Hyper-V cluster. The VM operating system and the VMM application components can easily migrate between hosts within the cluster in the event of a hardware failure. This design is recommended for deployments that can afford minor disruption to service as some downtime is required for failover to occur. This design offers a simpler setup and is the most often used.
- **Design 2** – Two VMM 2012 servers are used in a guest failover cluster configuration. This design offers VMM application layer HA, however it can also increase complexity during the initial setup of the environment. This design is more restrictive than host clustering, and also requires iSCSI shared storage, since

it is the only storage type supported by guest clusters. This design is recommended for critical deployments that can't afford any downtime due to a disruption of service.

Both of these designs are verified in the [Microsoft Fast Track Reference Architecture Guide](#) white paper, which also provides additional VMM 2012 architectural guidance.

Although Hyper-V 2008 R2 virtual machines can function without VMM, XenDesktop 7 relies on its availability to manage virtual machines. Therefore, it is imperative that the following VMM components are protected from failure:

- **VMM server** – If the VMM server is unavailable, XenDesktop 7 will not be able to manage the power state of the virtual machines that it manages or create additional virtual machines. Therefore, Microsoft failover clustering should be included in the design to ensure that the VMM server is highly available. For more information, please refer to the Microsoft blog – [Creating a Highly Available VMM server](#).
- **VMM library** – The VMM library is typically required for the creation of new desktops, servers hosting applications, and infrastructure servers. Therefore, the loss of the VMM Library is unlikely to impact the day-to-day activities of a standard XenDesktop deployment. However, if required, VMM supports the use of highly available library shares on a failover cluster.
- **SQL Server** – The VMM database contains all VMM configuration information. If the VMM database is unavailable, virtual machines will still be available but the VMM console will be unavailable. Therefore, Hyper-V failover clustering should be used to ensure that the SQL Server is highly available.

Decision: SCVMM Sizing

The following recommendations should be taken into consideration when determining the number of VMM servers required:

- A VMM 2012 server can manage up to 8,000 virtual machines.

When a VMM server is used to manage both virtual desktops and servers, Citrix has found that best performance is achieved when each VMM server is limited to managing 2,000 virtual desktops. There is no additional limitation placed on the number servers, or other control VMs also managed by each VMM server. For example, a VMM server could support 2,000 virtual desktops and 6,000 servers.

- While it is possible to run other applications on the VMM server, it is not recommended, especially other System Center 2012 applications because they tend to have heavy resource demands and could significantly impact VMM performance.

The following table provides recommended specifications for the key VMM infrastructure components

SCVMM recommendations

Requirement	Recommended	Recommended (>50 Hosts)
VMM Management Server		
Processor	Dual processor, dual core, 2.8GHz (x64) or greater	Dual processor, dual core, 3.6GHz (x64) or greater
Memory	4GB	8GB
Disk space (no local DB)	40GB	50GB
Disk space (local DB)	150GB	Use a dedicated SQL Server
Dedicated SQL Server		
Processor	Dual processor, dual core, 2GHz (x64) or greater	Dual processor, dual core, 2.8GHz (x64) or greater
Memory	4GB	8GB
Disk space	150GB	200GB
VMM Library Server		
Processor	Dual processor, dual core, 3.2GHz (x64) or greater	Dual processor, dual core, 3.2GHz (x64) or greater
Memory	2GB	2GB
Disk Space (no local DB)	Size will depend on what will be stored	Size will depend on what will be stored
VMM Admin Console		
Processor	Single processor, 1GHz or greater	Single processor, 2GHz or greater
Memory	2GB	4GB
Disk space	2GB	4GB

[Click here to provide feedback](#)

For more information on VMM 2012 system requirements, please refer to the following Microsoft TechNet articles:

- [System Requirements: VMM Management Server](#)
- [System Requirements: VMM Console](#)
- [System Requirements: VMM Database](#)
- [System Requirements: VMM Library](#)

Decision: Failover Clustering

A Hyper-V 2008 R2 failover cluster is a group of Hyper-V hosts (cluster nodes) that work together to increase the availability of hosted virtual machines. Virtual machines on a failed host are automatically restarted on another host within the same cluster, provided that sufficient resources are available.

The following should be considered when designing the Hyper-V 2008 R2 failover cluster:

- **Virtual machines to cluster** – The following virtual machines should be hosted on a failover cluster because they are critical to the successful operation of XenDesktop 7 on Hyper-V:
 - VMM servers
 - StoreFront servers
 - License servers
 - SQL servers
 - Delivery controllers
 - Provisioning Services
 - Server hosted applications
 - Dedicated desktops
 - Pooled desktops (with personal vDisk),
 - Hosted shared desktop servers

Pooled desktops and hosted shared desktops do not typically

need to be hosted on a cluster because they shouldn't have user-specific data.

Note: Cluster shared volume services depend on Active Directory for authentication and will not start unless they can successfully contact a domain controller. Therefore, always ensure that there is at least one domain controller that is not hosted on a cluster shared volume. For more information, please refer to the Microsoft knowledge base article KB888794 – [Things to Consider When You Host Active Directory Domain Controllers In Virtual Hosting Environments](#).

- **Number of failover clusters** – The number of Hyper-V 2008 R2 clusters required to support an implementation of XenDesktop varies depending on the following key factors:
 - **Infrastructure nodes** – When designing Hyper-V 2008 R2 clusters, consider separating the host server components such as AD, SQL and XenDesktop controllers from the dedicated virtual desktops by placing them on different Hyper-V clusters. This will ensure that dedicated virtual desktop resources are isolated from infrastructure resources for optimal performance and availability.
 - **Performance** – Many businesses will have dedicated desktops that require guaranteed levels of performance. As such, it is sometimes necessary to create dedicated clusters to meet the service-level agreements associated with these desktops.
 - **Application set** – It may be beneficial to dedicate clusters for large dedicated desktop groups as they share a common, predictable resource footprint and application behavior. Alternatively, grouping dedicated desktop groups together based on differing resource footprints could help to improve desktop density per host. For example, splitting processor intensive dedicated desktops across several clusters will help to distribute the impact from processor

saturation. For servers hosting applications, it is advisable to separate the servers onto a separate Hyper-V cluster to balance resource utilization.

- **Physical network** – Some environments may have complex network configurations that require multiple clusters to be deployed. For example, some delivery groups may need to be isolated onto specific subnets for reasons of security, whilst others may have requirements for network connectivity that can only be satisfied within specific datacenters or network locations.
- **Virtual network** – Depending on the environment, it may be overly complex to trunk every VLAN to every Hyper-V cluster. As such, it may be necessary to define clusters based on the VLANs to which they are connected.
- **Processor architecture** – Hyper-V requires that nodes within the same cluster have processors from the same family. Therefore, the Hyper-V design will need to consider the processor architecture available when determining the number of clusters required.
- **Nodes per failover cluster** – A Hyper-V 2008 R2 failover cluster supports up to 16 nodes, however consider reserving at least one node in the cluster for failover and maintenance.

A single Hyper-V 2008 R2 cluster can support up to a maximum of 1,000 virtual machines. It is unlikely that a cluster hosting virtual desktops would consist of 16 nodes as this would limit each node to 63 virtual machines, 66 if one host is dedicated to HA. It is far more likely that a Hyper-V 2008 R2 cluster supporting XenDesktop virtual desktops will consist of between 7 and 13 nodes. The following formula can be used to calculate the number of VMs per node:

$$\text{Number of Virtual Machines per Node} = 1000 / (\text{Number of Nodes} - 1)$$

- 7 nodes per cluster – 167 VMs per node with one

dedicated to HA

- 14 nodes per cluster – 76 VMs per node with one dedicated to HA

For more information, please see the Microsoft TechNet article – [Requirements and Limits for Virtual Machines and Hyper-V in Windows Server 2008 R2](#).

Note: When multiple virtual machines exist for each server role, ensure that they are not all hosted on the same physical host. This will help to ensure that the failure of a single virtualization host does not result in a service outage. In addition, the physical hosts supporting the core infrastructure components should ideally be located in different chassis/racks.

- **Cluster shared volumes** – A cluster shared volume (CSV) allows virtual machines that are distributed across multiple cluster nodes to access their virtual hard disk (VHD) files at the same time. The clustered virtual machines can all fail over independently of one another. CSVs are required for failover clustering and live migration functionality. Therefore infrastructure servers and dedicated desktops are typically hosted on CSVs.

The following recommendations should be considered during the cluster shared volume design:

- Microsoft recommends that the CSV communications take place over a different network to the virtual machine and management traffic.
- The network between cluster nodes needs to be low latency to avoid any lag in disk operations but doesn't need to be high bandwidth due to the minimal size of metadata traffic generated under normal circumstances.

Note: Since clustered shared volume communication occurs over the server message block (SMB) protocol, the Client for Microsoft Networks and File and Printer Sharing for Microsoft

Networks services must be enabled on the network adapters used for the cluster network. Disabling NetBIOS over TCP/IP is recommended.

For more information, please see the Microsoft TechNet article – [Requirements for Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2](#).

Hyper-V 2012 R2

This chapter provides design guidance for Citrix XenDesktop 7.x deployments that leverage Microsoft Hyper-V 2012 R2. It will give the reader a high-level understanding of the key implementation decisions when building a Hyper-V 2012 R2 environment

Hardware

Decision: Host Hardware

The hardware selected for the Hyper-V hosts will have a direct impact on the performance, scalability, and resilience of the XenDesktop solution. Hardware compatibility is important for maintaining a stable Hyper-V environment. Microsoft maintains a catalog of various hardware components that have been tested and certified to work with Windows Server 2012 R2. Any product that meets the requirements of the “Certified for Windows Server 2012 R2” logo is fully supported in Hyper-V R2 environments. Please refer to the [Windows Server Catalog](#) for a full listing of hardware products that have been certified for Windows Server 2012 R2.

Decision: Hyper-V Edition

Hyper-V 2012 R2 is available as a server role in Windows Server 2012 R2 editions: Standard, Datacenter, and as a stand-alone product Hyper-V Server 2012 R2. XenDesktop 7.x is only supported

on the Standard and Datacenter editions.

XenDesktop 7.x manages virtual machines through System Center Virtual Manager 2012 and 2012 R2, which is supported on Windows Server 2012/2012 R2 Standard and Datacenter editions, and Windows Server 2008 R2 Standard, Enterprise, and Datacenter editions.

As shown in the following table, the primary difference between Hyper-V editions is the number of virtual server image use rights supported. Hyper-V Server 2012 R2 is ideal for environments with existing server licenses and no additional licenses are required.

For small or test environments (approximately ten virtual servers or less), Windows Server 2012 R2 Standard will be the more cost effective solution.

For Hyper-V servers hosting approximately ten or more virtual servers, the Datacenter edition is the more cost effective solution since it provides an unlimited number of virtual server instances with each license.

Note: The cost of Microsoft products is subject to change, so please check with a Microsoft reseller for current pricing.

Hyper-V Editions comparison

Capability	Windows Server 2012 R2		Hyper-V Server 2012 R2
	Standard	Datacenter	
Virtual Server Image Use Rights	2	Unlimited	0
Sockets	64	64	64
Logical Processors	320	320	320
Memory	4TB	4TB	4TB
Virtual Desktop Infrastructure (VDI)	Yes	Yes	Yes
Host Clustering	Yes	Yes	Yes
Live Migration	Yes	Yes	Yes
Local Graphical User Interface	Yes	Yes	No
Supports additional server roles	Yes	Yes	No

For more information on Windows Server 2012 R2 licensing for Hyper-V refer to the [Microsoft Brief – Licensing Microsoft Server Products for Use in Virtual Environments](#).

[Click here to provide feedback](#)

Decision: Processor

The server hosting the Hyper-V role must meet the following processor requirements:

- A x64-based processor that's optimized for virtualization using either the Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor with a clock speed of at least 1.4 GHz. (2GHz or faster is recommended).
- Hardware-enforced Data Execution Prevention (DEP) must be enabled. Specifically, the Intel XD bit (execute disable bit) or the AMD NX bit (no execute bit).

Both settings are usually available in the computer BIOS. Check with the manufacturer of the computer for the specific location.

For improved performance Second Level Address Translation (SLAT) capable processors can be used, but are not a requirement. This technology provides hardware assistance for memory virtualization, memory allocation, and address translation between the physical and virtual memory. The reduction in processor and memory overhead improves scalability allowing a single Hyper-V host to run more virtual machines.

If live migration or failover clustering will be used, then the processors in the hosts must be from the same family type. It is not possible to migrate VMs from an AMD CPU host to a host with an Intel CPU.

Decision: Memory

A Hyper-V 2012 R2 host can support up to 4TB of memory, however most XenDesktop 7.1 deployments rarely allocate more than 256GB of RAM per desktop host. Determining how much memory is required will depend on different factors including:

- The number and type of virtual machines to be supported per host
- The operating system used by the virtual machines

The following calculation can be used to estimate the amount of memory required per Hyper-V 2012 R2 host:

Host Total Memory =

*Hypervisor Overhead + Memory for OS + (# of VMs * RAM per VM)*

Where:

- Hypervisor overhead – The hypervisor typically requires at least 300MB of memory. Create VMs and factor in another 32MB for the first GB of RAM allocated to each VM and 8MB for every additional GB of RAM allocated to each VM. For example, a VM with 2GB of RAM would use 40MB (32MB plus 8MB) of memory for virtualization-related overhead.
- Memory for OS – A minimum of 512 MB should be allocated for the host operating system.
- RAM per VM - Although Hyper-V supports up to 1TB of memory per VM most hosted VDI desktops typically require less than 4GB of memory. For hosted shared desktops, multiple users are simultaneously hosted on a single server. The number of users that can be hosted on a server depends upon user load, number of processors, and memory. A “typical” hosted shared desktop deployment with standard business applications installed like Microsoft Office will not require more than 16GB of memory.

Memory cannot be allocated directly to the host parent partition, therefore Microsoft recommends at least 2GB of RAM for the host parent partition (hypervisor overhead and host operating system).

For example, if sizing a host to support 50 VMs with 2GB of RAM each, a safe estimate for total memory would be:

$$102.5GB = (40MB * 50) + 512MB + (50 * 2GB)$$

Note: Microsoft does not recommend running applications other than management agents, backup agents and firewalls on the Hyper-V host.

Note: To improve the performance of virtual machines configured with a large amount of memory, Hyper-V 2012 R2 supports the Non-Uniform Memory Architecture (NUMA). NUMA is a computer architecture used in multiprocessor systems, where the memory access time depends on the memory location relative to the processor. The NUMA architecture divides memory and processors into nodes. For any single processor, memory that is in the same NUMA node is considered local, and memory that is contained in another NUMA node is considered remote. Since accessing local memory is faster than remote, Hyper-V 2012 R2 attempts to allocate all memory for the virtual machine from a single physical NUMA node. If the memory requirements for the virtual machine cannot be satisfied from a single node, Hyper-V allocates memory from another physical NUMA node. For more information on the NUMA architecture and how it relates to Hyper-V please refer to the TechNet article – [Hyper-V Virtual NUMA Overview](#).

Assigning dynamic memory to virtual machines can reduce the total amount of RAM required. Virtual machines would only consume the amount of RAM necessary for the current workload. When using dynamic memory however, it is very important to understand application requirements. Some applications are designed to use as much memory as is available and can exhaust the physical memory in a host.

Note: For information on the recommended minimum dynamic memory setting by operating system please refer to the Microsoft blog [Windows Server 2012 R2 Hyper-V Best Practices](#).

Note: It is possible for virtual machines to suffer a small performance hit when configured to use dynamic memory. If the automatic stop action of a virtual machine is set to “Save the virtual machine state” a .BIN file is created for that virtual machine. The .BIN file is used to store the memory contents of the virtual machine when it goes into a saved state. When used with virtual machines configured with dynamic memory, the size of the .BIN file will adjust in real-time as memory changes with the virtual machine. For more

information on .BIN files please see the [Storage Requirements](#) section of this document.

Decision: Storage Requirements

The amount of storage required will vary depending on whether Windows Server 2012 R2 or Hyper-V Server 2012 R2 is used:

- Windows Server 2012 R2 requires a minimum of 32GB for storage, however if the Hyper-V host will have 16GB or more memory, then more disk space is required in order to adequately store files for paging, and dump files. Please refer to [System Requirement and Installation Information for Windows Server 2012 R2](#) for more information.
- Hyper-V Server 2012 R2 requires a minimum of 27GB for storage for the Server Core mode installation with only the Hyper-V role enabled. For more information, please refer to the following Microsoft TechNet article – [Windows Server Installation Options](#).

The provisioning method used to deliver the virtual machines also plays a role in determining how much storage will be required:

- Virtual desktops will typically require between 15 and 40GB of storage depending on the operating system used and installed applications. Provisioning Services helps to reduce virtual desktop storage requirements to between 1 and 3GB.
- Hosted shared desktops typically require between 40 and 60GB of storage depending on the application set. Provisioning Services helps to reduce hosted shared desktop requirements to between 10 and 15GB.

To help determine the total storage required for virtual machines when using Machine Creation Services, the following formula can be used:

Total VM Storage =

$$\begin{aligned} & \# \text{ Master Images} * \text{ Size} + \# \text{ Virtual Machines} \\ & * (\text{Differencing Disk Size} + \text{Identity Disk Size}) \end{aligned}$$

For example, 500 Windows 8 desktops will be deployed from two 40GB master images. A 6GB differencing disk will be allocated to each desktop. The amount of storage required for the virtual desktops will be:

$$\text{Total VM Storage} = 2 * 40\text{GB} + 500 * (6\text{GB} + 16\text{MB}) = 3,088\text{GB}$$

A similar formula can be used to determine total storage when using Provisioning Services:

$$\begin{aligned} \text{Total VM Storage} &= \# \text{vDisk Images} * \text{Size} \\ &+ \# \text{Virtual Machines} * \text{Write Cache Disk Size} \end{aligned}$$

Using the same example as before:

$$\text{Total VM Storage} = 2 * 40\text{GB} + 500 * 6\text{GB} = 3,080\text{GB}$$

If Hyper-V will be configured to store the saved state of virtual machines automatically when a stop action occurs, then the .BIN files generated should be factored into the storage equation. The size of the .BIN file will match the amount of RAM assigned to the virtual machine. The formula for calculating VM storage when using Machine Creation Services will change to:

Total VM Storage =

$$\begin{aligned} & (\# \text{ Master Images} * \text{ Size} + \# \text{ Virtual Machines}) \\ & * (\text{Differencing Disk Size} + \text{Identity Disk Size}) \\ & + (\# \text{ Virtual machines} * \text{ Size of RAM}) \end{aligned}$$

Using the example as before, each virtual desktop will be assigned 2GB of RAM:

$$\text{Total VM Storage} = 3,088 + (500 * 2) = 4,088\text{GB}$$

For Provisioning Services the new formula becomes:

$$\begin{aligned} \text{Total VM Storage} = \\ (\#vDisk\ Images * Size + \#Virtual\ Machines * Write\ Cache\ Disk\ Size) \\ + (\#Virtual\ Machines * Size\ of\ RAM) \end{aligned}$$

Therefore, following the same example, the new storage required for Provisioning Services when taking into account the .BIN files will be:

$$\text{Total VM Storage} = 3,080 + (500 * 2) = 4,080\text{GB}$$

Note: .BIN files can consume a significant amount of disk space. Saving the virtual machine state for random (pooled) and hosted shared desktops is not necessary, therefore consider changing the automatic stop action to either “Turn off virtual machine” or “Shut down the guest operating system”. For more information on how to change the virtual machine automatic stop action please see the Microsoft TechNet blog [Using SCVMM Cmdlets to change VM startup and shutdown actions](#).

Decision: Network Requirements

At a minimum, Windows Server 2012 R2 requires a Gigabit (10/100/1000baseT) Ethernet adapter. However, in order to take advantage of the Hyper-V networking features, two or more physical network adapters are required. A minimum of two physical network adapters are also required to setup NIC Teaming, load balancing, and cluster failover. The two physical network adapters must operate at the same speed. For more information on Hyper-V network design decisions, please refer to the [Networking](#) section of this document.

Decision: Scale Up/Out

There are a number of environmental factors that need to be considered when determining whether the Hyper-V host hardware specification should be “scaled up” (reduced number of high-

performance hosts) or “scaled out” (increased number of less powerful hosts), including:

- **Data center capacity** - The data center may have limited space, power and/or cooling available. In this situation, consider scaling up.
- **Infrastructure and maintenance costs** - When determining the overall cost of the Hyper-V hosts, careful consideration should be taken when planning to include costs such as rack space, support contracts and the network infrastructure required.
- **Hosting costs** - There may be hosting and / or maintenance costs based on the number of physical servers used. If so, consider “scaling up” to reduce the long-term costs of these overheads.
- **Redundancy** - Spreading user load across additional less-powerful servers helps reduce the number of users affected from hardware or software failure on a single host. If the business is unable to accept the loss of a single high-specification server, consider “scaling out”.

Host Scalability

During the design, it is necessary to estimate the number of physical host servers that will be required to support the XenDesktop implementation. Ideally, scalability testing should be performed prior to the hardware being ordered, however this is not always possible.

In previous versions of Hyper-V there were limits on the number of virtual processors to logical processors in a host machine. In Windows Server 2012 R2 those ratios no longer apply. A Hyper-V 2012 R2 host can have up to 320 logical processors, 2,048 virtual processors, and 1,024 active VMs. The following scalability table applies to both Windows Server 2012 R2 Standard and Datacenter editions running Hyper-V.

Hyper-V 2012 R2 Scalability Limits

System	Resource	Max Number
Host	Logical Processors	320
	Physical Memory	4TB
	Virtual Processors	2,048
Virtual Machine	Virtual processors per VM	64
	Memory per VM	1TB
	Active VMs per host	1,024
	IDE devices	4
	SCSI disks	256
	Max disk size	64TB (vhdx) 2TB (vhd)
	Virtual NICs	12
	Virtual Fibre Channel adapters	4
Cluster	Nodes	64
	Virtual Machines	4,000

Networking

A proper network design is critical for ensuring optimal performance. This section covers important network design considerations for Hyper-V 2012 R2.

Decision: Networks

The following networks are typically required when hosting XenDesktop on Hyper-V:

- SCVMM/Hyper-V management
- Failover Clustering
- Live Migration
- VM traffic to external/public network
- Storage

A typical XenDesktop deployment on Hyper-V requires four different networks:

1. Storage

2. Clustering
3. Management
4. All other traffic

For networks that are bandwidth constrained it may be necessary to isolate Provisioning Services and / or live migration traffic to improve performance. This scenario is becoming increasingly rare as more and more datacenters adopt 10Gbps networks as the standard. In addition, some organizations may have security requirements which demand certain traffic be isolated.

Network isolation can be achieved on the Hyper-V host in the following ways:

- Hyper-V host is connected to a physical switch. Isolation occurs at the switch by configuring specific VLANs to different switch ports. The physical network adapters are connected to different switch ports. This solution requires a physical network adapter for each network.
- Hyper-V host is connected to a trunked switch port. Separate physical network adapters are tagged with a specific VLAN ID. This solution also requires a physical network adapter for each network.
- Hyper-V host is connected to a trunked switch port. Virtual network adapters are created on the Hyper-V host parent partition. Each virtual network adapter is tagged to a specific VLAN ID.

VLANs tag traffic with specific VLAN IDs. The Hyper-V host's physical network adapter is connected to a switch port that is usually trunked so that multiple VLANs can be routed through the physical network adapter. This configuration is ideal for most blade servers which have a limited number of network adapter slots available.

The following table summarizes the network isolation options available to Hyper-V. The table assumes that five networks will be

used in the solution and each network will be dedicated. For the sake of simplicity, the table will be based on a non-teamed network adapter.

Network isolation options

Network properties	Non-trunked switch port	Trunked switch port / physical adapters	Trunked switch port/ virtual adapters
Number of physical NICs required	Five	Five	One
Network management	Windows Network and Sharing Center	Windows Network and Sharing Center and Hyper-V Manager	Windows Network and Sharing Center and Hyper-V Manager
Benefits	Failure of one NIC will not effect the other networks	Failure of one NIC will not effect the other networks	Inexpensive to implement
Drawbacks	Physical NICs must connect to specific switch ports; change for human error increases. Expensive to implement due to the number of physical NICs and switch ports required.	Expensive to implement due to the number of physical NICs required for solution.	Failure of physical NIC will affect all virtual networks.
Best use case	Large datacenter deployments	Servers with expansion slots for additional network adapters.	Deployments on server hardware with limited number of physical adapters.

Decision: Physical NIC

When the Hyper-V role is enabled, the physical network adapters in the host server are detected and administrators must decide which adapters to use as virtual Hyper-V switches. The Hyper-V switch, also referred to as an extensible switch, is a Layer 2 virtual switch bound to a physical NIC for external networking traffic.

There are three types of virtual switches available in Hyper-V:

- External – Virtual machines can communicate with anything on the physical network. A physical NIC is required in the host server.
- Internal – Virtual machines can only communicate with the host

server and other virtual machines on the same server. It does not require a physical NIC in the host server.

- Private – Virtual machines can only communicate with each other virtual machines on the same server. It does not require a physical NIC in the host server.

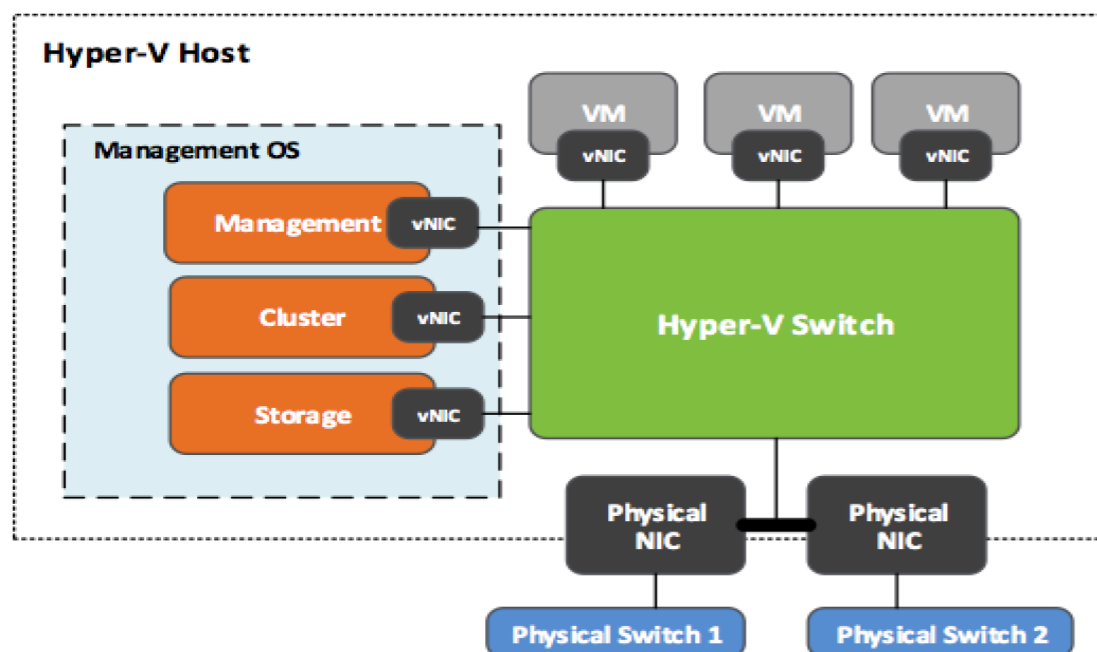
Note: XenDesktop has only been tested with the external virtual switch type.

The Hyper-V switch is also used to connect virtual NICs created in the parent partition to the external network. Virtual NICs will be covered in the [Virtual NICs – Hosts](#) section of this document.

Note: It is considered a best practice to dedicate at least one physical adapter for management access to the Hyper-V host. This will prevent high workload virtual machines from hindering administrative access to the host.

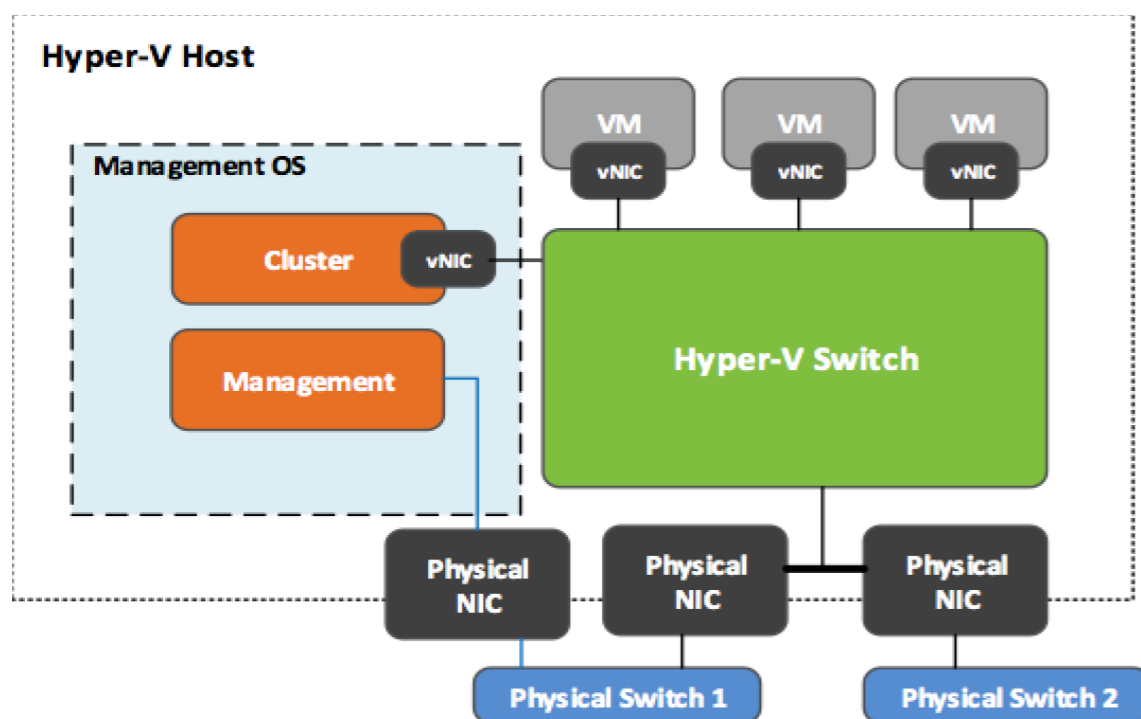
The following diagrams are examples of two Hyper-V network configurations. In the first diagram the Hyper-V switch is built on a pair of teamed physical NICs. Virtual NICs are created in the management OS for the management, storage, and cluster networks. A physical NIC dedicated for management may not always be available due to cost or physical limitations of the hardware, in which case using a virtual management adapter will suffice.

Hyper-V 2012 R2 Networking Example



In the second example, a physical NIC is dedicated for management use. All other traffic is routed through a pair of teamed physical adapters. Storage (not shown) is using a Host Bus Adapter (HBA) connection.

Hyper-V 2012 R2 Networking Example



Decision: NIC Teaming

The implementation of NIC teaming provides increased resiliency by allowing two or more network adapters to function as a single entity. In the event of one NIC failing, traffic is automatically routed to the next available NIC. In addition, when using the native teaming features of Windows Server 2012/2012 R2 outbound traffic can be distributed among the available links in one of two different ways:

- Hashing – This method keeps all packets from the same TCP stream on the same network adapter.
- Hyper-V Switch Port – This method distributes traffic based on the MAC address of the virtual machine and is advantageous for VDI environments because the switch can determine that the specific source MAC address is on one connected network adapter, and can balance the traffic from the switch to the computer across multiple links, based on the destination MAC address of the virtual machine. This method however does limit the virtual machine to the bandwidth available on a single network adapter.

Ideally, all network traffic should use teamed network adapters apart from the management and cluster networks because this traffic can be failed over to alternative networks in the event of failure.

Many servers today are supplied with network adapters offering two or more ports. As such, it is important that any teams created consist of connections from separate physical NICs so that a single card failure does not bring down the team.

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

Note: Virtual machines can be configured with virtual teamed NICs that are connected to more than one Hyper-V switch. This will

ensure that VMs still have connectivity if one of the physical network adapters becomes unavailable, but adds complexity to the solution. For most XenDesktop deployments creating the team at the physical layer will be sufficient.

Note: Since failover between network adapters in a virtual machine might result in traffic being sent with the MAC address of the other network adapter, each Hyper-V switch port that is associated with a virtual machine that is using NIC teaming must be set to allow MAC spoofing or must have the “AllowTeaming=On” parameter set using the `Set-VmNetworkAdapter PowerShell cmdlet`. For more information please see the Microsoft TechNet article [NIC Teaming Overview](#).

Note: Hyper-V 2012/2012 R2 supports NIC teaming with network adapters that support Single-Root I/O Virtualization (SR-IOV), however SR-IOV traffic does not go through the Hyper-V virtual switch and therefore cannot take advantage of the protection provided by a NIC team under a virtual switch. Citrix recommends enabling SR-IOV in Hyper-V only on XenDesktop networks that will be dedicated to Live Migration.

Decision: Virtual NICs – Hosts

In the previous release of Hyper-V, only one virtual NIC (vNIC) was supported in the parent partition. In Hyper-V 2012/2012 R2, multiple vNICs can be created and dedicated to specific networks including live migration, storage, and management. Each vNIC can be assigned to a different virtual Local Area Network (VLAN) as a way of isolating specific traffic coming across the physical adapter.

Before creating vNICs it is advisable to first create a NIC team for redundancy, although it is not required. The Hyper-V extensible switch is created on top of the NIC team allowing virtual NICs to be created on the Hyper-V host. The virtual NICs will appear in the GUI alongside physical NICs, and the IP configuration and NIC properties can be set just like a physical NIC. Hyper-V virtual NICs are identified as “Hyper-V Virtual Ethernet Adapter #” in Windows Server 2012 R2.

Network connections showing Hyper-V vNICs

The screenshot shows the Windows Network Connections window. The breadcrumb path is Control Panel > Network and Internet > Network Connections. The window displays a table of network connections with the following data:

Name	Status	Device Name	Connectivity	Network Category
VM Network	Enabled	HP NC362i Integrated DP Gigabit Server Adapter		
Management	Network	HP NC362i Integrated DP Gigabit Server Adapter #2	Internet access	Private network
vEthernet (Virtual Switch)	Network	Hyper-V Virtual Ethernet Adapter #2	Internet access	Private network
vEthernet (Storage)	Unidentified network	Hyper-V Virtual Ethernet Adapter #3	No Internet access	Public network

The following table can be used as a guide to help with the planning of networks based on the type of network and the number of physical adapters available. The configuration settings are recommendations only and may not be applicable to every XenDesktop network design.

Recommended Host Configuration

Number of physical adapters	10 Gbps+ recommended configuration	1 Gbps+ recommended configuration
2 adapters (Note: If using iSCSI storage, this configuration is only supported when using the native Windows 2012/2012 R2 teaming software.)	Configure 1 NIC team with 4 vNICs. NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Management vNIC2: Clustering vNIC3: Live Migration vNIC4: Storage (iSCSI) 	Configure 1 NIC team with 4 vNICs. NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Management (Windows QoS applied) vNIC2: Clustering vNIC3: Live Migration (Windows QoS applied) vNIC 4: Storage (iSCSI)
3 adapters	Configure 1 physical NIC, 1 NIC team with 3 vNICs. pNIC1: Management NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Clustering vNIC2: Live Migration vNIC3: Storage (iSCSI) 	Configure 1 physical NIC, 1 NIC team with 3 vNICs. pNIC1: Management NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Clustering vNIC2: Live Migration (Windows QoS applied) vNIC3: Storage (iSCSI)
4 adapters	Configure 2 physical NICs, 1 NIC team with 2 vNICs pNIC1: Management pNIC2: Live Migration (with SR-IOV) NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Clustering vNIC2: Storage (iSCSI) 	Configure 2 physical NICs, 1 NIC team with 2 vNICs pNIC1: Management pNIC2: Live Migration (with SR-IOV) NIC Team: VM traffic <ul style="list-style-type: none"> vNIC1: Clustering vNIC2: Storage (iSCSI)
5 adapters	Not necessary	Configure 1 physical NIC, 2 NIC teams with 1 vNIC pNIC1: Management NIC Team1: VM traffic <ul style="list-style-type: none"> vNIC1: Clustering NIC Team2: Live Migration <ul style="list-style-type: none"> vNIC2: Storage (iSCSI)
6 adapters	Not necessary	Configure 2 physical NICs, 2 NIC teams, 1 vNIC pNIC1: Management pNIC2: Live Migration NIC Team1: VM Traffic NIC Team2: Clustering <ul style="list-style-type: none"> vNIC1: Storage (iSCSI)

Decision: Virtual NICs – Guests

Hyper-V 2012 R2 Guest VMs support three types of network adapter: a standard network adapter (also referred to as a synthetic adapter), a legacy adapter, and a Fibre Channel adapter.

Standard Network Adapter

In Hyper-V when a virtual machine is created it is configured to use the standard network adapter by default. The standard network adapter performs better than the legacy network adapter because it uses network drivers to pass data through the Hyper-V VMBus. The standard driver must be loaded in order for the guest virtual machine to use so it cannot be used prior to Windows starting. This network adapter is recommended for XenDesktop 7.x deployments using Machine Creation Services (MCS).

Legacy Network Adapter

The legacy network adapter is an emulated multiport DEC 21140 10/100 Mbps adapter. The legacy network adapter supports Pre-Execution Environment (PXE) booting, and is required for XenDesktop 7 deployments using Provisioning Services (PVS) since it uses PXE to boot the desktops.

Note: The legacy network adapter is also required when using Boot Device Manager (BDM) to boot the desktops. The BDM bootstrap requires a network adapter with PXE extensions enabled.

The legacy network adapter is not as efficient as the standard network adapter, it also adds more overhead in the parent partition since every instruction has to be emulated. A new feature added to Provisioning Services 7 is the ability for virtual machines to switch over from the legacy adapter to the standard adapter after the virtual machine boots. In order for the switchover to occur the standard adapter must be on the same subnet as the legacy adapter.

When designing a XenDesktop 7.x deployment using Provisioning Services, administrators are faced with two design decisions:

- Only use the legacy network adapter for target devices.
- Configure the Provisioning Services targets with one standard and one legacy adapter. This configuration will potentially double the number of DHCP addresses required for the target devices.

Note: When using dual network adapters ensure the templates have the correct binding order with the legacy adapter first, so that provisioned virtual machines are able to boot successfully.

For the sake of simplicity Citrix recommends using only the legacy network adapter for XenDesktop 7.x deployments when using Provisioning Services, unless there will be applications involved that require the better performing standard network adapter.

Note: Hyper-V 2012 R2 Generation 2 virtual machines support PXE boot using the standard network adapter, however XenDesktop 7 does not support Hyper-V Generation 2 virtual machines.

Fibre Channel Adapter

The virtual Fibre Channel adapter allows guest operating systems to be clustered over Fibre Channel and to directly access virtual SANs. This feature is not available to virtual machines running Windows desktop operating system. Therefore, it is unlikely that it will be required for XenDesktop 7 deployments.

Decision: IP Addressing

IP addresses need to be assigned to the Hyper-V network interfaces and individual virtual machines. As such, the design must consider the IP addressing requirements for those components. If DHCP is used to provide the IP configuration for the Hyper-V hosts, ensure that reservations are created for the appropriate MAC addresses to prevent configuration issues when IP addresses change.

A XenDesktop design employing Provisioning Services will require virtual machines with the legacy network adapter to boot. When

designing a solution that will assign the legacy and standard network adapters, make sure that the DHCP scope is sufficient to allocate two unique IP addresses to each virtual machine created.

The Hyper-V network design should ensure that the Hyper-V traffic is routed via the appropriate virtual and physical networks. For example, shared storage traffic is routed via the parent partition and not directly from the virtual machine. Depending on the network architecture used, static routes may need to be added to the parent partition's routing table to ensure that the correct interface is used.

Network Performance Tuning

To improve the network performance of XenDesktop 7.x on Hyper-V hosts the following practices are recommended:

- Consider enabling Large Send Offload, TCP Checksum Offload, Receive Side Scaling, and Dynamic Virtual Machine Queue to reduce CPU usage of network I/Os from virtual machines. Most network adapters will have these features enabled by default. Consult the manufacturer for details on whether these features are supported and enabled in Windows Server 2012/2012R2
- To improve VLAN performance, use physical network adapters that support NDIS_ENCAPSULATION_IEEE_802_3_P_AND_Q_IN_OOB encapsulation for Large Send Offload and checksum offload support. Without this support, Hyper-V cannot use hardware offload for packets that require VLAN tagging and network performance can be decreased. For more information, please refer to the Microsoft article – [Performance Tuning Guidelines for Windows Server 2012 R2](#).
- Use a private dedicated network for live migration traffic. This will help minimize the time required to complete live migrations and ensure consistent migration times. If the network adapter supports it, then enable Single Root I/O Virtualization (SR-IOV)

on this dedicated network only. SR-IOV can help reduce network latency, CPU utilization, and increase network throughput. In addition, increasing the number of receive and send buffers on each network adapter that is involved in the migration can improve migration performance.

Note: When SR-IOV is enabled, the traffic bypasses the Hyper-V virtual switch and cannot take advantage of redundancy that a NIC team provides. Teaming SR-IOV enabled network adapters is not supported, therefore SR-IOV is only recommended on networks dedicated to live migration.

- Hyper-V supports creating multiple virtual network switches. Each switch is associated with a physical network adapter. If the Hyper-V host has multiple network adapters, virtual machines with network-intensive loads can benefit by having their workloads divided between the virtual switches for better use of the physical network adapters.

Decision: Security

For security, firewalls should be used to control traffic flow between the following components:

- Virtual desktops and the infrastructure servers
- End users and the virtual desktops
- End users and the infrastructure servers

The Hyper-V network design should ensure that port 8100 (WCF) is open between the XenDesktop Controllers and the VMM server to facilitate machine state queries and power management operations. For a full list of ports required by XenDesktop, please refer to the Citrix Knowledgebase article CTX101810 – [Communication Ports Used By Citrix Technologies](#).

Hyper-V Storage

Storage has a major impact on the performance, scalability, and availability of the XenDesktop implementation. The following storage solutions are available for XenDesktop 7.x on Hyper-V 2012:

- Local Storage
- Direct Attached Storage (DAS)
- Storage Area Network (SAN)
- SMB 3.0 File Shares
- Windows Storage Spaces

Some XenDesktop configurations may require a combination of two or more solutions depending on the type of desktops and applications being deployed.

Decision: Local Storage

Local storage is the simplest storage to implement but provides no redundancy in case of a server failure. This type of storage is best suited for Hyper-V servers hosting random (pooled) desktops, or hosted shared desktops. Failover Clustering is not supported, but it isn't necessary since the user data does not persist with these types of desktops. Live migration is supported; however this does not protect the virtual machines in case of server failure.

Decision: Direct Attached Storage

Direct attached storage (DAS) offers similar functionality to local storage, but can also be shared between multiple computers. Therefore, Failover Clustering and live migration are supported. DAS devices are generally much cheaper when compared to a SAN solution, but do not scale as well. As the number of virtual desktops per host increases, additional disks are required to accommodate the number of IOPs generated. DAS devices are best suited for small to medium sized XenApp and XenDesktop

deployments.

Decision: Storage Area Network

A Storage Area Network (SAN) solution supports the Failover Clustering feature, so that the host server is not a single point of failure for the hosted virtual machines. SAN solutions however are expensive to implement, and usually require a high level of SAN expertise in order to manage. SANs can scale much higher than local storage, which makes them ideal for medium to large XenDesktop deployments, and Hyper-V servers hosting static (dedicated) desktops. SANs are also ideal for storing virtual servers supporting the XenDesktop infrastructure.

Decision: SMB 3.0 File Shares

SMB 3.0 file shares are a new feature in Windows Server 2012/2012 R2. This solution presents an SMB 3.0 file share as shared storage to a Hyper-V host or cluster. This solution requires a Windows Server 2012/2012 R2 file server, or a non-Microsoft file server that supports the SMB 3.0 protocol. SMB 3.0 file shares offer excellent performance, failover clustering and live migration at a low cost point.

Note: A loopback configuration where the Hyper-V host is also used as the file server for virtual machine storage is not supported.

When implementing high availability using Windows file servers, separate failover clusters are required for the file servers and the Hyper-V hosts. For SMB file shares to work properly with System Center Virtual Machine Manager, the file server, whether stand-alone or clustered, must not be running the Hyper-V role or assigned as a managed host in VMM. For more information on how to configure Windows Server 2012/2012 R2 file server clusters with Hyper-V, please refer to the Microsoft TechNet article – [Deploy Hyper-V over SMB](#).

Decision: Windows Storage Spaces

Windows Storage Spaces is a technology specific to Windows Server 2012/2012 R2. It can be best described as a DAS that behaves like a SAN. It works by grouping industry-standard disks or JBODs (just-a-bunch-of-disks) into storage pools. It supports failover clustering, live migration, and can scale like a SAN by allowing additional disks to be added to the pool. Storage Spaces is a low cost solution better suited for small to medium deployments of XenDesktop. For more information on Storage Spaces, please refer to the Microsoft TechNet article – [Storage Spaces Overview](#).

The following table summarizes the storage options available to Hyper-V and rates their suitability for XenDesktop deployments.

Hyper-V storage feature comparison

Storage Properties	Local	DAS	SAN	SMB 3.0 File Shares	Storage Spaces
Implementation costs	Low	Medium	High	Low	Low
Administration	Low	Medium	High	Medium	Medium
Performance	High	Med-High	High	Med-High	Med-High
Redundancy	Low-Med	Low-Med	High	Med-High	High
Live Migration	Supported; non-clustered hosts	Supported	Supported	Supported	Supported
Failover Clustering support	No	Yes	Yes	Yes	Yes
Scalability	Low	Low-Med	High	Med-High	Med-High
Best Use Case	Small production and test environments	Small to medium production environments	Medium to large production environments	Small to medium production environments	Small to medium production environments

Decision: RAID Level

The RAID level of a storage sub-system can have direct impact on the performance of the applications and workloads that it supports. To choose the optimal RAID level, it is necessary to consider the IOPS and read/write ratio generated by a given application or workload in combination with the individual capabilities of the

different RAID levels available. The following table outlines the key attributes of the most commonly used RAID levels:

RAID levels

RAID	Capacity	Fault Tolerance	Read Performance	Write Performance	Minimum # of Disks
0	100%	None	Very High	High (Write Penalty 1)	2
1	50%	Single-drive failure	Very High	Medium (Write Penalty 2)	2
5	67%-94%	Single-drive failure	High	Low (Write Penalty 4)	3
6	50%-88%	Dual-drive failure	High	Low (Write Penalty 6)	4
10	50%	Single-drive failure in each sub array	Very High	Medium (Write Penalty 2)	4

RAID levels 0, 1, and 10 are optimized for writes, however RAID 0 should not be implemented in a XenDesktop production environment since it offers no fault tolerance. RAID 1 and 10 offer the best performance for read/write operations but will only make use of 50% maximum disk capacity. RAID 5 and 6 makes the best use of disk capacity, but has a high write penalty. RAID 5 and 6 is suitable for all aspects of the XenDesktop solution except for the Provisioning Services write-cache. The Provisioning Services write-cache is highly write intensive therefore it should be hosted on RAID 1 or 10 storage.

Note: Storage Spaces supports just three resiliency types: Mirror which is the equivalent of RAID 1, Parity which is the equivalent of RAID 5, and Simple (no resiliency) which is the equivalent of RAID 0.

Decision: Virtual Disks

Hyper-V provides two options for creating virtual disks for virtual machines: fixed sized or dynamic. Virtual machines with dynamic disks can be provisioned faster, and make more efficient use of physical disk space. Virtual machine performance however, is faster when using fixed-size disks.

When using dynamic disks with Machine Creation Services or

Provisioning Services, the first boot of the virtual desktop requires expansion of the disk to accommodate changes required during the setup. During a boot storm, the expansion on the SAN can require significant resources. However, subsequent reboots do not impact Provisioning Services because the same disk, which is already expanded, is still used. Machine Creation Services however, disconnects the old disk, creates a new disk, expands the disk, boots the virtual desktop, and eventually deletes the old disk. This impacts storage every time a desktop is rebooted.

Due to the impact disk expansion has on virtual machine write performance, it is recommended to use fixed-sized disks for all virtual desktops and servers hosting applications.

Virtual Machine Manager

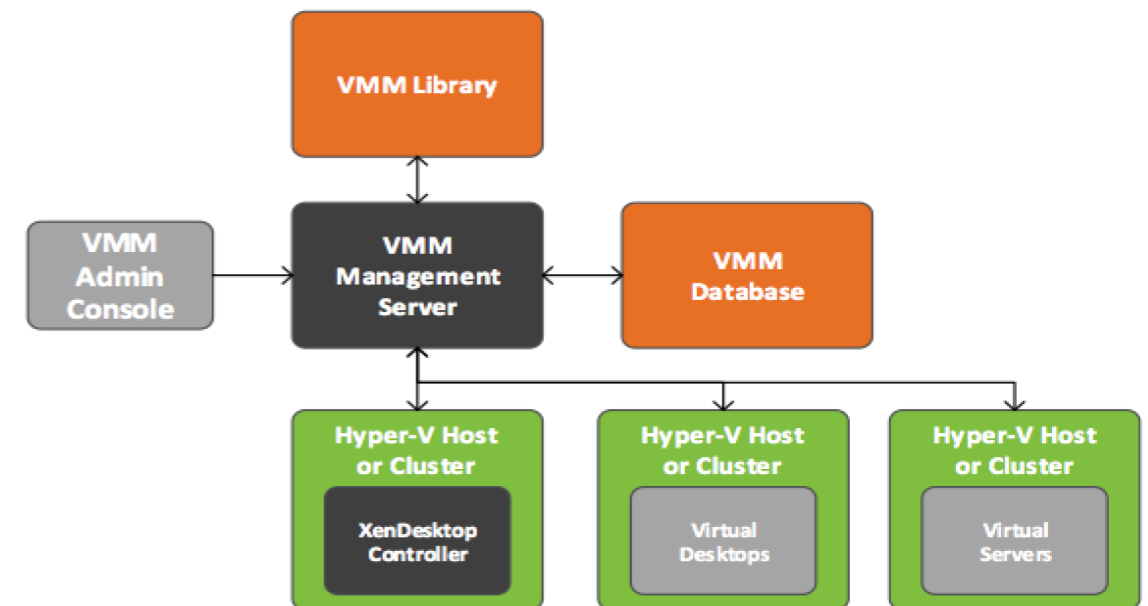
System Center 2012 R2 Virtual Machine Manager (VMM) is the key integration point between XenDesktop and Hyper-V 2012 R2. XenDesktop communicates directly with VMM to create, manage, and operate virtual machines.

There are four key components which make up the VMM role:

- Management server – The server on which the Virtual Machine Manager service runs. It is responsible for processing commands and controls communications with the VMM database, VMM library, and virtual machine hosts.
- Database – A Microsoft SQL database that stores the VMM configuration information such as virtual machine, service templates, and VMM profiles.
- Administrative console – The program that connects to the VMM management server to view and manage physical and virtual resources, such as virtual machine hosts, virtual machines, services, and library resources.
- VMM library – The catalog of resources used to deploy virtual

machines and services. Resources can include virtual hard disks, templates, and VMM profiles (hardware profiles, guest operating system profiles, application profiles, and SQL server profiles).

VMM Diagram



The VMM components can be installed on either physical or virtual servers. Within XenDesktop, a single desktop group can be associated with only one VMM server; however one VMM server can manage multiple desktop groups.

The VMM database is supported on SQL Server 2008 R2 Standard, Enterprise, and Datacenter editions and SQL 2012 Standard and Enterprise editions.

The VMM 2012 R2 administrative console can be installed on a Windows server or desktop operating system. The supported server operating systems are Windows Server 2012 and higher, and the supported desktop operating systems are Windows 7 and higher.

Decision: VMM Sizing and Scalability

The following considerations should be taken into account when determining the number of VMM 2012 R2 servers required to support a XenDesktop design:

- Citrix has found the best performance is achieved when each VMM 2012 R2 server is limited to managing 8000 virtual desktops. There is no limitation placed on the number of servers, or other control VMs also managed by each VMM server.
- While it is possible to run other applications on the VMM server, it is not recommended, especially other System Center 2012 R2 applications because they tend to have heavy resource demands and could significantly impact VMM performance.

The following table provides recommended specifications for the key VMM infrastructure components:

VMM recommendations

Component	Recommended	Recommended (>150 Hosts)
VMM Management Server		
CPU	Dual processor, dual core, 2.8GHz (x64) or greater	Dual processor, dual core, 3.6GHz (x64) or greater
RAM	4GB	8GB
Disk space (no local DB)	40GB	50GB
Disk space (local DB)	150GB	Use a dedicated SQL Server
Administrative Console		
CPU	Pentium 4, 1 GHz or greater	Pentium 4, Dual-Processor, 2 GHz or greater
RAM	2GB	4GB
Disk space	2 GB	4GB
Database		
CPU	Dual-Core 64-bit, 2GHz	Dual-Core 64-bit, 2.8GHz
RAM	4GB	8GB
Disk Space (local DB)	150GB	200GB
Library		
CPU	Dual-Core 64-bit, 3.2GHz or greater	Dual-Core 64-bit, 3.2GHz or greater
RAM	2GB	2GB
Hard disk space	Size will depend on what will be stored	Size will depend on what will be stored

[Click here to provide feedback](#)

Decision: VMM High Availability Design

Of the four key components which make up the VMM role, three of them (management server, database, VMM library) can be configured for high availability. A VMM server failure will not prevent users from connecting to their virtual desktops or applications, however, a VMM server failure will prevent the XenDesktop controllers from starting or provisioning new desktops.

When designing a highly available VMM solution the following considerations should be taken into account:

• VMM Management

- There can only be one implementation of a highly available VMM management server on a given failover cluster.
- The management server can be installed on as many as sixteen nodes on a failover cluster, but there can only be one node active at a time.
- During a planned failover, ensure that there are no active tasks running on the VMM management server. Any running tasks will fail during a failover, and any failed jobs will not start automatically after a failover.

• VMM Database

- Use a highly available installation of SQL server. SQL Server 2012 supports AlwaysOn Failover Cluster Instances, AlwaysOn Availability Groups, Database Mirroring, and Log Shipping. Database mirroring will be removed in a future version of Microsoft SQL Server. For more information, please refer to the Microsoft TechNet article – [High Availability Solutions \(SQL Server\)](#).
- Microsoft recommends implementing the SQL HA installation on a failover cluster separate from the failover cluster on which the VMM management server is installed. For more information, please refer to the Microsoft TechNet

article – [Installing a Highly Available VMM Management Server](#).

- **VMM Library**

- For a complete HA configuration consider placing the VMM library on a highly available file server. A clustered group of file servers is recommended.

Note: XenDesktop does not require the VMM library to be highly available to function. XenDesktop will not be able to provision new virtual machines if the VMM library is not available, however existing virtual machines will not be impacted.

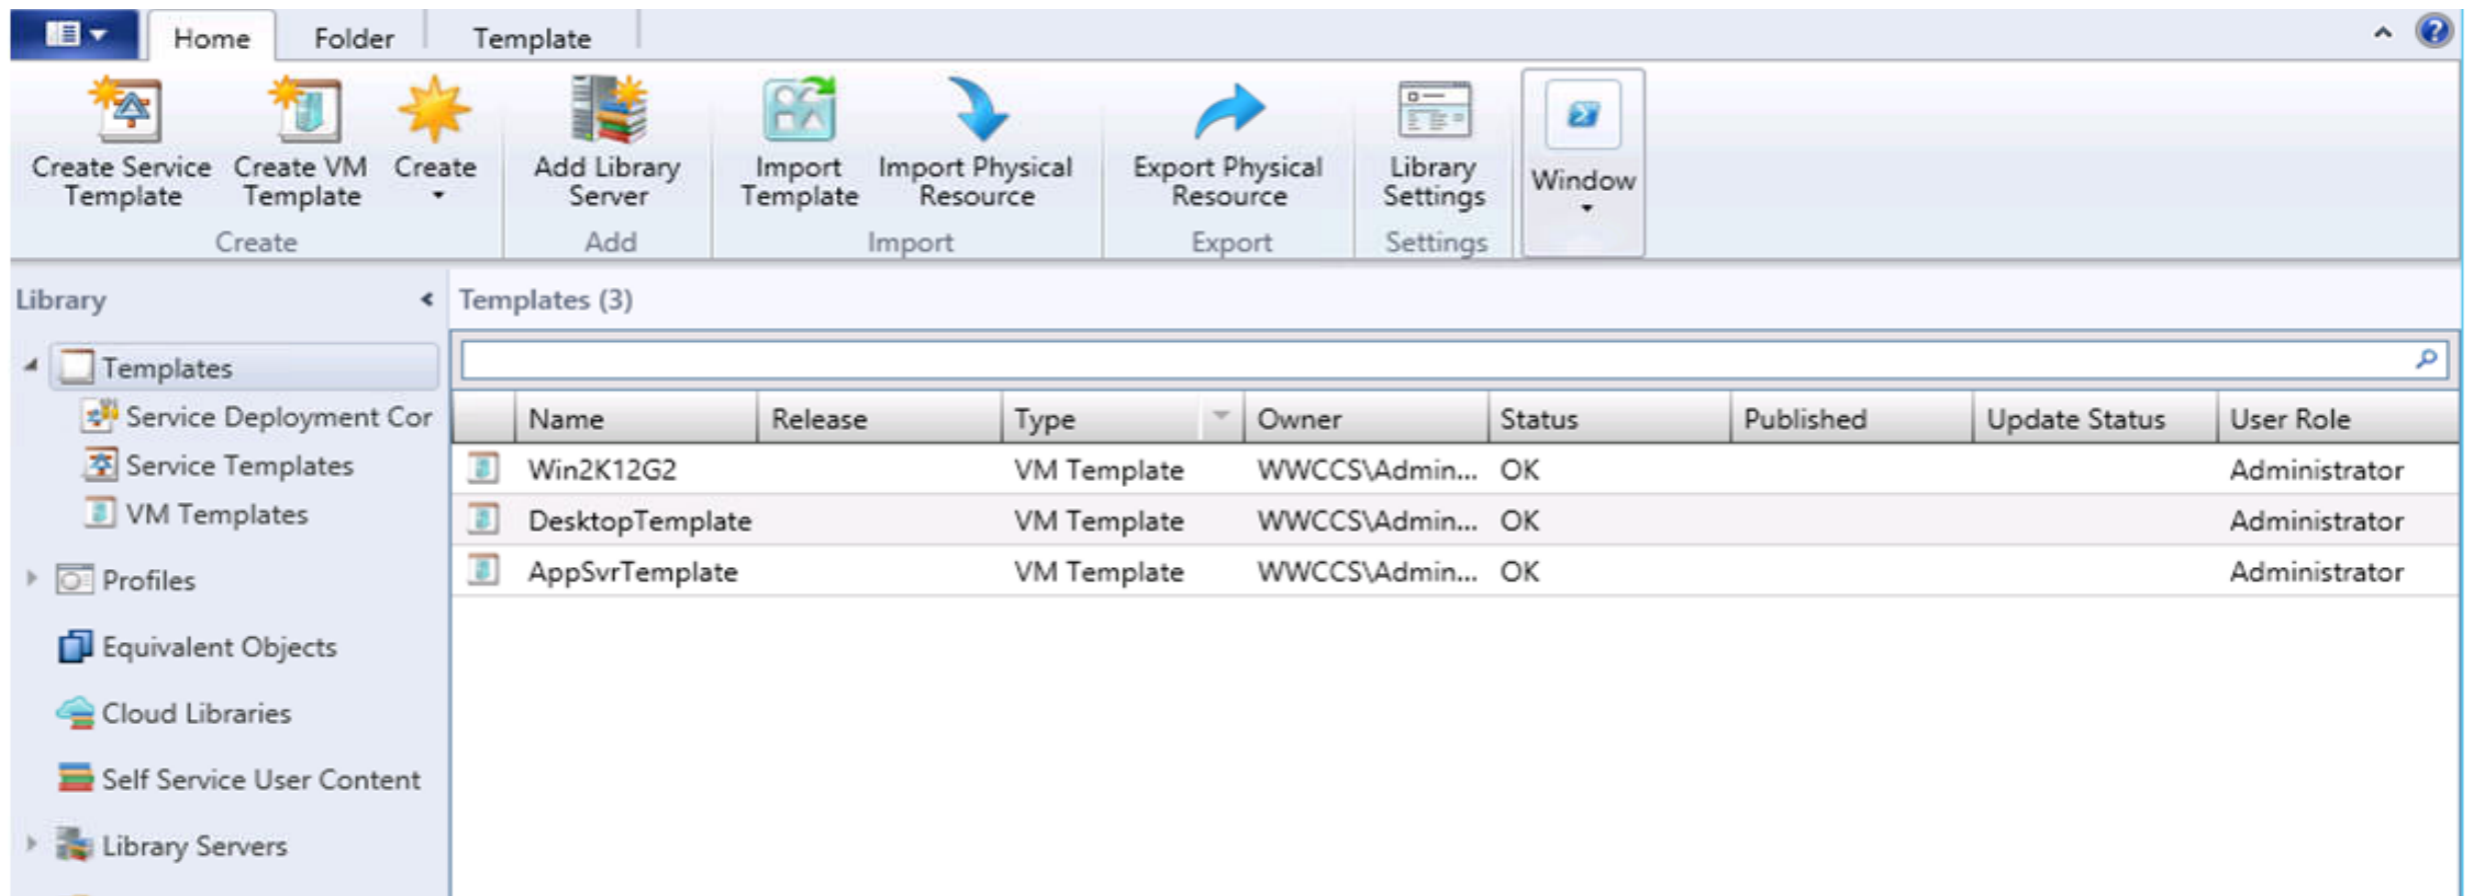
VM Provisioning

This section covers the decisions when considering the XenDesktop provisioning methods on Hyper-V 2012 R2.

Decision: Templates

XenDesktop uses templates to create virtual machines. The templates are stored within the VMM library to simplify deployment across all managed hosts. When new virtual machines are created from the template, they receive a unique identity including a MAC address, IP address, and computer name. An example of the templates in a VMM Library is shown below.

Templates in VMM Library



Note: Citrix Consulting recommends creating templates for servers that will support the XenDesktop infrastructure (e.g. Delivery Controllers, Provisioning Servers, StoreFront servers, etc.) This will allow servers to be provisioned quickly and consistently.

When building templates for use with Provisioning Services or Machine Creation Services, the following should be taken into consideration:

- Ensure that sufficient disk space is allocated for each desktop type required. Separate templates will be required if different NICs per desktop are required. The Provisioning Services Setup Wizards will allow for the adjustment of vCPUs, RAM, and write cache disk size to be specified based on a single template.
- Use fixed disks instead of dynamic for optimal performance. Using fixed disks will also prevent the Cluster Shared Volume from running out of disk space as the VHDs expand.
- It is easier to expand a disk that is too small than to shrink a disk that is too big. Therefore, start with a small disk and manually expand it later if necessary.

Note: To assist with automating the build process, Citrix developed the XenDesktop 7.1 Service Template for System Center 2012 Virtual Machine Manager. The service template can be used to deploy XenDesktop environments more rapidly and consistently. For more information about the service template please refer to the [Introducing the XenDesktop 7.1 Service Tech Preview](#) blog.

High Availability

Any downtime within a XenDesktop solution may result in lost revenue and productivity. Therefore, it is important that all components of the infrastructure are considered when designing a highly available XenDesktop solution. XenDesktop is built around fault tolerant components that can be further enhanced with Citrix

NetScaler to provide disaster recovery and business continuity.

Decision: Failover Clustering

Hyper-V Failover Clustering functionality provides support for specific virtual machines and applications that are mission critical and require high availability and reliability. A Hyper-V failover cluster allows a group of host servers to be managed as a single physical entity. The failover cluster operates by moving resources between hosts to maintain service if hosts become unavailable.

When a failure occurs on one host, the resources are moved without any user or management intervention. During the failover process, the end user may notice a slight pause with their XenDesktop session for a few seconds, but once the new host establishes management of the resource, the end user will be able to continue working as normal.

In order to setup failover clustering in Hyper-V 2012 R2 the following is required:

- Windows Server 2012 R2 with the Hyper-V role and Failover Clustering feature enabled, or Hyper-V Server 2012 R2 with Failover Clustering feature enabled.
- A network infrastructure that supports a cluster configuration. There should be no single point of failure in the infrastructure, e.g. multiple data paths available to shared storage, etc.
- A storage solution that is accessible to all hosts in the cluster.
- Host servers with the same or similar hardware components are recommended.
- If using Serial Attached SCSI or Fibre Channel in the host servers, all elements of the storage stack should be identical. It is recommended that the host bus adapter (HBA), HBA drivers, and HBA firmware be identical. If dissimilar HBAs are used, then verify with the storage vendor that the configuration is supported.

- If using iSCSI, the network adapters used to connect to the iSCSI target should be identical. The network used for iSCSI should not be used for any other communication.

If sufficient capacity exists, these key XenDesktop infrastructure components should be configured for Hyper-V failover clustering to ensure that a single host failure does not affect the availability or performance of the XenDesktop environment:

- StoreFront / Web Interface
- XenDesktop Controllers
- XenApp Controllers
- License Server
- Provisioning Services
- SQL Servers
- NetScaler VPX

In addition, static (dedicated) desktops, and random (pooled) desktops using Personal vDisk (PvD) should reside within the Hyper-V failover cluster. For random (pooled) desktops not using PvD, or hosted shared desktops, it is not necessary to place them within a Hyper-V failover cluster since they do not contain any user specific data or applications.

Decision: Cluster Shared Volumes

Cluster Shared Volumes (CSV) allow nodes within a failover cluster to have simultaneous access to the same file system. By using CSV, multiple virtual machines can use the same disk, and fail over independently of one another. CSVs are required for Failover Clustering and Live Migration functionality. Virtual servers supporting the XenDesktop infrastructure and dedicated desktops are typically hosted on CSVs.

The following recommendations should be considered during the Cluster Shared Volume design:

[Click here to provide feedback](#)

- For fault tolerance consider using multiple cluster networks to carry CSV traffic, or consider using teamed network adapters.
- Microsoft recommends that the CSV communications take place over a dedicated network not shared by other traffic.
- All network adapters used to carry cluster communication should have the following properties enabled:
 - Client for Microsoft Networks
 - File and Printer Sharing for Microsoft Networks
 - Microsoft Failover Cluster Virtual Adapter Performance Filter
- To make the best use of CSV for virtual machines, create separate CSVs; one for virtual machine system files and page file, and another for data files.
- When planning the storage configuration for virtual machines, consider the disk requirements of the service and applications that the virtual machine will support. Understanding these requirements will help avoid poor performance due to disk contention.

For more information on CSVs, please refer to the Microsoft TechNet article – [Use Cluster Shared Volumes in a Failover Cluster](#).

Decision: CSV Cache

The Cluster Shared Volumes cache is a new feature in Windows Server 2012. The CSV cache uses RAM to cache read-only unbuffered I/O, which improves Hyper-V performance since it conducts unbuffered I/O when accessing a VHD or VHDX file. This greatly improves the performance of Machine Creation Services since the base VHDX file can be read and cached in memory. By default the CSV cache is disabled, therefore consider enabling it for XenDesktop deployments.

There are considerations when planning to deploy the CSV cache:

- Windows Server 2012 R2 will allow you to allocate up to 80% of the total physical RAM for CSV write-through cache, which is consumed from non-paged pool memory. Microsoft testing has found that 512MB is a good starting point for the CSV cache allocation. Since system memory is a contended resource on a Hyper-V cluster, it is recommended to keep the CSV cache size to 512MB, 1GB, or 2GB.
- The cache size can be modified without requiring a server reboot in Windows Server 2012 R2, however enabling the CSV cache on an individual disk requires that the Physical Disk resource is taken offline then back online, for the change to take effect.
- CSV cache will be disabled on:
 - Tiered Storage Space with heat map tracking enabled.
 - De-duplicated files using in-box Windows Server Data De-duplication.
Note: Data will instead be cached by the dedup cache.
 - Resilient File System (ReFS) volume with integrity streams enabled
Note: NTFS is the recommended file system for virtual machines in a production environment.
 - When the CSV is in File System Redirected Mode.

Monitoring

Designing and implementing an appropriate monitoring solution for Hyper-V will help to maintain the consistency, reliability, and performance of the XenDesktop infrastructure.

[Click here to provide feedback](#)

Decision: Management Tool

It is imperative that the monitoring tool selected includes support for Hyper-V. System Center Virtual Machine Manager supports Dynamic Optimization and Power Optimization on Hyper-V host clusters. Dynamic Optimization performs load balancing within the host clusters that support Live Migration. It will automatically migrate virtual machines between hosts once a set threshold is triggered. This improves load balancing among hosts and helps to ensure that a single host is not constrained.

Power Optimization allows Virtual Machine Manager to save energy by turning off hosts that are not needed to meet resource requirements within a host cluster, and turns them back on again when they are needed. This feature requires the host server to have a baseboard management controller (BMC), and Dynamic Optimization must be enabled. For more information please see the Microsoft TechNet article – [Configuring Dynamic Optimization and Power Optimization in VMM](#).

Virtual Machine Manager can also be integrated with System Center Operations Manager (SCOM) to monitor the health and performance of virtual machine hosts and their virtual machines. It can also monitor the health and availability of the VMM management server, VMM database server, and VMM library servers. For more information please see the Microsoft TechNet article – [Configuring Operations Manager Integration with VMM](#).

Note: SCOM Management Packs for XenDesktop are available through Comtrade. For more information please see Comtrade's [Management Packs for Citrix](#) webpage.

Hyper-V Hardware

Vendor supplied tools should be used where possible to capture and analyze bespoke hardware events.

The high availability feature of Hyper-V can be used to mitigate

the impact from a single host failure. However, there will still be an element of downtime while the virtual machines are restarted on alternative hosts. Effective proactive monitoring can help to avoid the unnecessary invocation of high availability and the restart of virtual workloads.

Hyper-V Performance

The overall performance of each Hyper-V host will affect the performance of the virtual machines that it supports. It is critical that the Hyper-V hosts are monitored for bottlenecks within the processor, memory, disk and network subsystems. For more information please refer to Microsoft blog post [Hyper-V Performance Monitoring](#).

Backup and Recovery

An important piece of the design process will involve having a strategy in place for backup and recovery of the XenDesktop environment.

Decision: Backup Method

There are several factors to consider including, the different types of backups you can make, the state of the virtual machines, and the type of storage being used by the virtual machines. Backups can occur at two levels on the hypervisor:

- Backup the Hyper-V host – This method is preferred because it captures more data than performing the backup from within the guest operating system, including the configuration of virtual machines and snapshots associated with virtual machines.
- Backup within the guest virtual machines – A backup application runs within the virtual machine. Use this method when backing up data from storage that is not supported by the Hyper-V Volume Shadow Copy Service (VSS) writer.

Decision: Components to Backup

The following Hyper-V and XenDesktop components should be backed up so that it is possible to recover from a complete failure:

- XenDesktop database
- XenApp database
- Provisioning Services database
- Provisioning Services vDisks (virtual desktops and hosted app servers)
- Hyper-V VMM database
- Hyper-V VMM library
- Hyper-V host system state
- Hyper-V host local disks
- Hyper-V Cluster quorum
- Static (dedicated) virtual desktops
- StoreFront configuration
- License files
- User profiles / home folders

Note: It is assumed that there is a fast automated rebuild process in place for the servers supporting the XenDesktop infrastructure (XenDesktop controller, StoreFront server, Provisioning Server, etc.). If this assumption is not true then all infrastructure servers must also be backed up.

Virtual networks are not included in a full server backup. You will need to reconfigure the virtual networking by recreating the virtual networks and then reattaching the virtual network adapters in each virtual machine to the appropriate virtual network. Make sure the virtual network configuration and all relevant settings are documented as part of the backup process.

Storage

This chapter provides guidance for planning and working with various storage architectures and technologies. This chapter does not provide architectural blueprints because of the wide variety of storage systems and solutions available on the market. Therefore it is strongly advised to involve storage vendors or specialists in the planning process.

Decision: Storage Architecture

Four primary storage architectures are available:

- **Local Storage** - Uses hard disks directly attached to the computer system. The disks cannot be shared with other computer systems, but if the computer is hosting pooled or hosted shared desktops, a shared storage solution is not necessary. In many cases local storage can perform as well as shared storage. Scalability is limited to the number of drive bays available in the computer system. Many blade servers for example have just two drive bays, so using local storage to support a XenDesktop deployment may not be optimal.
- **DAS** - Storage sub-system directly attached to a server or workstation using a cable. It uses block-level storage and can be a hard disk local to the computer system or a disk shelf with multiple disks attached by means of external cabling. Unlike local disks, disk shelves require separate management. Storage shelves can be connected to multiple servers so the data or disks can be shared.
- **NAS** - Provides file-level storage to computer systems through network file shares. The NAS operates as a file server, and NAS systems are networked appliances which contain one or more hard drives, often arranged into logical, redundant storage containers or RAID arrays. Access is typically provided using standard Ethernet and network file sharing protocols such as

NFS, SMB/CIFS, or AFP.

Note: NAS can become a single point of failure. If the network share becomes unavailable, all target devices streamed from the disk will be unavailable as well.

- **SAN** - Dedicated storage network that provides access to consolidated, block-level storage. SANs allow computers to connect to different storage devices, so no server has ownership of the storage subsystem enabling data to be shared among multiple computers. A SAN will typically have its own dedicated network of storage devices that are generally not accessible through the network by standard means. In order to connect a device to the SAN network a specialized adapter called the Host Bus Adapter (HBA) is required. SANs are highly scalable with no noticeable change in performance as more storage and devices are connected. SANs can be a costly investment both in terms of capital and the time required to learn, deploy and manage the technology.
- **Hybrid** - A NAS head refers to a NAS which does not have any on-board storage, but instead connects to a SAN. In effect, it acts as a translator between the file-level NAS protocols (NFS, CIFS, etc.) and the block-level SAN protocols (Fibre Channel and iSCSI). Thus it can combine the advantages of both technologies and allows computers without Host Bus Adapters (HBA) to connect to centralized storage.

The following table summarizes the storage options available and rates their suitability for XenDesktop deployments.

Storage feature comparison

Storage Properties	Local	DAS	NAS	Recommended (>150 Hosts)
Implementation costs	Low	Medium	Medium	High
Administration	Low	Medium	Medium	High
Performance	High	Med-High	Med-High	High
Redundancy	Low-Med	Medium	Med-High	High
Scalability	Low	Low-Med	Med-High	High
Typical use case	Small to medium production and test environments.	Small to medium production environments.	Small to medium production environments.	Medium to large production environments.

Note: Hyper-V 2008 R2 does not support NAS technology. Hyper-V 2012/2012 R2 only supports NAS solutions that support the SMB 3.0 protocol. For more information please refer to the [Hyper-V 2008 R2](#) and [Hyper-V 2012 R2](#) sections of the handbook.

Local storage is best suited for storing virtual machines which do not have high availability requirements or persistent data attached such as random (pooled) desktops or hosted shared desktops. Local and DAS is suited for storing user data and home directory files. If using Machine Creation Services, master images as well as any updates must be replicated to each server.

NAS and SAN storage is best suited for infrastructure servers supporting the XenDesktop environment, and virtual machines with persistent data such as static (dedicated) desktops, and random (pooled) desktops with Personal vDisks.

Decision: RAID Level

To choose the optimal RAID level, it is necessary to consider the IOPS and read/write ratio generated by a given application or workload in combination with the individual capabilities of a RAID level. For hosting read intensive workloads, such as the Provisioning Services vDisk store, RAID levels that are optimized for read

operations such as RAID 1, 5, 6, 10 are optimal. This is because these RAID levels allow read operations to be spread across all disks within the RAID set simultaneously.

For hosting write intensive workloads, such as Provisioning Services write cache and Machine Creation Services differencing disks, RAID levels such as RAID 1 or 10 are optimal, as these are optimized for writes and have a low write penalty.

The following table outlines the key quantitative attributes of the most commonly used RAID levels:

RAID levels

RAID	Capacity	Fault Tolerance	Read Performance	Write Performance	Minimum # of Disks
0	100%	None	Very High	High (Write Penalty 1)	2
1	50%	Single-drive failure	Very High	Medium (Write Penalty 2)	2
5	67-94%	Single-drive failure	High	Low (Write Penalty 4)	3
6	50-88%	Dual-drive failure	High	Low (Write Penalty 6)	4
10	50%	Single-drive failure in each sub array	Very High	Medium (Write Penalty 2)	4

Note: The write penalty is inherent in RAID data protection techniques, which require multiple disk I/O requests for each application write request, and ranges from minimal (mirrored arrays) to substantial (RAID levels 5 and 6).

Decision: Numbers of Disks

To determine the number of disks required it is important to understand the performance characteristics of each disk, the characteristics of the RAID level and the performance requirements of the given workload. The basic calculation for determining the total number of disks needed is:

$$\text{Total \# of Disks} = \frac{(\text{Total Read IOPS} + (\text{Total Write IOPS} * \text{RAID Penalty}))}{\text{Disk Speed IOPS}}$$

For example, a disk manufacturer is reporting that a particular disk array which they have developed has a total workload IOPS of 2000. The raw IOPS per disk is 175. To determine how many disks are required to support a workload with 20% read operations and 80% write operations on RAID 10:

$$\text{Total \# of Disks} = \frac{((20\% * 2000) + (80\% * 2000) * 2)}{175} = 20.57 \text{ or } 21 \text{ Disks}$$

Based on the previous example, the following table shows how the disk count will vary based on the RAID level and the read/write ratio.

Example of how disk count changes per RAID level and R/W ratio

RAID	RAW IOPS (per disk)	Workload IOPS	Read %	Write %	Disk count
0	175	2000	20%	80%	12
	175	2000	80%	20%	12
1/10	175	2000	20%	80%	21
	175	2000	80%	20%	14
5	175	2000	20%	80%	39
	175	2000	80%	20%	19

Decision: Disk Type

Hard disk drives (HDDs) are the traditional variation of disk drives. These kinds of disks consist of rotating platters on a motor-driven spindle within a protective enclosure. The data is magnetically written to and read from the platter by read/write heads.

Different implementations of this technology are available on the market, which differ in terms of performance, cost and reliability.

- Serial ATA (SATA) disk transmit data serially over two pairs of conductors. One pair is for differential transmission of data and the other pair is for differential receiving of data. SATA drives are widely found in consumer desktop and laptop computers. Typical SATA drives are transfer speeds ranging from 1500-

6000 Mbps and support hot-swapping by design.

- Small Computer Systems Interface (SCSI) disks use a buffered, peer to peer interface that uses handshake signals between devices. Many SCSI devices require a SCSI initiator to initiate SCSI transactions between the host and SCSI target. SCSI disks are common in workstations and servers and have throughputs ranging from 40 – 5120Mbps. iSCSI (Internet Small Computer System Interface) is a mapping of the regular SCSI protocol over TCP/IP, more commonly over Gigabit Ethernet.
- Fibre Channel (FC) disk is the successor to the parallel SCSI disk and is common in SAN storage devices. Fibre Channel signals can run on an electrical interface or fibre-optic cables. Throughput can range from 1 – 20Gbps, and connections are hot-pluggable.
- Serial Attached SCSI (SAS) disk uses a new generation serial communication protocol to allow for higher speed data transfers than SATA disks. Throughput can range from 2400 – 9600Mbps.

In contrast to traditional hard disks, Solid State Disks (SSDs) use microchips to retain data in either NAND non-volatile memory chips (flash) or DRAM and contain no moving parts. SSDs are less susceptible to physical shock, have lower access times and latency and have higher I/O rates. SSDs have significantly higher random read performance. An SSD drive can attain anywhere from 5,000 to 20,000 random reads per second. SSDs are also more expensive per gigabyte (GB) and typically support a limited number of writes over the life of the disk.

Flash memory-based SSDs can be either based on multi-level cells (MLC) or single-level cells (SLC). SLC devices only store one bit of information in each cell. MLC devices can store multiple bits of information with each cell. Flash based SSDs cost lower than DRAM based SSDs but perform slower. DRAM based SSD devices are used primarily to accelerate applications that would otherwise be held back by the latency of flash SSDs or traditional HDDs.

SSDs were previously not viable for enterprise storage solutions because of the high cost, low capacity and fast wear of the drives. Improvements in SSD technology and lowering costs are making them more favorable over HDDs. Solid state hybrid drives (SSHD) combine the features of SSDs and HDDs, by containing a large HDD drive with an SSD cache to improve performance of frequently accessed data.

Comparing SSDs and HDDs is difficult since HDD benchmarks are focused on finding the performance aspects such as rotational latency time and seek time. As SSDs do not spin, or seek, they may show huge superiority in such tests. However SSDs have challenges with mixed reads and writes and their performance may degrade over time.

The following table compares the transfer rates of some of the more common storage types available on the market today.

Some common disk types and transfer rates

Technology	Rate (bit/s)
iSCSI over Fast Ethernet	100 Mbps
Ultra-2 wide SCSI (16 bits/40 MHz)	640 Mbps
iSCSI over Gigabit Ethernet	1,000 Mbps
SATA rev 3	6,000 Mbps
SAS 3	9,600 Mbps
FCoE over 10GbE	10,000 Mbps
SATA rev 3.2 – SATA Express	16,000 Mbps
iSCSI over Infiniband	32,000 Mbps

SCSI and SATA disks are best suited for storing data that does not have high performance requirements like the PVS vDisk store. SAS, Fibre Channel, or SSD drives are best suited for storing data that have high performance requirements like the PVS write cache.

Decision: Storage Bandwidth

Storage bandwidth is the connectivity between servers and the storage subsystem. Understanding bandwidth requirements can help determine the proper hardware for delivering data and applications at speeds for a positive end user experience. For most datacenters 10Gbps Ethernet or 10Gbps FCoE is sufficient for storage connections. Smaller environments however may only need 1Gbps bandwidth. In virtualized environments it is not just important to look at the bandwidth requirements of the physical host and storage subsystem, but determining how much bandwidth is required for each virtual machine plays a factor too.

In order to plan for the required bandwidth, it is necessary to determine the throughputs for every individual system that uses a shared component or network path. For example, the following information is provided for an environment with 100 similar virtual machines (hosted on 10 virtualization hosts and connected to one NAS head)

Throughput example

	Average	Peak
Throughput per VM	10 Mbps	30 Mbps
Throughput per host	100 Mbps (10 VMs * 10 Mbps)	300 Mbps (10 VMs * 30 Mbps)
Throughput per storage	1 Gbps (10 hosts * 100 Mbps)	3 Gbps (10 hosts * 300 Mbps)

The NIC used for storage communication needs to be a 1Gbps adapter in order to handle the peak load. The NAS head as well as its network connection need to support 3Gbps worth of data traffic in order to support the peak load of all systems.

Decision: 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.

For example, a XenDesktop implementation using tiered storage may look like the following:

- **Tier 1 storage group** – Write intensive files such as the write cache and differencing disks are placed in a storage group consisting of SSDs.
- **Tier 2 storage group** – Mission critical data, or data that requires high availability such as Personal vDisks, are placed in a storage group consisting of less expensive high performing drives.
- **Tier 3 storage group** – Seldom used data files, read-only files, or other non-mission critical data placed in a storage group consisting of low cost and lower performing drives.

Decision: Thin Provisioning

Thin provisioning allows more storage space to be presented to the virtual machines than is actually available on the storage repository. This lowers storage costs by allowing virtual machines access to disk space that is often unused. This is particularly beneficial to Machine Creation Services which uses a linked-clone approach to provisioning virtual machines. Thin provisioning minimizes the storage space required for the master image copies used to build virtual machines. Thin provisioning is possible at the physical storage layer, a feature usually available with most SAN solutions, and at the virtual layer. NFS based storage solutions will usually have thin provisioning enabled by default.

At the physical storage layer it is important to ensure that sufficient storage is available to prevent the risk of virtual machines not being available in a storage “overcommit” scenario when available disk space is exhausted. Organizations should decide if the cost savings thin provisioning provides outweighs the associated risk and consider enabling if the storage solution supports it.

Note: Virtual machines may not function if disk space is exhausted, so it is important to have a process in place, either through alert or notifications that will give administrators enough time to add more disks to the storage solution so that the XenDesktop environment is not impacted.

Decision: Data De-Duplication

Data de-duplication is a data compression technique whereby duplicate data is replaced with pointers to a single copy of the original item. This reduces storage requirements and costs by improving storage utilization, however it can impact storage performance. There are two implementations of de-duplication available:

- **Post-process de-duplication** – The de-duplication is performed after the data has been written to disk. Post-process de-duplication should be scheduled outside business hours to ensure that it does not impact system performance. Post Process de-duplication offers minimal advantages for random desktops as the write-cache/difference disk is typically reset on a daily basis.
- **In-line de-duplication** – Examines data before it is written to disk so that duplicate blocks are not stored. The additional checks performed before the data is written to disk can sometimes cause slow performance. If enabled, in-line duplication should be carefully monitored to ensure that it is not affecting the performance of the XenDesktop environment.

If the storage solution supports it, enabling post-process data de-duplication is recommended for minimal impact to XenDesktop performance.

Disaster Recovery

Decision: Datacenter Utilization

XenApp and XenDesktop deployments can leverage multiple datacenters to improve user performance and the availability of resources. When deploying multiple datacenters, a critical decision is determining whether the datacenters will be in an active/active or active/passive configuration.

An active/active configuration allows resources in each datacenter to be used efficiently by allowing different user groups to access resources in each datacenter. In contrast, an active/passive datacenter configuration consists of an active and cold standby site, which users access in the event of a failure in the primary datacenter. Having a completely passive datacenter for disaster recovery is not ideal since it would require a major investment in hardware that would only be utilized in the event of a disaster.

Due to the complexities of actively consuming user data and backend database resources between datacenters, an active/active configuration with site affinity is recommended. This allows for both datacenters to be utilized, although individual users are tied to a specific location.

Decision: Datacenter Connectivity

An active/active datacenter configuration should utilize GSLB (Global Server Load Balancing) to ensure users will be able to establish a connection even if one datacenter is unavailable.

Although GSLB can dynamically assign users between datacenters based on proximity to the datacenter or load metrics, it is recommended to statically assign users to a specific datacenter. Dynamic distribution of users is only possible with users that have no user data or in cases where the datacenters are linked with high speed connections that allow the same file server to be actively used between locations.

Therefore, having an active/active configuration with users statically assigned to a specific location is the most optimal solution. This will allow users to be split between datacenters based on categories such as location, functional group, or applications. A separate repository for user data such as Windows profiles and documents would be located in each datacenter. Although user data mirrored on different file servers cannot be actively consumed between datacenters, it can be mirrored in a passive configuration to another datacenter for failover purposes.

Affinity to a specific datacenter can be accomplished using features such as Group Extraction on NetScaler Gateway or Optimal Routing and User Mapping with StoreFront. In the event of a failure, users would be automatically redirected to the other datacenter where a different set of resources could be made available.

In any multi-datacenter configuration, it is recommended that a separate XenApp/XenDesktop site is created in each location along with its own set of infrastructure components. This includes separate SQL servers, Delivery Controllers, and NetScaler appliances. This reduces the risk and complexity of sharing resources between datacenters. Citrix does not officially support a single XenDesktop or XenApp site that spans between multiple datacenters. Certain components such as StoreFront and the Citrix License server, which are not directly tied to a specific XenDesktop or XenApp site can be shared between datacenters. For component level high availability, components such as StoreFront servers can be load balanced between multiple datacenters.

The following items should also be considered before deciding upon an active/active configuration:

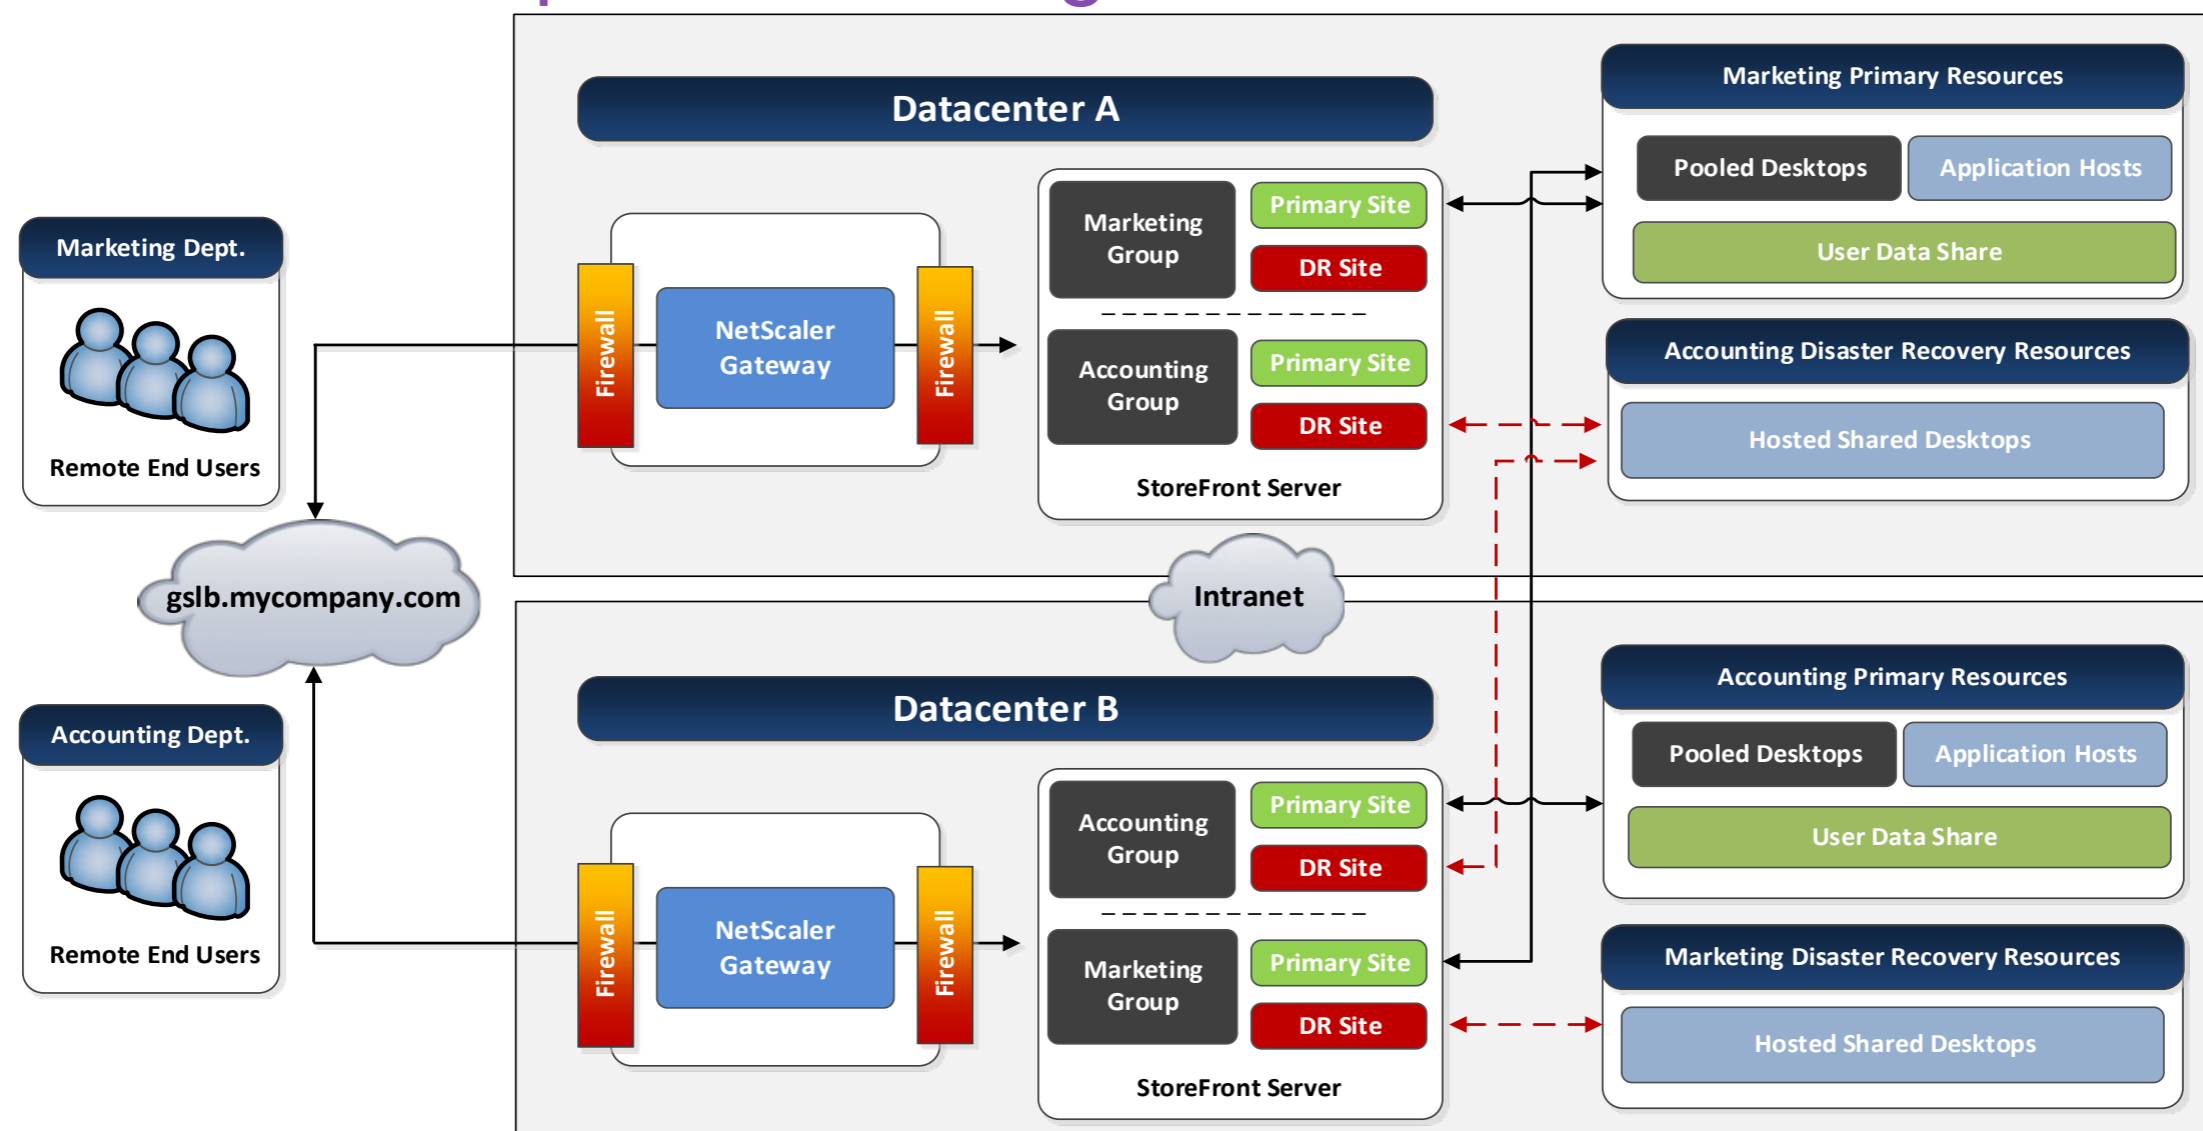
- **Datacenter Failover Time** – Although automatic failover to resources in another datacenter can be configured, the decision must be made whether those resources will constantly be running or only placed online in the event of a failover. If the resources must be manually provisioned, the failover time may

be hours rather than minutes.

- **Application Servers** – It must be determined if the application servers are able to support an active/active configuration between datacenters. If an active/active configuration is supported, a seamless failover to secondary datacenter is possible. If the backend application servers only support a passive secondary location, there must be manual intervention to ensure the application is properly restored in the secondary datacenter increasing the failover time.
- **StoreFront Optimal Routing** – The optimal NetScaler Gateway routing feature allows for resources located at separate physical

locations that are aggregated inside StoreFront to be routed through the NetScaler Gateway appliance in their respective datacenter. This allows the Citrix session to be routed in the most efficient way possible therefore providing a better user experience. By default, the connection for the resource will be routed through the NetScaler Gateway that the user initially connected through. This means without the optimal routing configuration, the Citrix session would transverse the link between datacenters instead of going through the NetScaler Gateway in the same physical location. More information about this feature can be found on the eDocs article — [Configure optimal NetScaler Gateway](#).

StoreFront Optimal Routing



Experience from the Field

Manufacturing – A large manufacturing company has two datacenters setup in an active/active configuration each with their own separate Citrix infrastructure. An active/active configuration is chosen to more efficiently utilize resources in each datacenter and mitigate the risk if a failure occurs in a single datacenter. Since the Pooled Desktops and Application servers require persistent configurations be stored inside the Windows profile, users are statically assigned to a specific datacenter based on their department. Affinity to a specific datacenter is accomplished by using the user mapping & failover features of StoreFront which allows users assigned to a specific Active Directory security group to only access applications from their primary datacenter. In the event of a failure and the primary resources are unavailable, secondary resources will be made available.

In the diagram above, GSLB allows users of the Marketing and Accounting departments to land in either Datacenter A or Datacenter B. Since the primary resources for each group only resides in a single datacenter, the StoreFront user mapping & failover functionality along with NetScaler Optimal Routing is used to ensure that no matter which datacenter the user initially connects to, only the resources in their primary datacenter will be presented and the data flow will return through the NetScaler Gateway in the datacenter where the desktop or application is hosted.

For example, if the Pooled Desktops & Application Servers made available to the Marketing department are in Datacenter A are unavailable, the failover functionality inside StoreFront will allow disaster recovery resources, in this case Hosted Shared Desktops in Datacenter B to be presented to the end user. While there will be no personalization available, the Hosted Shared Desktops will have a base set of applications that will allow the Marketing Department users to function until access to their primary resource is restored.

Decision: Capacity in Secondary Datacenter

Establishing the capacity of the secondary datacenter is vital for a smooth transition in the event of a failure. While ideal, having a secondary datacenter with the same capacity as the primary datacenter may not be feasible. The deterring factors to this decision are primarily cost and management to support full capacity in each datacenter. Therefore, it should be determined what percentage of total users or what percentage of users for a specific application the secondary datacenter will facilitate. This value can be determined by conducting an analysis of the impact to business for different user groups, and determining which groups will be affected the most if access to their resources are lost.

If it is determined that only a subset of users will gain access to the failover datacenter, measures should be taken to limit access to only the selected users. This will prevent the failover datacenter from being overloaded with resource requests it cannot handle. This can be accomplished by assigning users to specific Active Directory groups and then limiting access on the NetScaler Gateway in the failover datacenter. This can also be accomplished by using the Group Extraction functionality of the NetScaler Gateway. For example, it might be determined that only five of the fifty applications delivered via XenApp are mission critical in the event of an outage. For each of the five applications, it could be determined that only a small percentage of users actually need access in the event of a failover. Focusing on just these core applications and users along with their backend databases rather than providing access to every application will make the failover process much easier.

Once the capacity of the secondary datacenter is designed, the type and amount of resources that will be made available in a

failover scenario must be determined. If users have resources such as dedicated desktops or desktops with personal vDisks assigned to them in the primary datacenter, it is not possible to seamlessly failover these resources. Since the time spent in the secondary datacenter should be minimal, users only require the essential resources necessary for their role. It is recommended that another FlexCast model such as Pooled Desktops or Hosted Shared Desktops is selected for the secondary datacenter in the event of a failure. For example, if dedicated desktops are the primary resource, a Hosted Shared desktop can be provided in the secondary datacenter with the most business critical applications installed. User personalization settings will not be available, but productivity can continue.

Experience from the Field

Healthcare – A medical group has two datacenters that are configured in an active/active configuration with site affinity based on Active Directory user groups. The medical group has hundreds of applications published, although there is only a set of five core applications that are critical to the business in the event of a datacenter failure. One of the five applications is an Electronic Medical Record (EMR) application which has the ability to support the full user base in the secondary datacenter in the event of a failure. The other four applications have been designed to only support thirty percent of the required user capacity. It was decided that only thirty percent of the user base for each of the other four applications would be provided in the secondary datacenter due to the capacity needed to support all users.

The medical group used the user mapping & failover features of StoreFront to limit access to applications in the failover datacenter. This was accomplished by creating a new Active Directory group and assigning the thirty-percent of users they require to have access to the applications. Access to the applications in the secondary datacenter would then be filtered to only users who were a member of the newly created group.

[Click here to provide feedback](#)

Monitor Quick Access Links

Monitor Overview.....	186
Support.....	186
Operations.....	194
Monitoring.....	203

Monitor Overview

During the Monitor phase of the Citrix Consulting methodology, the Citrix solution is integrated into an organization's existing support structure, testing processes, monitoring system and operational processes.

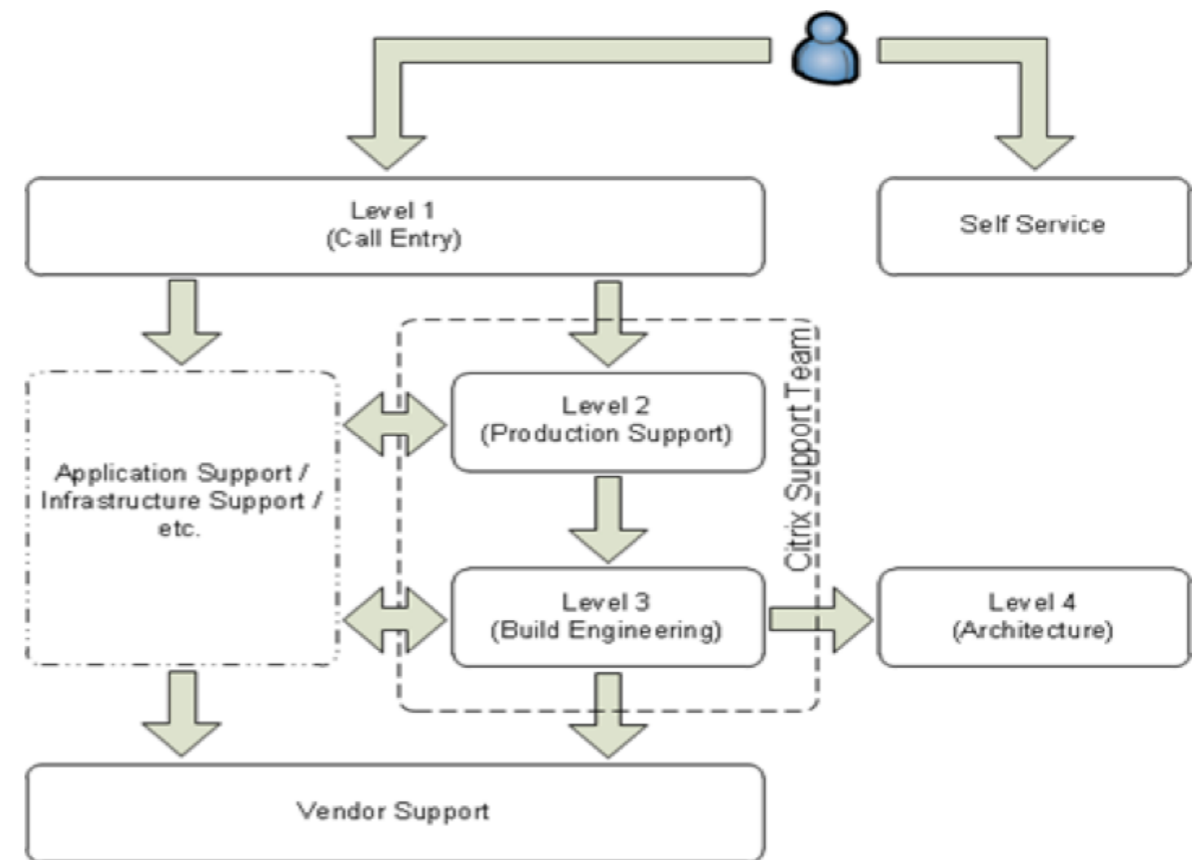
Support

When problems arise, technical support is the first point of contact. This section addresses the proper staffing, organization, training, delegated administration and tools that should be used to maintain the Citrix deployment.

Decision:Support Structure

Multiple support levels have been found to be the most effective means of addressing support issues. Low criticality, low complexity or frequently occurring issues should be managed and resolved at the lower support levels. High criticality and complex issues are escalated to more experienced architects or infrastructure owners. The diagram below outlines a common multi-level support structure.

Support Structure



If a user encounters an issue, Level-1 support (helpdesk) is the entry point to the support system. Level-1 should resolve 75% of all issues encountered, of which a majority will be routine problems that only require a limited knowledge of the Citrix environment. At this level, issues are quickly resolved and some may be automated (self-service), for example password resets and resource provisioning.

Non-routine problems that exceed Level-1's abilities are escalated to Level-2 (Operators). This support level is generally comprised of the administrators supporting the production Citrix environment. Information on the end user's problem and attempted troubleshooting steps are documented at the first level allowing Level-2 technicians to immediately begin addressing the problem. Level-2 technicians should handle only 20% of the support tickets and are highly knowledgeable on the Citrix environment.

Complex issues that exceed Level-2's abilities should be escalated

to Level-3 (Implementers). Level-2 and Level-3 support may often both be members of the Citrix Support Team, with Level-3 comprising the senior staff maintaining the Citrix environment. Level-3 issues are complicated and often mission critical requiring expert knowledge of the virtual desktop and application environment. Level-3 support tickets should amount to no more than 5% of all support issues.

The final level, Level-4 (Architects), is focused on the strategic improvements for the solution, testing new technologies, planning migrations, and other high level changes. Generally, Level-4 is not involved in active support of a production environment.

Should support discover an issue is related to an application or underlying infrastructure, the ticket is handed to the appropriate team for troubleshooting. If a program bug is discovered, the issue is then re-escalated and a ticket is established with the appropriate vendor.

Decision:Support Responsibilities and Skill Set

The following table highlights the recommended characteristics or each support level.

Support Levels

Support Level	Description	Responsibilities	Skill Set
Level-1 (Help Desk)	Provide first-line support of reported issues. Initially, servicing support messages and phone calls. This level needs to perform initial issue analysis, problem definition, ticket routing, and simple issue resolution. Often processes requests for application access or support with configuring plugins.	<ul style="list-style-type: none"> Perform issue definition, initial analysis and basic issue resolution Perform initial troubleshooting to determine the nature of the issue Create ticket, collect user information, and log all troubleshooting steps performed Resolve basic Citrix related issues, connectivity problems and application related issues using existing knowledge base articles Escalate issue to Level-2 if it advanced skills or elevated permissions are required If the issue is related to particular applications or other technologies and not the Citrix infrastructure, escalate the issue to the appropriate application or technology support units If the issue is deemed important enough, escalate directly to Level-3 Generate requests for additional issue resolution guides as necessary Follow up with end users when a support ticket is closed to ensure the problem has been satisfactorily resolved 	<ul style="list-style-type: none"> General Citrix XenApp/XenDesktop Knowledge General Windows Knowledge

Support Level	Description	Responsibilities	Skill Set
Level-2 (Operations)	<p>Primarily supporting day-to-day operations of the Citrix environment; may include proactive monitoring and management. In addition, this role should also perform intermediate level troubleshooting and utilize available monitoring or troubleshooting tools. Assist with resolving issues escalated by Level-1 Support.</p>	<ul style="list-style-type: none"> • Perform intermediate issue analysis and resolution • Identify root cause of issues • Respond to server alerts and system outages • Create weekly report on number of issues, close rate, open issues, etc. • Review vendor knowledge base articles • Respond to out-of-hours helpdesk calls • Respond to critical monitoring alerts • Generate internal knowledge base articles and issue resolution scripts and maintain Level-1 troubleshooting workflows • Perform basic server maintenance and operational procedures • Manage user profiles and data • Escalate ticket to Level-3 or appropriate technology owner if advanced skills or elevated permissions are required • Generate requests for additional issue resolution scripts and knowledge base articles as necessary 	<ul style="list-style-type: none"> • Experience with Microsoft Windows Server including but not limited to: <ul style="list-style-type: none"> • Configuring operating system options • Understanding Remote Desktop Services policies and profiles • Using Active Directory • Creating users/managing permissions and administrator rights • Creating and modifying Active Directory group policies • Basic administration skills, including: <ul style="list-style-type: none"> • An understanding of protocols (TCP) • An understanding of firewall concepts • An understanding of email administration and account creation • An understanding of Remote Desktop Services policies and profiles • The ability to create shares and give access to shared folders/files • Experience performing the following: <ul style="list-style-type: none"> • Managing, maintaining, monitoring and troubleshooting Citrix solutions • Backing up components in Citrix environments • Updating components in Citrix environments • Creating reports for trend analysis
Level-3 (Implementer)	<p>Central point for implementing, administering and maintaining Citrix desktop and application virtualization infrastructure. This role focuses on deploying new use cases and leading lifecycle management initiatives. Generally, one Implementer could focus on one use-case at a time. For example, three new concurrent use cases would require three Implementers. Escalates issues to software vendor specific Technical Support and notifies Level-4 about this issue.</p>	<ul style="list-style-type: none"> • Perform advanced issue analysis and resolution • Perform maintenance and environment upgrades • Addresses high severity issues and service outages • Manage the Citrix environment • Oversee and lead administrative tasks performed by Level-2 • Manage network and storage infrastructure as it relates to the Citrix environment (depending on size of company or Citrix environment). • Review periodic reports of server health, resource usage, user experience, and overall environment performance • Review vendor knowledge base articles and newly released updates • Perform policy-level changes and make Active Directory updates • Review change control requests that impact the Citrix environment • Perform advanced server and infrastructure maintenance • Review knowledge base articles and issue resolution scripts for accuracy, compliance, and feasibility • Create knowledge base articles and issue resolution scripts to address Level-2 requests 	<ul style="list-style-type: none"> • Knowledge of how the following Windows components integrate with Citrix technologies: <ul style="list-style-type: none"> • Active Directory Domain Services • Active Directory Certificate Services • Policies • Domain Name System (DNS) • Dynamic Host Configuration Protocol (DHCP) • Group Policy Objects (GPOs) • NTFS Permissions • Authentication and Authorization • Knowledge of IIS • Microsoft Windows Operating Systems <ul style="list-style-type: none"> • Windows 8.1 • Windows 7 • Windows Server 2012 R2 • Windows Server 2008 R2 • Roles and features of Windows Server • Knowledge of SQL 2008 R2 and newer • Knowledge of SQL clustering and mirroring • General networking skills (i.e. routing, switching) • Knowledge of hypervisors • Knowledge of shared storage configuration and management

Support Level	Description	Responsibilities	Skill Set
Level-4 (Architect)	The Level-4 team has minimal exposure to administrative tasks but focuses on scoping, planning and executing Citrix-specific service and project requests. An architect translates business requirements into a technical design.	<ul style="list-style-type: none"> • Provide technical leadership for upcoming projects • Lead design updates and architecture revisions • Address high severity issues and service outages • Oversee technology integration workflows • Review periodic reports of server health, resource usage, user experience, and overall environment performance to determine next steps and upgrade paths • Initiate load testing to determine capacity of environment • Review frequently recurring helpdesk issues • Ensure technical specifications continue to meet business needs • Update design documentation 	<ul style="list-style-type: none"> • Advanced architectural assessment and design skills for: <ul style="list-style-type: none"> • Citrix XenApp • Citrix XenDesktop • Citrix XenServer / VMware vSphere / Microsoft Hyper-V • Citrix Provisioning Services • Citrix NetScaler • Citrix StoreFront • Active Directory • Storage solutions • Networking • Application delivery • Disaster recovery • Policies/policy structures and security restrictions • Licensing • Methodology • Intermediate knowledge of: <ul style="list-style-type: none"> • General networking skills • Change control process • Project management • Risk assessment
Vendor Support	Vendor assistance may be necessary should defects in a program be discovered. At this stage, Level-3 engineers need to establish a support tickets with the appropriate vendor to assist with finding a solution.		
Self Service	A self-service portal should be utilized for noncritical tasks such as application access, permissions, password resets, etc. The portal can range from a simple FAQ page to a fully automated process requiring no human interaction. The purpose of the self-service portal is to add an additional touch point for end users to address basic issues, preventing the creation of new support tickets.		

Decision: Certifications and Training

The following table details the recommended training, certifications and experience for each support level.

Training Recommendations

Role	Recommended Training	Recommended Course(s)	Recommended Certification	Relevant Experience
Help Desk (Level-1)	<p>Level-1 support personnel should be provided with basic training on Citrix XenApp, Citrix XenDesktop and supporting technologies. This can include internal training from subject matter experts or from a Citrix Authorized Learning Center. The training provided should focus on the following topics:</p> <ul style="list-style-type: none"> • High-Level overview of the XenApp and XenDesktop implementation • Using Citrix Director to manage user sessions • Troubleshooting Citrix XenApp and XenDesktop sessions • Troubleshooting methodology <p>In addition, regular training should be provided to the Tier-1 team members on the latest troubleshooting recommendations from the Level-2 and Level-3 teams as well as details on any relevant changes to the environment. This will help to ensure a good baseline knowledge level across the team and consistent customer service.</p>	CXD-104: Citrix XenDesktop 7 Help Desk Support	N/A	1+ years (Entry level also acceptable)
Operator (Level-2)	<p>Level-2 personnel should conduct regular team training sessions to refine administrative skills and ensure a baseline knowledge level across the team.</p> <p>Formalized trainings are also essential when there are architectural updates to the environment and the Level-2 team is working with unfamiliar technologies. All members of the Level-2 team should achieve the Citrix Certified Associate (CCA) certification for Citrix Desktops and Apps. Advanced training on Windows concepts will also be essential for Level-2 team members who do not have desktop or server support experience.</p> <p>Finally, on-the-job training along with close integration with Level-3 administrators is essential as the Level-2 roles are formalized and responsibilities are handed over from Level-3 to Level-2.</p>	CXD-203: Managing App and Desktop Solutions with Citrix XenDesktop 7	Citrix Certified Associate - Virtualization	2-3 years
Implementer (Level-3)	<p>Level-3 support team members hold a minimum of three years of enterprise experience implementing and supporting XenApp, XenDesktop, Provisioning Services and Windows operating systems.</p> <p>Level-3 staff should also complete the Citrix Certified Professional (CCP) certification track as this will prepare them to proactively manage the user community and implement Citrix solutions according to Citrix best practices.</p>	CXD-300: Deploying App and Desktop Solutions with Citrix XenApp and XenDesktop 7.5	Citrix Certified Professional - Virtualization	3-4 years
Architect (Level-4)	<p>Experience is essential for Level-4 staff. A qualified Level-4 resource should have a minimum of five to six years of experience implementing, supporting, and serving in a technology architect role for a XenApp and/or XenDesktop environment as well as additional administrative experience with integrated technologies such as application and profile management solutions. The ideal candidate will have served in such a capacity at two or more environments for purposes of product exposure and in at least one environment of over 1,200 concurrent users. A Citrix Certified Expert (CCE) certification or comparable training and experience should be a prerequisite of the role.</p>	CXD-400: Designing App and Desktop Solutions with Citrix XenDesktop 7	Citrix Certified Expert - Virtualization	5+ years

Decision: Support Staffing

The following table provides guidance on the recommended number of support staff.

Staffing Recommendations

Role	Small Environment Sites: 1 Users: <500 Images: 1-2	Mid-size Environment Sites: 1-2 Users: 1000-5000 Images: 3-5	Large Environment Sites: 2+ Users: >5000 Images: 5+
Help Desk (Level -1)	3	5-10	15-20
Operator (Level -2)	1-2	2-3	4-5
Implementer (Level-3)	1	1-2	2-3
Architect (Level-4)	1	1	1-2

Note: This table should only be used as a baseline. Support staffing decisions should be evaluated against the defined requirements, projected workloads, and operational procedures of an organization. Multiple levels can be combined, for example there may be insufficient design projects to have a dedicated architect role, a more senior member of the Citrix team can act as an Operator and Implementer, or help desk members may support more technologies than just Citrix.

Decision: Job Aids

General Support Tools

The following table details tools that should be made available to all support levels.

Recommended General Helpdesk Tools

Tools	Details
Ticket Management System	Used to document customer information and issues. A typical ticket management system provides the following functionality: <ul style="list-style-type: none"> Monitoring the queue of tickets Setting a limit of the number of open tickets Establishing thresholds such as how long a certain type of ticket should take to be answered Identifying a group of users or individuals who require high priority assistance Informing the user when their ticket is open, updated, or closed
Call Scripts	The first contact help desk personnel should have documented scripts to ensure that all relevant data is captured while the user is on the phone. This practice also assists in proper triage and allows the next support level to perform research prior to customer contact. A sample call script is provided for reference.
Shadowing Tool	Shadowing is a useful tool when troubleshooting user issues. Support technicians and administrators can remotely observe a user's actions.
Knowledge Base	Documentation should be created and maintained in a knowledge base or library of known issues. Articles should be searchable for quick recovery. Knowledge bases help support staff to quickly resolve known issues and reduce the need to perform time consuming research.

Citrix Support Tools

The following table provides recommendations on the Citrix support tools that should be made available to each support level.

Support Citrix Tool Assignment

Tools	Details	Products				Support Level			
		XD	XA	PVS	XS	L1	L2	L3	L4
Citrix Director	Citrix Director provides an overview of hosted desktops and application sessions. It enables support teams to perform basic maintenance tasks and to monitor and troubleshoot system issues.	X	X			X	X	X	X
Citrix Studio	Citrix Studio enables administrators to perform configuration as well as maintenance tasks for a XenDesktop/XenApp site and associated virtual desktops or hosted applications.	X	X				X	X	X
Citrix Insight Services	Run from a single Citrix Delivery Controller to capture key data points and CDF traces for selected computers followed by a secure and reliable upload of the data package to Citrix Technical Support for escalation	X	X	X				X	X
HDX Monitor	HDX Monitor is a tool to validate the operation of the Citrix ICA/HDX stack of a user session. HDX Monitor provides information about client capabilities, network performance/activity, session settings and many more items.	X	X			X	X	X	X
HDX Insight	HDX Insight is a component of NetScaler Insight Center. It captures data about the ICA traffic that flows between the clients and the servers, generates records by doing deep inspection of the data, and presents the records as visual reports.	X	X				X	X	X
Provisioning Services Console	The Provisioning Services Console enables administrators to perform configuration and maintenance tasks for a Provisioning Services farm.			X			X	X	X
XenCenter	XenCenter enables administrators to perform configuration and maintenance tasks for a XenServer Resource Pool.				X			X	X

Citrix Insight Services

Administrators can utilize [Citrix Insight Services](#) to simplify the support and troubleshooting of the Citrix environment. Citrix Insight Services is run locally to collect environment information. Online analysis capabilities analyze that information and provide administrators recommendations based on their Citrix environment

[Click here to provide feedback](#)

and configuration. Additional information regarding Citrix Insight Services can be referenced in the Citrix Support article [CTX131233 – FAQ: Citrix Insight Services](#).

A full list of the available tools provided by Citrix Support to assist with troubleshooting can be referenced in Citrix Support article [CTX126294 – Complete List of Citrix Support Troubleshooting Tools](#).

Call Script

The following call script can be used as an initial baseline for a Citrix Help Desk team. Citrix Consulting recommends reviewing this sample call guide and adding any environment specific details that may need to be collected.

Step	Details
1.	What is the name and location of the user? This question will identify if the user is accessing the environment from an external or internal network location.
2.	Do any other users at the site/location experience the same issue? Can they have a colleague logon from same and/or different workstation? These questions help to determine if this is a workstation issue or a user issue.
3.	What type of endpoint device is the user utilizing? (Corporate device, BYOD, thin client, pc, laptop, etc.) This question will help determine if the issue is related to the user's endpoint.
4.	What is the Citrix Receiver version and connection information ? This question will verify if the user is using the right version of Receiver (the latest Receiver version or the version standardized by the company).
5.	Can the user see the StoreFront authentication page? This question helps to identify network issues.
6.	What is the name of the application (or virtual desktop) the user is attempting to use? Does the user see the appropriate application or desktop icon on the StoreFront site? These questions help to determine if there is an issue with user access and/or group membership.
7.	Does the application (or desktop) launch when the icon is selected? Does the application logon screen appear (if applicable)? These questions help to determine if a connection is made into the Citrix XenDesktop infrastructure.
8.	Can the user authenticate into the application (if applicable)? Does the issue occur after the application is launched? This question helps to determine if the issue is with the application rather than the application delivery infrastructure.
9.	What is the specific error seen (if applicable)? This question identifies the specific error. The user should be requested to provide a screenshot, if available.

Decision: Delegated Administration

Each support level must be provided with sufficient rights to effectively perform their role. The following tables provide guidance on the recommended privileges per support level.

XenDesktop/XenApp Delegated Rights

Admin Role	Support Level
Help Desk Administrator	Level-1
Delivery Group Administrator	Level-2
Full Administrator	Level-3
Read-Only Administrator	Level-4

For further information about delegated rights within a XenDesktop/XenApp Site, please refer to Citrix eDocs – [XenApp and XenDesktop Delegated Administration](#).

Provisioning Services Delegated Rights

Admin Role	Support Level
N/A	Level-1
Site Administrator	Level-2
Farm Administrator	Level-3
N/A	Level-4

For further information about delegated rights within a Provisioning Services Site, please refer to Citrix eDocs – [Provisioning Services Managing Administrative Roles](#).

StoreFront Delegated Rights

Admin Role	Support Level
N/A	Level-1
N/A	Level-2
Local Administrator on StoreFront Server	Level-3
N/A	Level-4

Users with local administrator rights have access to view and manage all objects within StoreFront or Web Interface. These users can create new sites and modify existing ones.

Citrix License Server Delegated Rights

Admin Role	Support Level
N/A	Level-1
N/A	Level-2
Administrator	Level-3
N/A	Level-4

By default, the account used during the installation of the license server becomes the administrator for the console. Often the accounts used for the installation are not the intended accounts for the regular administration tasks. For the steps to change the default administrator please reference CTX135841 – [How to Change the Default Administrator for the Citrix Licensing Server Version 11.10](#). All users created through this process are full administrators of the Citrix License Server.

XenServer Delegated Rights

Admin Role	Support Level
N/A	Level-1
Virtual Machine Operator	Level-2
Pool Administrator	Level-3
Read-Only	Level-4

For further information about delegated rights within a XenServer Resource Pool, please refer to CTX137828 – [XenServer 6.2 Administrators Guide](#) (see chapter Role Based Access Control).

Daily Operations

Component	Tasks	Description	Responsible
Generic	Review Citrix Director, Windows Performance Monitor, Event Log, and other monitoring software alerts.	<p>Check for warnings or alerts within Citrix Director, event logs, or other monitoring software. Investigate the root cause of the alert if any.</p> <p>Note: A computer and monitor can be set up to display the Citrix Director dashboard to create a Heads up Display for the Citrix department. This ensures the status of the environment is clearly visible.</p> <p>Monitoring recommendations for XenDesktop and XenApp 7.x are included in the Monitoring section of the Virtual Desktop Handbook.</p>	Operators
Generic	Verify backups completed successfully	<p>Verify all scheduled backups have been completed successfully. This can include but is not limited to:</p> <ul style="list-style-type: none"> • User data (user profiles / home folders) • Application data • Citrix databases • StoreFront configuration • Web Interface configuration • Provisioning Services vDisks (virtual desktops and application servers) • XenServer VM/Pool metadata (or equivalent for other hypervisors) • Dedicated virtual desktops • License files 	Operators
Generic	Test environment access	<p>Simulate a connection both internally and externally to ensure desktop and application resources are available before most users log on for the day. This should be tested throughout the day and may even be automated.</p>	Operators

Operations

This section defines routine operations for the Citrix environment that help to improve stability and performance.

Decision: Administrative Tasks

The Citrix Support Team should perform regular operations and maintenance tasks to ensure a stable, scalable Citrix environment.

Each operations is categorized by the associated component of the solution as well as the frequency of the operation (ongoing, daily, weekly, and yearly). Tasks have been aligned to the roles described within [Decision: Support Responsibilities and Skill Set](#).

Daily Periodic Tasks

The following table outlines the tasks that should be performed by the Citrix Support Team on a daily basis.

Component	Tasks	Description	Responsible
XenApp/XenDesktop	Virtual machine power checking	Verify that the appropriate number of idle desktops and application servers are powered on and registered with the Delivery Controllers to ensure availability for user workloads.	Operators
XenApp/XenDesktop	Perform incremental backup of Citrix related databases	Perform incremental-data backups of the following Citrix databases: <ul style="list-style-type: none"> • Site Database • Configuration Logging Database • Monitoring Database 	Operators, Database team (if Citrix environment is using a shared SQL)
Provisioning Services	Check Citrix Provisioning Server utilization	Check the number of target devices connected to the Citrix Provisioning Servers and balance the load across servers, if required.	Operators
Provisioning Services	Perform incremental backup of Citrix PVS database	Incremental backup of Citrix Provisioning Server database hosted on SQL Server infrastructure.	Operators, Database team (if Citrix environment is using a shared SQL)

Weekly Periodic Tasks

The following table outlines the tasks that should be performed by the Citrix Support Team on a weekly basis.

Weekly Operations

Component	Tasks	Description	Responsible
Generic	Review latest hotfixes and patches	Review, test, and deploy the latest Citrix hotfixes and ascertain whether the Delivery Controllers and Server-Based OS / Desktop-Based OS virtual machines require them. Note: Any required hotfixes should be tested using the recommended testing process prior to implementation in production.	Operators, Implementers (review process)
Generic	Create Citrix environment status report	Create report on overall environment performance (server health, resource usage, user experience) and number of Citrix issues (close rate, open issues, and so on).	Operators
Generic	Review status report	Review Citrix status report to identify any trends or common issues.	Implementers, Architect
Generic	Maintain internal support knowledge base	Create knowledge base articles and issue resolution scripts to address Level-1 and Level-2 support requests. Review knowledge base articles and issue resolution scripts for accuracy, compliance, and feasibility.	Operators (Level-1 requests), Implementers (Level-2 requests, and review process)
XenApp/ XenDesktop	Check Configuration Logging reports	Confirm that Citrix site-wide changes implemented during the previous week were approved through change control.	Implementers
XenApp/ XenDesktop	Perform full backup of Citrix related databases	Perform full-data backups of the following Citrix databases: <ul style="list-style-type: none"> • Site Database • Configuration Logging Database • Monitoring Database 	Operators, Database team (if Citrix environment is using a shared SQL)

Component	Tasks	Description	Responsible
Provisioning Services	Check storage capacity (only prior to updating a vDisk)	Review storage utilization, used and free storage space, for vDisk store and each vDisk. Note: Lack of space within the vDisk repository will be an issue only when the vDisks are updated using versioning or when a vDisk is placed in private mode during an update procedure. Storage utilization within vDisk should also be investigated. For example a 20GB vDisk may only have 200MB of free storage. If the vDisk itself is limited for storage then it needs to be extended. Citrix does not support resizing of a VHD file. Refer to the Microsoft link Resize-VHD for information on resizing a VHD file.	Operators
Provisioning Services	Perform vDisk updates (as necessary)	Perform a full backup of the vDisk before implementing any updates. Update the master vDisk image files and apply the following: <ul style="list-style-type: none"> Windows software updates and patches Operating system and application changes Anti-virus pattern and definitions updates <p>Note: Updates should be tested using the recommended testing process prior to implementation in production.</p>	Operators
Provisioning Services	Check auditing reports	Review the Citrix Provisioning Services auditing Logs. Note: Provisioning Server auditing is off by default and can be enabled to record configuration actions on components within the Provisioning Services farm. To enable auditing refer to the Citrix eDocs article Enabling Auditing Information .	Implementers
Provisioning Services	Perform full backup of Citrix PVS database	Backup of Citrix Provisioning Server database hosted on SQL Server infrastructure.	Operators

Monthly Periodic Tasks

The following table outlines the tasks that should be performed by the Citrix Support Team on a monthly basis.

Monthly Operations

Component	Tasks	Description	Responsible
Generic	Perform capacity assessment	Perform capacity assessment of the Citrix environment to determine environment utilization and any scalability requirements. Note: Recommendations for performing a capacity assessment are included in Decision: Capacity Management in the Monitoring section of the Virtual Desktop Handbook.	Architect

Yearly Periodic Tasks

The following table outlines the tasks that should be performed by the Citrix Support Team on a yearly basis.

Yearly Operations

Component	Tasks	Description	Responsible
Generic	Conduct Citrix policy assessment	Review Citrix policies and determine whether new policies are required and existing policies need to be updated.	Implementers
Generic	Review software upgrades	Review and assess the requirement for new Citrix software releases or versions.	Architect
Generic	Perform Business Continuity Plan (BCP)/ Disaster Recovery (DR) test	Conduct functional BCP/DR test to confirm DR readiness. This plan should include a yearly restore test to validate the actual restore process from backup data is functioning correctly.	Architect
Generic	Perform application assessment	Review the usage of applications outside and within the Citrix environment. Assess the validity of adding additional applications to the Citrix site, removing applications that are no longer required, or upgrading the applications to the latest version.	Architect
Generic	Archive audit reports	Perform an archive of the Citrix Provisioning Server Audit Trail Information for compliance requirements.	Implementers

Decision: Backup Location

The location of backups directly effects the recovery time and reliability of the Citrix environment. It is recommended to store backups of critical data both onsite and at an offsite location. If offsite backups are not possible due to costs associated or sensitivity of the data, backups should be placed at separate physical locations within the same datacenter.

Each backup option is discussed further below.

- **Onsite Backups** – Onsite backups should be located on a storage device in the datacenter that will allow the data to be recovered quickly in the event of a failure. Onsite backups are ideal for issues that only affect a small subnet of hardware in the datacenter. Backups can also be stored on a cold storage solution such as tape. While this medium is slower to recover from, it provides additional protection since it is only active during the backup process.
- **Offsite Backups** – Although the time to recover is much higher,


offsite backups provide additional protection in the event of a disaster. Offsite backups may require transferring data over the Internet to a third party provider or they are created onsite and then transported to a remote location on storage mediums such as tape. It is typical to put a limited number of backups offsite. For example one backup a week or month.

Decision: Testing Process

Regular updates and maintenance are an everyday part of IT operations. Standard processes must be followed to ensure updates do not negatively impact the production environment. This includes maintaining a dedicated testing infrastructure where modifications can be validated prior to being implemented in production.

Since changes to Citrix infrastructure can impact thousands of virtual desktop and application users, multi-phase testing is critical for the reliability and performance of the environment. As such, the process for testing should resemble the following:

Testing Process



- **Development** - The development infrastructure exists outside of the production network. Typically, it consists of short-lived virtual machines whose configuration matches production as closely as possible. The purpose of the development phase is to provide change requestors a non-production environment to perform proof of concepts, determine integration requirements and perform iterative testing as part of a discovery phase. Proposed changes should be documented so they can be applied in the test phase.
- **Test** - The test environment is a standalone 1:1 copy of the production infrastructure and is used to confirm that the proposed changes can be easily repeated prior to the pre-production staging environment. The changes made should follow documentation from the development stage. If testing fails within the testing stage, the architect must determine the severity of failure and determine whether minor updates to documentation is sufficient or a full development cycle is needed.
- **Pre-production** - The staging environment should mimic the current production environment. The goal of staging is to implement the proposed changes with little risk or uncertainty. It is expected that any changes made to the staging infrastructure have been tested and documented for repeatability. There should not be any iterations or adjustments required within this phase. During this phase and within this environment User Acceptance Testing (UAT) should be performed.
- **Production** - The production environment is a fully redundant

and scalable solution designed for normal usage by end users. There should be minimal changes to the environment. If possible, all approved changes should be rolled out in stages to the production environment. This process is known as a staged rollout and mitigates risk by allowing changes to be rolled back, if necessary, without impacting the entire environment.

Decision: Change Control

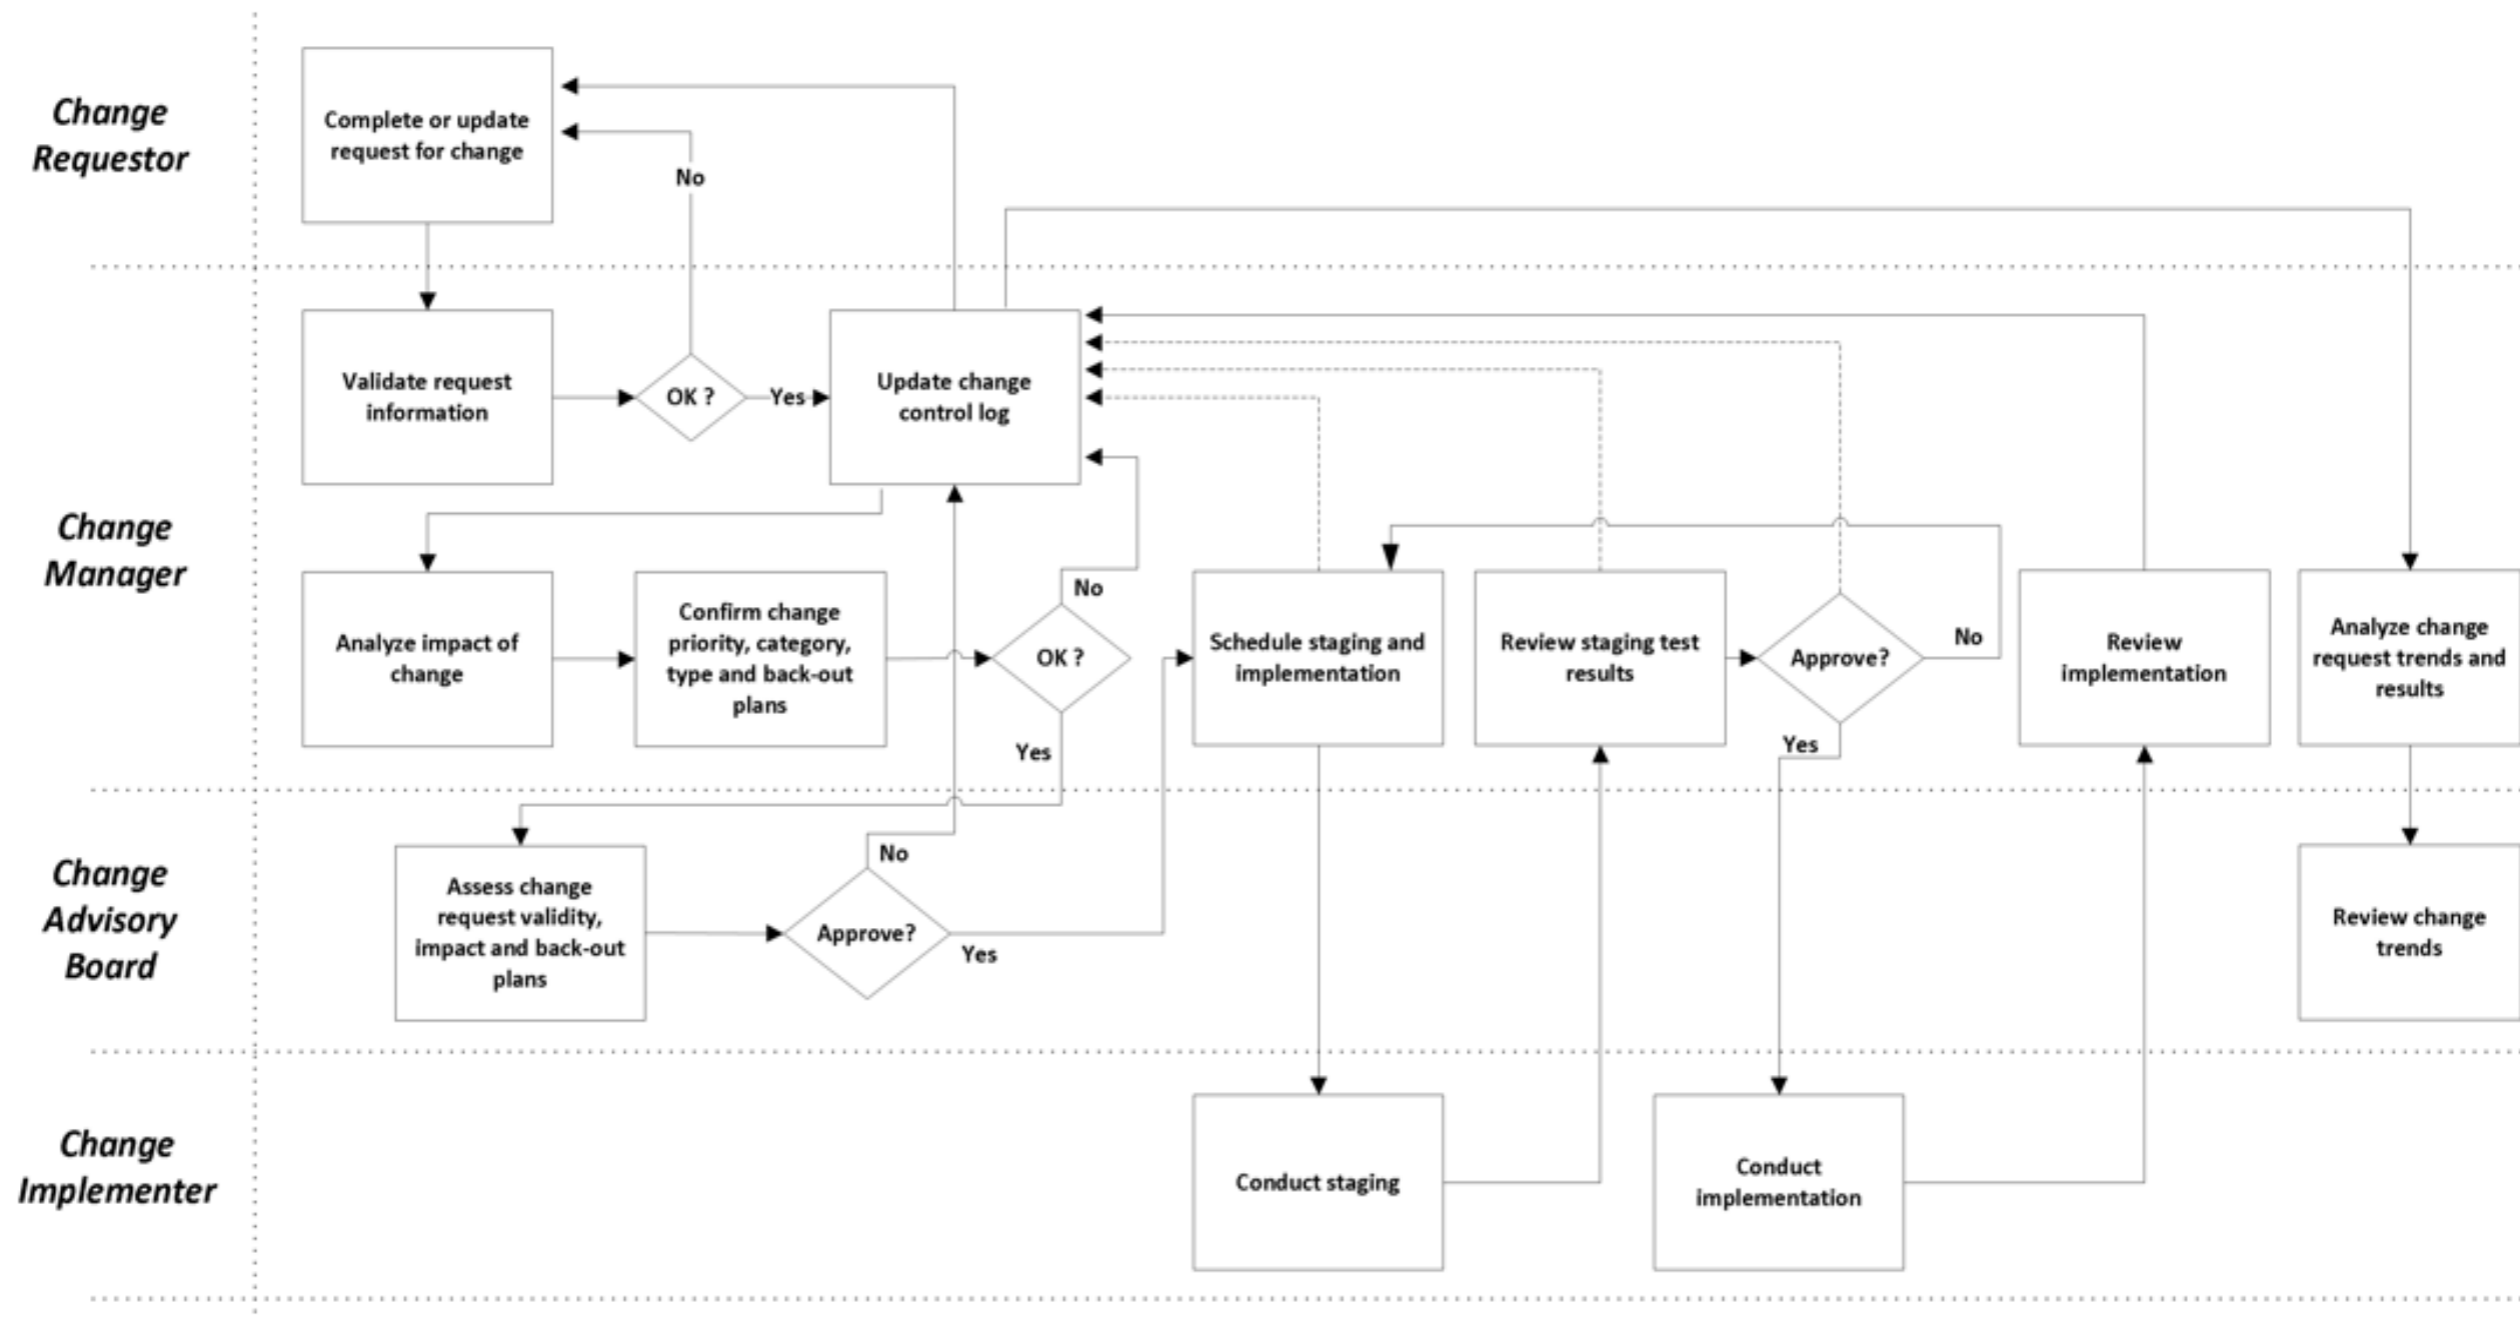
Standardized processes that manage changes throughout a system's lifecycle are necessary to ensure consistent and accountable performance. The following change control leading practices should be considered.

- Use a change control window so that all applicable parties know when there might be downtime.
- Make sure that all teams are represented in the Change Advisory Board (CAB).
- Every change should have a roll back plan.
- If a change fails have a "hot wash" to determine what went wrong.
- Always use an automated change control system so that support staff can quickly and easily identify changes.
- When available, ensure configuration logging is enabled to track any changes made to the Citrix environment.

The change control process should be closely followed starting with a change request. A change request form should be filled out detailing changes requested, reasons for the change, and intended timeframes for the action. This is then reviewed and edited if required by a Change Manager and advisory board. When the change request has gone through the entire change approval process it is given to a change implementer who stages the change for testing, and finally conducts the implementation in production.

A sample change control process, including detailed steps, is provided in the diagram below:

Change Control Process



The process is as follows:

1. The Request for Change (RfC) form is completed by any person requesting a change.
2. After appropriate manager approvals have been acquired, the RfC is forwarded to the appropriate Change Manager(s).
3. The Change Manager validates the RfC for completeness and logs the RfC information into the Change Control Log for tracking. Incomplete change requests are returned to the requestor for update and re-submission.
4. The Change Manager assesses the impact of the change in conjunction with subject matter experts and/or managers of the teams associated/affected by this change.
5. The Change Manager works with the associated/affected teams as well as the change requestor in order to confirm the priority, category and type of the change as well as the proposed rollback plan.
6. If the change is approved by the Change Manager, the RfC is forwarded to the CAB for approval. If the change is rejected, the Change Control Log is updated with the current status as well as the reason of the rejection and the RfC is send back to the requestor.
7. The CAB reviews and validates the change in detail, and discusses and evaluates purpose, reasons, impact, cost and benefits. Each board member represents their department and provides guidance on the change requests. The CAB also reviews multiple requests to coordinate implementations and “package” requests into a single release schedule.
8. Upon approval the change is sent back to the Change Manager to schedule the change for implementation into the staging environment.
9. The change is implemented and tests are conducted. The results are sent back to the Change Manager.

10. If the staging implementation and testing are successful, the change is scheduled for production implementation. In case the staging phase was not successful another staging iteration will be conducted.
11. If possible, the change is rolled out in stages to the production environment. This process is known as a staged rollout and mitigates risk by allowing changes to be rolled back, if necessary, without impacting the entire environment. A rollback plan should be in place if there is an issue implementing a change in the production environment.
12. The Change Manager reviews the implementation and finally updates the Change Control Log.
13. On a periodic basis, the Change Manager reviews the Change Control Log to identify trends on type, frequency and size of changes and forwards the results to the CAB for review.

In an emergency, the processes may be expedited. Should an issue be declared an emergency, a change request form is still filled out and delivered to the appropriate change management representative. When approved, the requested change is immediately implemented and the advisory board notified.

Decision: Availability Testing

Availability testing is focused on ensuring resources are still available in the instance of a component failure. These tests are essential to ensuring users always have access to business critical resources. The testing should be conducted during nonbusiness hours or during a scheduled maintenance weekend when appropriate notice has been given to end users to make them aware if any unforeseen issues arise.

The following is a list of the key components that should be tested on a regular basis.

- **StoreFront** – StoreFront should be load balanced and health checked by a NetScaler or other load balancing device. To

validate its configuration, all but one of the StoreFront servers should be shutdown. This will validate that the load balancing device is detecting the failure and directing users to the functioning server.

- **SQL** – SQL Server should be in a high availability configuration. To validate the configuration, the primary SQL server should be taken offline and then the Citrix Studio console should be opened. Since Citrix Studio will not be accessible without a functioning SQL server, it will validate that the SQL server failover mechanisms are functioning properly.
- **Delivery Controllers** - Resources deployed should be configured with a list of multiple Delivery Controllers. If one is made unavailable, desktops and application hosts will automatically establish a connection to another server in the list. To validate this, shutdown one of the Delivery Controller hosts and determine if the resources initially connected to it automatically register to another server. This can be determined by viewing the registration status of the resources inside Citrix Studio.

Sample Availability Testing Workflows

The following availability testing workflows can be used as a starting point for integrating Citrix availability testing into standard operational procedures. A successful availability test is defined as:

Citrix Provisioning Services

Prerequisites and configuration requirements:

- Hypervisor, XenApp, and XenDesktop services are up and running.
- At least two PVS servers are installed and configured, providing the streamed disk image.
- Resilient networking and storage infrastructure with multiple links to each server.

- Test users are active on the XenApp or XenDesktop machines.

Steps	Expected Results
PVS Server Outage <ul style="list-style-type: none"> • Shutdown one of the Provisioning Servers • Validate PVS continues to function • Restart PVS Server • Validate connections rebalance between PVS Servers 	<ul style="list-style-type: none"> • Existing XenApp/XenDesktop machines connect to another PVS server. • There is limited to no impact to the users utilizing that server. • New XenApp/XenDesktop machines can be booted and start correctly. • SCOM reports that the PVS server is down / not available. • Live connections are rebalanced between both PVS servers once both PVS servers are made available again.
PVS Bond Disruption <ul style="list-style-type: none"> • Disable/unplug a NIC in the PVS Streaming Bond on the PVS server 	<ul style="list-style-type: none"> • Provisioning Server continues to stream over remaining NICS in PVS Streaming Bond.
SQL Server PVS Database Mirror Failover <ul style="list-style-type: none"> • Admin logs on to Principle SQL Server. • Initiate failover of PVS database. • Validate PVS continues to function. • Initiate failback of PVS database. • Validate PVS continues to function. 	<ul style="list-style-type: none"> • PVS continues to function.
SQL Service Outage <ul style="list-style-type: none"> • Admin reboots both Principle & Mirror SQL Servers simultaneously. • Validate PVS continues to function, but that administration is not possible. • Wait for the SQL Server to come back online. • Validate PVS administrative functions are once again possible. 	<ul style="list-style-type: none"> • PVS continues to function. • PVS administrative functions are no longer available. • PVS administrative functions are available once the SQL services are restored.

Citrix XenDesktop and XenApp Services

Prerequisites and configuration requirements:

- Hypervisor, XenDesktop, and StoreFront services are up and running.
- Network and storage services available.
- Provisioning Services is providing the streamed disk images.
- Test users are active on the virtual machines.
- SQL (Mirroring) and XenDesktop servers are up and running.
- Ensure multiple StoreFront servers are running.
- NetScaler load balancing services.

Steps	Expected Results
<p>XenApp/XenDesktop 7.x Delivery Controller Citrix Broker Service Outage:</p> <ul style="list-style-type: none"> • Stop the Citrix Broker Service on one of the Delivery Controller servers. • Validate virtual desktops or applications can still be enumerated and launched. • Start the Citrix Broker Service on the Delivery Controller server. • Shutdown one of the Desktop Controllers. • Validate virtual desktops or applications can still be enumerated and launched. • With a desktop launched, determine which Controller owns the host connection. Shut the Controller down and verify that another Controller takes over the session. 	<ul style="list-style-type: none"> • StoreFront correctly identifies service as being unavailable and redirects connections to remaining Delivery Controller. • Desktops continue to be enumerated and launch successfully. • Launched desktop can be supported if a hosting Controller goes down.
<p>SQL Server Database Mirror Failover:</p> <ul style="list-style-type: none"> • Admin logs on to principle SQL Server. • Initiate failover of XenApp/XenDesktop database. • Validate XenApp/XenDesktop continues to function. 	<ul style="list-style-type: none"> • The database should failover and the Citrix Studio should pick up the failover database with no issues. • Existing sessions are not impacted. • New sessions are possible. • Administrative functions are possible.

Steps	Expected Results
<p>SQL Service Outage:</p> <ul style="list-style-type: none"> • Admin restarts both principle & mirror SQL Servers simultaneously. • Validate XenApp/XenDesktop continues to function, but that administration is not possible. • Wait for the SQL Service to come back online. • Validate administrative functions are once again possible. 	<ul style="list-style-type: none"> • Existing XenDesktop sessions are not impacted • Recently used applications, hosted shared desktops and assigned VDI can be accessed due to connection leasing. New sessions are not possible if they do not meet the criteria for connection leasing. <p>For more information on connection leasing, reference Citrix eDocs – Connection leasing.</p> <ul style="list-style-type: none"> • XenDesktop Administrative functions are not possible • XenDesktop Administrative functions are possible once SQL service is available.

Citrix Licensing Services

Prerequisites and configuration requirements:

- Citrix Licensing Server up and running (with valid licenses installed).
- Hypervisor, XenApp/XenDesktop and StoreFront services are up and running.
- Users are active on the Server OS or Desktop OS machines.

Steps	Expected Results
<p>Service continuity during complete failure of the Citrix Licensing Server:</p> <ul style="list-style-type: none"> • Shutdown the Citrix Licensing server. • Reboot an existing Server OS machine. • Logon to the Citrix StoreFront and launch a published application. • Reboot an existing Desktop OS machine. • Logon to the Citrix StoreFront and launch a virtual desktop. 	<ul style="list-style-type: none"> • License Server connectivity error posted in Event Log. • Provisioned Server OS boots successfully. • Users are able to launch published applications. • Provisioned Desktop OS boots successfully. • User is able to launch a virtual desktop. • Administrators will have 30 days grace to recover the Citrix Licensing Server.

Monitoring

By having an in-depth understanding of current and expected behavior of the Citrix environment and its components, administrators are better equipped to discover an issue before it impacts the user community. Furthermore the data tracked during normal operations can be used for trending and capacity planning. This section defines how a Citrix environment should be monitored, as well as some common tools that can be used.

Decision: Performance Monitor Metrics

Monitoring the performance of the overall environment is crucial towards making sure all components are available and performing effectively to ensure users have a high quality experience.

Different components within the overall solution require monitoring of unique metrics with appropriately set thresholds. The metrics and thresholds presented are based on real world experience but may not apply to all environments. Organizations will need to perform their own baselining, validity testing and validation before implementing within a production environment.

Note: Some hypervisors, such as VMware vSphere and Hyper-V, provide specific performance counters for tracking CPU and Memory utilization within virtual machines (i.e. "VM Processor \ % Processor Time"). These performance counters should be used in addition to the general counters listed below.

Generic

These performance counters should be used to monitor the key performance metrics of the Citrix infrastructure, application servers, and virtual desktops.

Recommended Metrics to Monitor for all Virtual Machines

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
Processor - % Processor Time	% Processor Time is the percentage of elapsed time that the processor spends to execute a non-Idle thread. It is calculated by measuring the duration of the idle thread is active in the sample interval, and subtracting that time from interval duration. (Each processor has an idle thread that consumes cycles when no other threads are ready to run). This counter is the primary indicator of processor activity, and displays the average percentage of busy time observed during the sample interval. It is calculated by monitoring the time that the service is inactive and subtracting that value from 100%.	80% for 15 minutes	95% for 15 minutes	<p>Identify the processes/services consuming processor time using Task Manager or Resource Monitor.</p> <p>If all processes/services work within normal parameters and the level of CPU consumption is an expected behavior it should be considered to add additional CPU resources to this system in the future.</p> <p>If a process/service can be identified which works outside normal parameters, the process should be killed. Please note that killing a process can cause unsaved data to be lost.</p>

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
System - Processor Queue Length	Processor queue length is the number of threads in the processor queue. Unlike the disk counters, this counter shows ready threads only, not threads that are running. There is a single queue for processor time even on computers with multiple processors. Therefore, if a computer has multiple processors, you need to divide this value by the number of processors servicing the workload. A sustained processor queue of less than ten threads per processor is normally acceptable, dependent of the workload.	5 (per core) for 5 minutes or 6 (per core) for 15 minutes	10 (per Core) for 10 minutes or 12 (per core) for 30 minutes	A long CPU queue is a clear symptom of a CPU bottleneck. Please follow the steps outlined for counter " Processor - % Processor Time".
Memory - Available Bytes	Available memory indicates the amount of memory that is left after nonpaged pool allocations, paged pool allocations, process' working sets, and the file system cache have all taken their piece.	<30% of total RAM or 20% of physical memory over 6 minutes	<15% of total RAM or 5% of physical memory over 6 minutes	Identify the processes/services consuming memory using Task Manager or Resource Monitor. If all processes/services work within normal parameters and the level of memory consumption is an expected behavior it should be considered to add additional memory to this system in the future. If a process/service can be identified which works outside normal parameters, the process should be killed. Please note that killing a process can cause unsaved data to be lost.
Paging File - %Usage	This is the percentage amount of the Page File instance in use.	>40% or 80% over 60 minutes	>70% or 95% over 60 minutes	Review this value in conjunction with " Memory - Available Bytes" and " Memory - Pages/sec" to understand paging activity on the affected system.
LogicalDisk/Physical Disk - % Free Space	% Free Space is the percentage of total usable space on the selected logical disk drive that is free.	<20% of physical disk or 20% reported after 2 minutes	<10% of physical disk or 15% reported after 1 minute	Identify which files or folders consume disk space and delete obsolete files if possible. In case no files can be deleted, consider increasing the size of the affected partition or add additional disks.
LogicalDisk/Physical Disk - % Disk Time	% Disk Time marks how busy the disk is.	>70% consistently or 90% over 15 minutes (_Total)	>90% consistently or 95% over 15 minutes (_Total)	Identify the processes / services consuming disk time using Task Manager or Resource Monitor. If all processes/services work within normal parameters and the level of disk consumption is an expected behavior it should be considered to move the affected partition to a more capable disk subsystem in the future. If a process/service can be identified which works outside normal parameters, the process should be killed. Please note that killing a process can cause unsaved data to be lost.

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
LogicalDisk/PhysicalDisk – Current Disk Queue Length	Current disk queue length provides a primary measure of disk congestion. It is an indication of the number of transactions that are waiting to be processed.	>=1 (per spindle) consistently or 3 over 15 minutes (_Total)	>=2 (per spindle) consistently or 10 over 30 minutes (_Total)	A long disk queue length typically indicated a disk performance bottleneck. This can be caused by either processes/services causing a high number of I/Os or a shortage of physical memory. Please follow the steps outlined for counter “ LogicalDisk/PhysicalDisk - % Disk Time” and counter “ Memory – Available Bytes”
LogicalDisk/PhysicalDisk – Avg. Disk Sec/Read – Avg. Disk Sec/Write – Avg. Disk Sec/Transfer	The Average Disk Second counters show the average time in seconds of a read/write/transfer from or to a disk.	>=15ms consistently	>=20ms consistently	High disk read or write latency indicates a disk performance bottleneck. Systems affected will become slow, unresponsive and application or services may fail. Please follow the steps outlined for counter “LogicalDisk/PhysicalDisk - % Disk Time”
Network Interface – Bytes Total/sec	Bytes Total/sec shows the rate at which the network adaptor is processing data bytes. This counter includes all application and file data, in addition to protocol information, such as packet headers.	< 8 MB/s for 100 Mbit/s adaptor <80 MB/s for 1000 Mbit/s adaptor or 60% of NIC speed inbound and outbound traffic for 1 min.	70% of NIC speed inbound and outbound traffic for 1 min.	Identify the processes / services consuming network using Task Manager or Resource Monitor. If all processes/services work within normal parameters and the level of bandwidth consumption is an expected behavior it should be considered to move the respective process/service to a dedicated NIC (or team of NICs). If a process/service can be identified which works outside normal parameters, the process should be killed. Please note that killing a process can cause unsaved data to be lost.

XenApp/XenDesktop

These performance counters are specific to the Delivery Controllers.

Recommended XenApp/XenDesktop Metrics

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
Database Avg. Transaction Time	The time on average, in seconds, taken to execute a database transaction. A baseline needs to be established in the environment in order to accurately establish threshold values.	Based on baseline values	Based on baseline values	In case the reported values exceed the baseline response time constantly, a potential performance issue needs to be investigated at the SQL server level.
Database Connected	Indicates whether this service is in contact with its database. (1 is connected; 0 is not connected).	0	0 (for over 30 minutes)	Both values report connectivity issues of the XenDesktop Broker service with the database. In case issues are reported, SQL server and network availability needs to be verified.
Database Transaction Errors/sec	The rate at which database transactions are failing.	None	>0	Both values report connectivity issues of the XenDesktop Broker service with the database. In case issues are reported, SQL server and network availability needs to be verified.

StoreFront

These performance counters are specific to the StoreFront servers.

Recommended StoreFront Metrics

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
ASP.NET – Request Queued	The number of requests waiting to be processed by ASP. A baseline needs to be established in the environment in order to accurately establish threshold values.	Based on baseline values	Based on baseline values	In case the queue length exceeds the critical limit requests may be rejected. In this case it should be considered to add additional StoreFront or Web Interface servers to the load balancing team in order to distribute the load across more nodes.
ASP.NET – Requests Rejected	The number of requests rejected because the request queue was full.	None	≥ 1	When this limit is exceeded, requests will be rejected with a 503 status code and the message “Server is too busy.” Please follow the steps outlined for counter “ASP.NET – Request Queued”

Citrix License Server

These performance counters are specific to the Citrix License Server.

Recommended Citrix License Server Metrics

Metric	Description	Warning (Yellow)	Critical (Red)	Troubleshooting/Remediation
Citrix Licensing – Last Recorded License Check-Out Response Time	Displays the last recorded license check-out response time in milliseconds.	> 2000 ms	> 5000 ms	If the reported values exceed the 5000 ms response time, a potential performance issue needs to be investigated in the Citrix License Server.
Citrix Licensing – License Server Connection Failure	Displays the number of minutes that XenDesktop has been disconnected from the License Server.	> 1 minute	> 1440 minutes	Both values report connectivity issues with the License Server. In case issues are reported, License Server and network availability needs to be verified.

Decision: Services Monitoring

Windows services that are critical to basic server functionality should be automatically monitored to ensure that they are running properly. The following table provides a list of the common Windows services that should be monitored. When any of these services are restarted or stopped a warning (Yellow) or critical (Red) alert should be assigned respectively. The recommended recovery actions for the services listed below are as follows:

- First failure: Restart the Service
- Second Failure: Restart the Service
- Subsequent Failures: Put the server in maintenance mode and investigate the root cause

XenApp/XenDesktop

XenApp/XenDesktop 7.x Services

Service	Functionality	Administration Risk
Citrix Diagnostic Facility COM Server Service	Manages and controls Citrix diagnostic trace sessions on the system. Dependencies: <ul style="list-style-type: none"> • RPC Service 	This service has no impact on the production environment. It is used to generate CDF trace files which aid in troubleshooting issues.
Citrix Environment Test Service	Manages tests for evaluating the state of a XenDesktop Site.	If this service is stopped administrators will be unable to establish new connections to Citrix Studio. Administrators will also be unable to check the status of the Citrix site configuration, machine catalogs, and delivery groups by running the tests under "Common Tasks" in the Citrix Studio administration console.
Citrix Host Services	Manages host and hypervisor connections. Dependencies: <ul style="list-style-type: none"> • WMI Service 	Administrators will be unable to create new Machine Catalogs or control virtual machine power settings via Citrix Studio. Administrators will be unable to establish new connections to Citrix Studio. Users may experience issues connecting to virtual desktops when this service is not available. If this service is stopped existing connections are not affected.
Citrix Machine Creation Service	Creates new virtual machines. Dependencies: <ul style="list-style-type: none"> • WMI Service 	Administrators will be unable to create new or modify existing Machine Catalogs or establish new connections to Citrix Studio. Administrators will be unable to establish new connections to Citrix Studio.
Citrix Monitor Service	Monitors the FlexCast system.	If this service is stopped XenApp/XenDesktop will be unable to communicate with the Monitoring Database. Citrix Director will be unable to retrieve any data on the environment. Administrators will be unable to establish new connections to Citrix Studio.
Citrix StoreFront Service	Manages deployment of StoreFront.	Administrators will be unable to establish new connections to Citrix Studio.

Delivery Controller Services Monitoring in Citrix Director

The **Infrastructure** pane within the Citrix Director dashboard provides status of the services running on the Delivery Controllers and will provide warning indications if a service or Controller is unavailable. These alerts can be accessed by clicking the **Alert** hyperlink within the **Infrastructure** pane.

Citrix Director Infrastructure Pane

Status	Services	Site Database	License Server	Configuration Logging Database	Monitoring Database
✓ Online	⚠ 1 Alert	✓ Connected	✓ Connected	✓ Connected	✓ Connected

Provisioning Services

Provisioning Server Services

Service	Functionality	Risk
Citrix PVS PXE Service	Provides the PVS PXE Boot Server functionality.	On failure of this service target devices may not be able to boot successfully if PXE booting is leveraged.
Citrix PVS Stream Service	Streams contents of the vDisk to the target device on demand.	If this service stopped it will not be possible to stream vDisk images.
Citrix PVS SOAP Service	Provides framework for external or existing solutions to interface with Provisioning services.	If this service fails PVS Server to PVS Server communication as well as PVS Console to PVS Server communication is not possible.
Citrix PVS TFTP Service	Provides the TFTP Server functionality.	On failure of this service target devices may not be able to boot if this server is used as TFTP server for the bootstrap.
Citrix PVS Two-Stage Boot Service	Provides the bootstrap functionality for devices booting by means of a BDM ISO file.	On failure of this service target devices may not be able to boot if a BDM ISO file is used.

StoreFront

StoreFront Services

Service	Functionality	Risk
Citrix Cluster Join Service	Provides Server Group join services.	This service is started when adding additional StoreFront servers to a Server Group. If this service does not start or is interrupted when this process is initiated the additional server will be unable to join the indicated Server Group and the process will result in an error.
Citrix Configuration Replication	Provides access to Delivery Services configuration information.	This service only exists on the primary StoreFront server of a Server Group. If this service is stopped additional StoreFront servers will be unable to join the Server Group and any changes made to the primary StoreFront server will not be replicated to other servers. This can result in servers within the Server Group being out of sync.

Service	Functionality	Risk
Citrix Credential Wallet	Provides a secure store of credentials. Dependencies: <ul style="list-style-type: none"> Citrix Peer Resolution Service 	If this service is stopped users will be unable to login to access their desktops or applications. Users logged into StoreFront will be unable to launch new application or desktop sessions. Existing application or desktop sessions are unaffected.
Citrix Default Domain Services	Provides authentication, change password, and other domain services.	If this service is stopped users will be unable to login to access their desktops or applications. Users currently logged in will not be affected.
Citrix Peer Resolution Service	Resolves peer names within peer-to-peer meshes.	On failure of this service both the Citrix Credential Wallet and Citrix Subscriptions store are stopped generating the risks associated with those services.
Citrix Subscriptions Store	Provides a store and replication of user subscriptions. Dependencies: <ul style="list-style-type: none"> Citrix Peer Resolution Service 	If this service is stopped Citrix Receiver cannot add, remove, and reposition applications within StoreFront. Users will need to re-add applications and all changes made to their selection of applications within the StoreFront store will not be saved or replicated to other sessions. Original user configuration will be restored once the service is restarted.
World Wide Web Publishing Service	Provides web connectivity and administration through the Internet Information Services Manager. Dependencies: <ul style="list-style-type: none"> HTTP RPC Service 	Access to published applications or published desktops will not be available through StoreFront. Users will be unable to resolve the Receiver for Web login page. Users logged into StoreFront will be unable to launch new application or desktop sessions and will need to reenter credentials when the service is restarted. Existing application or desktop sessions are unaffected.

Web Interface

Web Interface Services

Service	Functionality	Risk
World Wide Web Publishing Service	Provides web connectivity and administration through the Internet Information Services Manager. Dependencies: <ul style="list-style-type: none"> HTTP RPC Service 	Access to published applications or published desktops will not be available through Web Interface if the WWW service is not available.

Citrix License Server

Citrix License Server Services

Service	Functionality	Risk
Citrix Licensing Service	Provides licensing services for Citrix products.	Licensing mode changes to grace period when service is stopped or License Server cannot be contacted. If not monitored, functionality of Citrix products will cease after grace period expires.
Citrix Licensing Support Service	This account controls reading the license files and updating strings with license trailers (data dictionary functionality).	None
Citrix Licensing WMI	The Citrix License Management Console collects license data information using the WMI service.	None

Decision: Events Monitoring

Monitoring the Windows Event Log for unknown or critical events can help to proactively discover issues and allow administrators to understand event patterns:

- **Licensing** – Errors in the Event Log dealing with Remote Desktop licensing should be investigated. This might be a result of the installed Citrix product not being able to contact the Remote Desktop Licensing Server or the Citrix Licensing Server. If errors in the Event Log are not reviewed, users might eventually be denied access because they cannot acquire a valid license.
- **Hardware Failure** – Any event notification that relates to a hardware failure should be looked at immediately. Any device that has failed will have an impact on the performance of the system. At a minimum, a hardware failure will remove the redundancy of the component.
- **Security Warnings** – Customers should investigate security warnings or audit failure events regarding failed logons in the security log. This could be an indication that someone is attempting to compromise the servers.
- **Disk Capacity** – As the drives of a Windows system reach 90% of capacity, an event error message will be generated. To ensure continuous service, customers should poll these event errors. As the system runs out of hard disk space, the system is put at severe risk. The server might not have enough space left to service the requests of users for temporary file storage.
- **Application / Service errors** – Any event notification that relates to application or services errors should be investigated.
- **Citrix errors** – All Citrix software components will leverage the Windows Event Log for error logging. A list of the known Event Log warnings and errors issued by Citrix components can be found at the following links:

- [Event Codes Generated by PVS](#)
- [XenDesktop 7 - Event Log Messages](#)

It is important to periodically check the Event Viewer for Citrix related warnings or errors. Warnings or errors that repeatedly appear in the logs should be investigated immediately, because it may indicate a problem that could severely impact the Citrix environment if not properly resolved.

In multi-server environments it becomes easier to administer the servers when logs can be collected and reviewed from a central location. Most enterprise grade monitoring solutions provide this functionality. More sophisticated monitoring solutions enable an administrator to correlate event information with other data points such as performance metrics or availability statistics. In case the selected monitoring solution does not provide this functionality the Windows Server 2008 R2 or Windows Server 2012/2012 R2 Event Log subscription feature can be used. This feature allows administrators to receive events from multiple servers and view them from a designated collector computer. For more information please refer to the Microsoft TechNet article – [Manage Subscriptions](#).

XenServer is also capable of sending its logs to a central syslog server. The administrator sets the IP address of the syslog daemon server in the properties of each XenServer in the pool. This configuration allows administrators to capture real-time activity across multiple XenServer hosts. Further information can be found within the [XenServer Admin Guide](#).

Decision: Capacity Management

In addition to the day-to-day monitoring of system-level metrics, performance metrics should be tracked from a historical perspective to help plan for future growth as more users access the environment.

A baseline of the environment performance should be taken so that

it can be compared against performance over time. For example, if a user complains of poor performance, this baseline can be used for comparison purposes to identify if the issues are related to the user load exceeding the capacity of the environment.

An example of baseline performance metrics for capacity management would include historical data for CPU, Memory, and network utilization on the Delivery Controller and application servers or desktops.

Citrix Director

Administrators can utilize the Trends view within Citrix Director to track different parameters of the Citrix XenApp/XenDesktop deployment over time. These parameters can be leveraged for capacity planning of the Citrix environment.

From the Trends view, administrators can see historical data that is broken up into several categories including:

- **Sessions** – Provides the concurrent session usage over time enabling the ability to size the environment appropriately.
- **Connection Failures** – Gives an overview of the different types of connection failures that have occurred across different Delivery Groups.

- **Failed Desktop OS Machines** – Gives an overview of the different problems associated with failures in desktop machines.
- **Failed Server OS Machines** – Gives an overview of the different problems associated with failures in server machines.
- **Logon Performance** – Shows how long it takes for users to log on to their applications and desktops.
- **Load Evaluator Index** – Provides various performance counter-based metrics, including CPU, Memory, and Disk Usage for Server OS machines.
- **Hosted Application Usage** – Details all applications published in the site and can provide usage information about each individual applications in detail (concurrent instances, launches, usage duration, and so on).

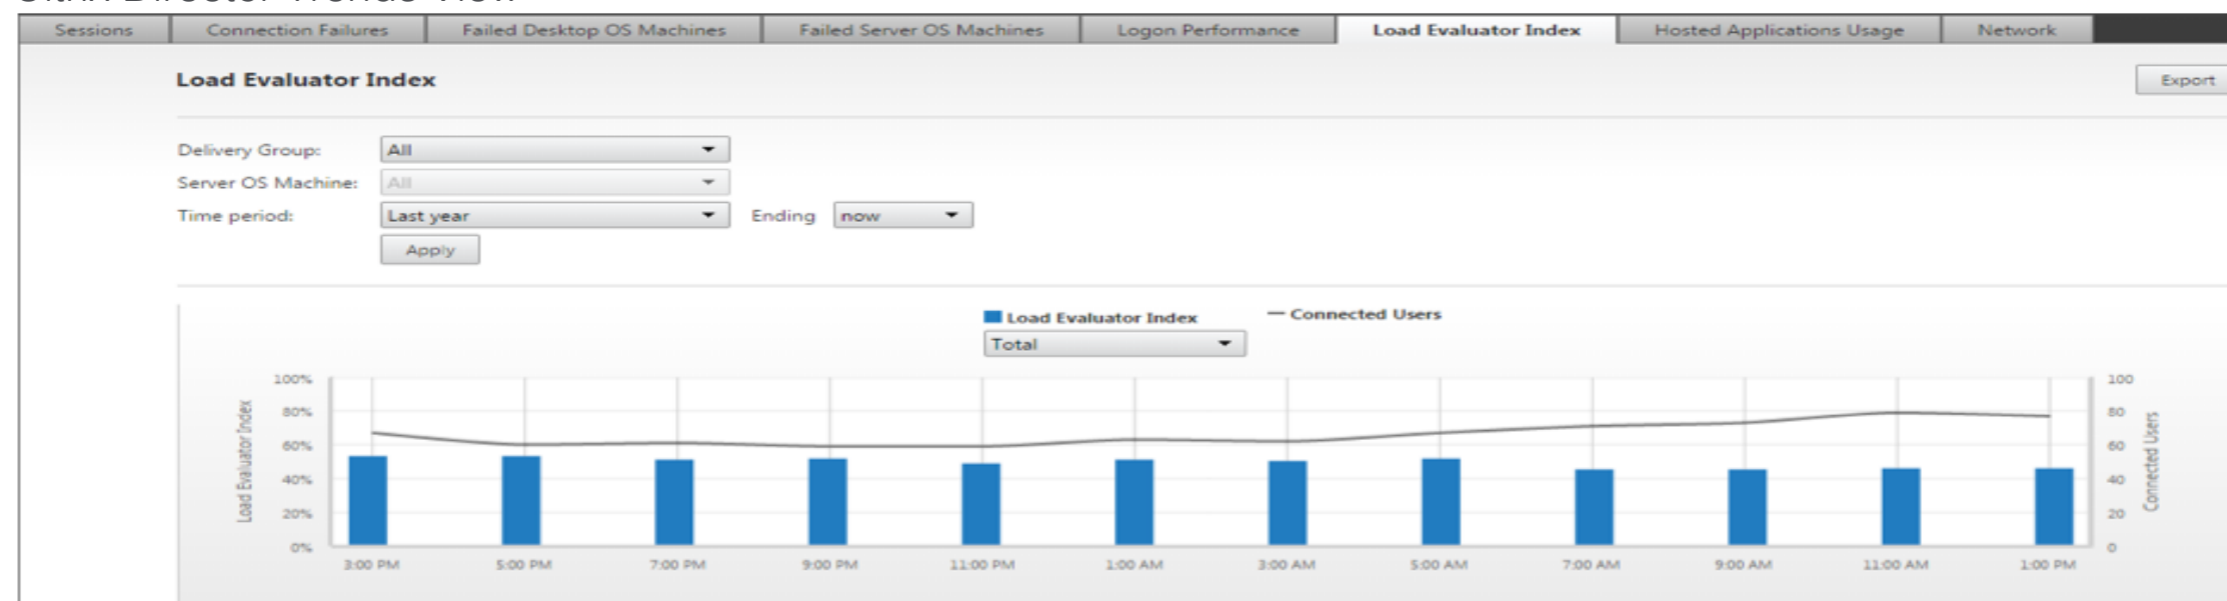
Note: Requires XenApp or XenDesktop Platinum licensing

- **Network** – Network analytics provided through NetScaler HDX Insight.

For more information on Citrix Director Trends, please refer to the following.

- Citrix Blogs – [Citrix Director: Trends Explained](#)
- Citrix Support – [CTX139382 Best Practices for Citrix Director](#)

Citrix Director Trends View



[Click here to provide feedback](#)

Appendix Quick Access Links

Glossary & Abbreviations	213
Baseline Policy Reference.....	213
Profile Management Policy	214
Microsoft Windows Policy	215
Folder Redirection Policy	217
Revision History.....	218
Authors	220
Contributors	220
Special Thank You.....	220

Glossary and Abbreviations

Bring your own Device (BYOD) – A diverse set of initiatives aiming to enhance employee satisfaction and morale by enabling the use of personal devices to supplement corporate endpoints. BYOD can replace a corporate-owned device entirely. Whatever approach an organization chooses to take, a complete, well thought-out policy is essential for embracing BYOD without increasing risk. The lack of a structured BYOD policy leaves many organizations exposed to problems such as security gaps, compliance issues and rising IT complexity.

Consumerization – A trend impacting businesses globally where a new generation of younger workers that have seen technology as a natural part of daily life and are demanding its usage at work. Unlike prior generations of employees who learned about technology primarily through the workplace, these users enter the workforce already primed by consumer technology and commonly own and use consumer products at least as sophisticated if not more so than the tools they're provided at work. These workers are demanding rather than requesting the use of consumer technologies including hardware and applications. Organizations are being forced to embrace this trend or risk alienating this new generation of workers.

Desktop virtualization – The concept of isolating then separating the operating system from the underlying client hardware used to access it. Several models exist, ranging from hosted virtual desktop infrastructure to client side hypervisor virtualization.

Baseline Policy Reference

The following sample outlines the details of initial policy configurations recommended for Profile Management, Microsoft Windows, and Folder Redirection.

The [Citrix Policy Reference](#) spreadsheet includes a baseline recommendation for Citrix User Policy and Citrix Computer Policy. Additionally, it provides a description for all Citrix policy settings available for XenDesktop.

Each policy configuration may contain the following policy settings:

Enabled – Enables the setting. Where applicable, specific settings are detailed.

Disabled – Disables the setting Note: Disabling the policy overrides lower priority policies settings.

Allow – Allows the action controlled by the setting. Where applicable, specific settings are detailed.

Prohibit – Prohibits the action controlled by the setting Note: Prohibiting a feature or functionality overrides lower priority policies settings.

Not configured – Unless specifically set, un-configured policies use default settings.

Profile Management Policy

User Policy Settings	
ICA\Printing\Client Printers	
Printer properties retention	Retained in profile only
Profile Management	
Enable Profile Management	Enabled
Process Groups	Configure groups
Path to User Store	UNC Path
Profile Management\ Advanced Settings	
Process Internet cookie files on logoff	Enabled
Profile Management\ File System	
Exclusion list – directories	AppData\Local
	AppData\LocalLow
	Java
	Local Settings
	UserData
	Downloads
	Saved Games
Profile Management\ File System\ Synchronization	
Directories to Synchronize	AppData\Local\Microsoft\Credentials
Files to Synchronize	(Example Synchronized Files for Microsoft Outlook and Google Earth)
	AppData\Local\Microsoft\Office*.qat
	AppData\Local\Microsoft\Office*.officeUI
	AppData\LocalLow\Google\GoogleEarth*.kml
Profile Management\ Profile handling	
Migration of existing profiles	Disabled

Note: For Citrix UPM 5.x

Microsoft Windows User Policy

User Policy			
Policy Path	Setting	Description	Applies To
Control Panel\ Prohibit Access to the Control Panel	Enable	Disables all control panel programs	Server OS, Desktop OS
Control Panel\ Personalization\ Enable screen saver	Enable	Enables the use of a Screen Saver	Server OS, Desktop OS
Control Panel\ Personalization\ Force specific screen saver	Enable scrnsave.scr	Forces the use of the "blank" screen saver in Windows	Server OS, Desktop OS
Control Panel\ Personalization\ Password protect the screen saver	Enabled	Forces password protection on the screen saver	Server OS, Desktop OS
Control Panel\ Personalization\ Screen saver timeout	Enabled X Minutes (default 15)	Sets the amount of time in minutes that elapse before the screen saver is activated	Server OS, Desktop OS
Desktop\ Don't save settings on exit	Enabled	Prevents users from changing some desktop configurations such as the size of the taskbar or the position of open windows on exit	Server OS
Desktop\ Hide Network Locations icon on desktop	Enabled	Removes the Network Locations icon from the desktop	Server OS
Desktop\ Prohibit user from manually redirecting Profile Folders	Enabled	Prevents users from manually changing the path to their profile folders	Server OS, Desktop OS
Desktop\ Remove Recycle Bin icon from desktop	Enabled	Removes most occurrences of the Recycle Bin icon	Server OS, Desktop OS
Start Menu and Taskbar\ Change Start Menu power button	Enabled Log Off	Set Start Menu power button functionality to Log Off user	Server OS, Desktop OS
Start Menu and Taskbar\ Prevent changes to Taskbar and Start Menu settings	Enabled	Removes the Taskbar and Start Menu settings from Settings on the Start Menu	Server OS
Start Menu and Taskbar\ Remove and prevent access to the Shut Down, Restart, Sleep and Hibernate commands	Enabled	Prevents user from performing these commands from the Start Menu or the Windows Security screen	Server OS
Start Menu and Taskbar\ Remove links and access to Windows Update	Enabled	Prevents users from connecting to the Windows Update website.	Server OS, Desktop OS
Start Menu and Taskbar\ Remove network icon from the Start Menu	Enabled	Removes the network icon from the Start Menu	Server OS, Desktop OS
Start Menu and Taskbar\ Remove Run menu from the Start Menu	Enabled	Removes the Run command from the Start Menu, Internet Explorer, and Task Manager	Server OS
System\ Prevent access to registry editing tools	Enabled	Disables the Windows Registry Editor	Server OS, Desktop OS
System\ Prevent access to the Command Prompt	Enabled	Prevents users from running the interactive command prompt "cmd.exe"	Server OS
System\ Ctrl+Alt+Del Options\ Remove Task Manager	Enabled	Prevents users from starting Task Manager	Server OS
System\ Folder Redirection\ Do not automatically make redirected folders available offline	Enabled	Prohibits redirected shell folders Contacts, Documents, Desktop, Favorites, Music, Pictures, Videos, Start Menu and AppData\Roaming from being available offline	Server OS, Desktop OS
System\ User Profiles\ Exclude Directories in Roaming Profile	Citrix, Contacts, Desktop, Downloads, Favorites, Links, Documents, Pictures, Videos, Music, Saved Games, Searches	Excludes the specified directories from the Roaming Profile	Server OS, Desktop OS
Windows Components\ Windows Update\ Remove access to use all Windows Update features	Enabled	Removes all Windows Update functions	Server OS, Desktop OS
Windows Explorer\ Do not move deleted files to the Recycle Bin	Enabled	Prohibits deleted files from being placed in the Recycle Bin. All files are permanently deleted.	Server OS, Desktop OS
Windows Explorer\ Hide these specified drives in My Computer	Enabled Local hard drives	Hides local hard drives from My Computer	Server OS
Windows Explorer\ Prevent access to drives from My Computer	Enabled Local hard drives	Prevents access to local hard drives from My Computer	Server OS

Microsoft Windows Machine Policy

Machine Policy			
Policy Path	Setting	Description	Applies To
Internet Communication settings\ Turn off Windows Customer Improvement Program	Enabled	Turns off the Windows Customer Improvement Program for all users	Server OS, Desktop OS
System\ Group Policy\ User Group Policy loopback processing mode	Merge or Replace	Applies alternate user settings when a user logs on to a computer affected by this setting	Server OS, Desktop OS
System\ Power Management\ Select an active power plan	High Performance	Specifies a power plan from a list of available plans.	Server OS, Desktop OS
System\ System Restore\ Turn off System Restore	Enabled	Turns off Windows System Restore features	Server OS, Desktop OS
System\ User Profiles\ Add the Administrators security group to the roaming users profiles	Enabled	Adds the Administrator security group to the users roaming profile path when it is first created.	Server OS, Desktop OS
System\ User Profiles\ Do not check for user ownership of Roaming Profile folders	Enabled	Disables security check for roaming profile folders	Server OS, Desktop OS
Windows Components\ AutoPlay Policies\ Turn off AutoPlay	Enabled	Turns off AutoPlay for removable devices.	Server OS
Windows Components\ Internet Explorer\ Turn off reopen last browsing session	Enabled	Disables ability to reopen the user's last browsing session	Server OS
Windows Components\ Remote Desktop Services\ RD Licensing\ License server security group	XenApp server security groups	Specifies the servers to which RDS will provide licenses	Server OS
Windows Components\ Remote Desktop Services\ Remote Desktop Session Host\ Licensing\ Set the Remote Desktop licensing mode	Per User or Per Device	Specifies the licensing mode used by Remote Desktop Server	Server OS
Windows Components\ Remote Desktop Services\ Remote Desktop Session Host\ Licensing\ Use the specified Remote Desktop license servers	Specified servers	Specifies the preferred license servers for Remote Desktop Services	Server OS
Windows Components\ Windows Update\ Configure Automatic Updates	Disabled	Specifies whether the computer system will receive automatic updates through the Windows Update process.	Server OS, Desktop OS

Folder Redirection Policy

User Policy\Windows Settings\Security Settings\Folder Redirection

Folder	Setting	Options
AppData (Roaming)	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Contacts	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Desktop	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Documents	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Disabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Downloads	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Favorites	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Links	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Music	Follow the Documents Folder	
Pictures	Follow the Documents Folder	
Saved Games	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Searches	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Start Menu	Basic	Grant User Exclusive Rights: Disabled Move Contents to new location: Enabled Apply Policy to Windows 2000, Windows XP, Windows 2003: Disabled Policy Removal Behavior: Leave Contents
Videos	Follow the Documents Folder	

Revision History

Revision Number

- | | | | |
|-----|---|-----|--|
| 1.0 | XD7 Handbook Released
<i>By Andy Baker - Oct 2013</i> | 4.0 | Multiple updates and restructuring of the Virtual Desktop Handbook including the new Monitor Section by Kevin Nardone.
Added License Server Chapter by Rafael Jose Gomez.
Added Database updates by Ed Duncan.
<i>By Rafael Jose Gomez - Aug, 2014</i> |
| 2.0 | Restructured Assess phase
Added PVS 7.1 chapter to Design phase.
Added Bandwidth chapter to Design phase.
Updated Introduction.
<i>By Andy Baker - Nov 2013</i> | 5.0 | Added Scalability updates, including Hardware Sizing and Resource Allocation by Amit Ben-Chanoch.
Added Disaster Recovery sections by Roger LaMarca.
<i>By Rafael Jose Gomez - Sep 15, 2014</i> |
| 2.1 | Minor updates for 7.1
<i>By Amit Ben-Chanoch - Dec 2013</i> | 5.1 | Updated HDX Encoding Methods and Session Bandwidth sections by Amit Ben-Chanoch.
Restructuring of several sections of the handbook.
<i>By Rafael Jose Gomez - Sep 17, 2014</i> |
| 3.0 | Added Hyper-V 2012 R2 chapter to Design Phase by Ed Duncan.
Added StoreFront 2.5 chapter to Design Phase by Rafael Jose Gomez.
<i>By Rafael Jose Gomez - Apr 2014</i> | 5.2 | Updated and added the Active Directory section by Rafael Jose Gomez
<i>By Rafael Jose Gomez - Sep 25, 2014</i> |
| 3.1 | Updated the PVS chapter to Design Phase by Ed Duncan.
Added Storage chapter to Design Phase by Ed Duncan.
<i>By Rafael Jose Gomez - May 2014</i> | 5.3 | Minor updates to Desktop Host Sizing, PVS Write Cache Destinations and the Resource Layer.
<i>By Rafael Jose Gomez - Sep 26, 2014</i> |
| 3.2 | Added Updated StoreFront 2.5 chapter to Design Phase by Rafael Jose Gomez.
<i>By Rafael Jose Gomez - June 2014</i> | 6.0 | Updated sections for XenDesktop 7.6 release by Ed Duncan.
Updated the IOPS Requirements by Workload table to show new data for Windows 7 with PVS medium and heavy workloads as well as MCS with a heavy workload.
Updated the Folder Redirection Matrix to show Application Data is recommended to be redirected for roaming and hybrid profiles.
Updated the User Policy Settings table for UPM 5.x. Many |
| 3.3 | Added Printing section to the Design Phase by Ed Duncan.
Minor updates through the handbook.
<i>By Rafael Jose Gomez - July 2014</i> | | |

policy settings are now handled by UPM and have been removed from the table.

Minor revisions.

By Rafael Jose Gomez - Oct 01, 2014

6.0.1 Added information about Unauthenticated user access through StoreFront.

By Rafael Jose Gomez - Oct 02, 2014

6.0.2 Updated the SQL High Availability options to show that the XenClient Database is not supported with SQL AlwaysOn.

By Rafael Jose Gomez - Oct 21, 2014

6.1 Updated section “Printers - Auto-creation” to include “Auto-create the client’s default printer only”.

Updated all of the equations used for:

1. Estimation of number of physical cores required for XenDesktop workload.
2. Physical cores required for XenDesktop Sizing example.
3. Estimation of number of physical cores required for XenApp.
4. Hardware Sizing Example- Physical cores required with dual socket machines with XenDesktop and XenApp.
5. Number of hosts required with a 16 core host.

By Rafael Jose Gomez - Oct 30, 2014

7.0 Added two new sections to the Monitor section - Operations and Monitoring by Kevin Nardone.

Minor revisions.

By Rafael Jose Gomez - Nov 08, 2014

7.1 The Disaster Recovery section has been properly merged into the Monitor - Operations section.

Corrected a typo where the decision “Availability Testing” was mislabeled as “Change Control”.

Updated the support staffing table and verbiage.

By Rafael Jose Gomez - Nov 10, 2014

7.2 “Web Interface Features not Supported by StoreFront 2.5” table has been updated to StoreFront 2.6.

1. Folder view is now the default view of the mandatory store.
2. The Store Customization SDK (Resource Filtering SDK) is now available.
3. Kerberos Constrained Delegation (KCD) is supported in StoreFront 2.6.
4. User notification is allowed.
5. Session Timeout is available in the Admin UI with SF 2.6.

Minor revisions.

By Rafael Jose Gomez - Nov 16, 2014

Authors

The creation of the handbook is a time consuming process and requires real deployment experience across many scenarios. Citrix Worldwide Consulting would like to thank the authors who wrote the chapters within this handbook:

- **Daniel Feller**
- **Andy Baker**
- **Thomas Berger**
- **Rich Meesters**
- **Matthew Brooks**
- **Adeel Arshed**
- **Martin Zugec**
- **Roger LaMarca**
- **Ed Duncan**
- **Kevin Nardone**
- **Amit Ben-Chanoch**
- **Rafael Jose Gomez**

Contributors

Citrix Worldwide Consulting would like to thank all of the individuals that contributed to the Citrix Virtual Desktop Handbook:

- **Michael Havens**
- **Maria Chang**
- **Uzair Ali**
- **Ryan F. Robott**
- **Pablo Legorreta**
- **Steven Krueger**
- **Josh Fu**

Special Thank You

Citrix Worldwide Consulting would also like to thank the following individuals for their assistance with the Virtual Desktop Handbook.

- **Jose Caceres**
- **Alee Abbas**