High Availability for Citrix XenDesktop

Enhancing XenDesktop Availability with NetScaler Reference Architecture

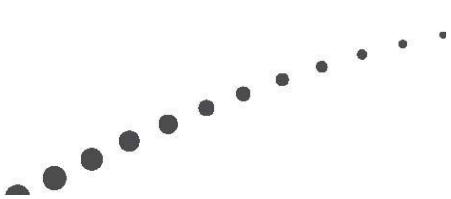

Contents

Introduction

As organizations simplify desktop management activities through the use of desktop virtualization, creating a highly available architecture becomes ever more important. Historically, users were minimally impacted by desktop outages. In many situations, a desktop outage typically impacted a single user, as one desktop device belonged to a single user. In a virtualized desktop operating model, there is the potential for an outage to impact hundreds or thousands of users if the environment is not designed appropriately.

This white paper focuses on the fault tolerant and high-availability options integrated with Citrix XenDesktop 5.5 as they align with three different focus areas:

- **Layer 1 – Desktop Availability:** Users should always be able to work with an available desktop.
- **Layer 2 – High Availability:** Failures within a site should not impact desktop availability.
- **Layer 3 – Business Continuity/Disaster Recovery:** The loss of an entire site should not prevent users from accessing their desktop.

By focusing on three distinct layers, an organization can feel confident in the operational availability of the architecture, even in the event of a catastrophic failure.

Desktop Availability

When delivered via XenDesktop, a user's desktop environment is built around four different layers:

- Hardware
- Operating System
- Applications
- Personalization

In order to provide a user with the correct desktop, each one of the desktop layers must be highly-available. The focus of this whitepaper is on delivering high availability for hosted virtual desktops. Subsequent whitepapers will discuss application availability and load balancing provisioning services.

Hardware

With XenDesktop, the hardware becomes a non-issue for providing availability. Previously, a user's physical desktop contained their operating environment. The loss of the desktop meant the loss of work until the device was repaired.

In a virtualized desktop model, the user's operating system, applications and personal settings are abstracted from the core hardware. Essentially, the user is able to jump across different machines without impacting the usability of their desktop. This provides fault tolerance in the following ways:

- Endpoint Failure: If the user's physical endpoint fails, any new endpoint can be used to gain access to the virtual desktop.
- Hosted Virtual Desktop Failure: If the server delivering the hosted virtual desktop fails, the user can immediately initiate a new connection and the XenDesktop connection broker will direct the user to a virtual desktop with all of their applications and personalization settings intact. As with a physical desktop failure, the user will lose any open and unsaved edits, but recovery to a new virtual desktop will be much faster.
- As can be seen, the act of transforming the desktop into a virtual desktop overcomes productivity loss associated with desktop device failure.

Operating System

The second major reason why users of physical desktops have a loss of productivity is due to the corruption of the operating system, which results in failure to boot, blue screens or system crashes. With XenDesktop, this area of availability is mitigated by delivering a standardized, proven, optimized image to the virtual desktop through the FlexCast for Hosted VDI Desktop models. This paper focuses on delivery through the pooled and dedicated models delivered through Machine Creation Services. For more information on Citrix FlexCast models, visit flexcast.citrix.com.

In the pooled desktop model, a snapshot of the original master disk is created. The combination of this snapshot as well as a difference disk and a unique identity disk create a virtual desktop where users interact normally. Changes made to the actual image are stored in the difference disk. Once the user is done with the desktop session and reboots, the changes made to the actual operating system image are erased. When the user makes a new connection to their virtual desktop, they receive a brand new, optimized environment. Any changes made to the user portion of the desktop are managed via the personalization layer.

The dedicated desktop model operates similar to the pooled model except changes made to the desktop that are stored within the difference disk remain persistent across reboots. A corruption in the virtual desktop would require the user to start from a clean image.

There are scenarios where an admin-level change must be made in the environment, like installing a new hotfix. In the pooled model, hotfixes can be applied to the catalog master image, and delivered to the users through a reboot of the virtual desktops. If an admin-level change results in an issue, the administrator can revert back to a previous desktop image. In the case of dedicated virtual desktops, admin level changes are applied and managed using traditional desktop management tools.

By using the Machine Creation Services component of XenDesktop as well as image snapshots, users are assured of an optimized desktop while minimizing the impact due to corruption.

Applications

Applications can be delivered into the virtual desktop in three different ways. Users can be protected from failures within the application due to corruption, deletion, or other means via the processes used to deliver those applications, particularly when dealing with pooled desktops, or dedicated desktops with streamed or hosted applications. Making the core services required to deliver streamed and hosted applications highly available will be discussed in a subsequent whitepaper.

- Installed: Installed applications are part of the base virtual desktop image. In a pooled model, the image is delivered as read-only, where user-level changes are stored in a temporary cache. When the virtual desktop is restarted, the changes are lost and the virtual desktop reverts to the base image. This keeps the operating system and the installed applications in a pristine format, free from any corruptions or misconfigurations by the user. Dedicated virtual desktops manage installed applications through traditional systems management model.

- Streamed: Streamed applications are delivered to the virtual desktop over the network as requested. The applications are stored on a file server (Application Hub) as an application profile. When users launch the application, portions are delivered to the virtual desktop. The streaming process verifies the files exist in the correct state during launch. If not, the correct files are streamed from the Application Hub automatically and seamlessly.

- Hosted: Hosted applications are virtualized and executed on a XenApp server; therefore the application does not impact the virtual desktop. Hosted applications can either be installed or streamed. If the applications are streamed to the XenApp servers, they will self-heal due to the file checks built into the application streaming process.

High Availability

Without any modifications, users have a base level of fault tolerance with XenDesktop. However, to provide a greater level of availability for the delivery of virtual desktops, the following areas can be augmented based on the following areas of focus:

- Virtual desktop hosting platform
- Operating system delivery
- Desktop delivery
- Desktop controllers

As can be seen in the following figure and subsections, these areas of focus provide a highly-available virtual desktop solution.

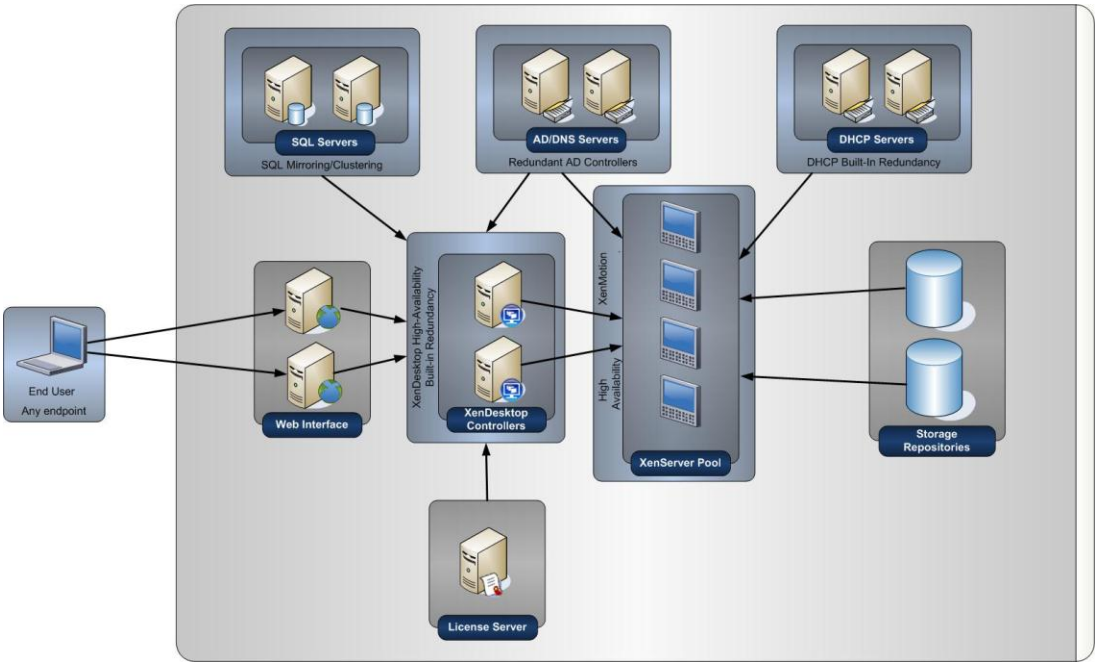

Figure 1: XenDesktop High-Availability

Virtual Desktop Hosting Platform

The first step is to provide a highly-available virtualization infrastructure to host the virtual desktops. The platform used must be able to accommodate planned and unplanned outages. The system must be able to recover seamlessly. Using XenServer as the underlying infrastructure platform allows for the following levels of fault tolerance:

- Hypervisor Failure: Hypervisor failure needs to be considered in order to avoid situations where the loss of a single physical server can affect users running virtual desktops from the configuration. When configured for high availability, the loss of any one server within the virtualization pool will not impact the other servers within the pool. While a XenServer pool has a dedicated master server responsible for the proper functioning of the pool, any other server in the pool can become the master so the loss of any server does not create a single point of failure.

- Imminent Failures: Imminent failures are issues with the physical hardware that slowly degrades service. These types of failures can be overcome without any user interruption with XenServer's XenMotion technology, along with proper storage design. XenMotion allows any running virtual machine to be migrated to another XenServer within the same resource pool without disrupting the user. In many XenDesktop architectures, only the infrastructure components are protected with XenMotion while the virtual desktops are not, as XenMotion cannot be used for some desktop delivery infrastructures, such as

Machine Creation Services with IntelliCache. Of course this is a business decision and impacts the overall design.

- Critical Failures: A critical failure is an unforeseen failure that causes the physical server to crash. These types of failures typically have no warnings, which results in locally running sessions being lost. XenServer incorporates high-availability features to allow for the auto-restart of different systems based on priority levels. This type of functionality is not needed for pooled virtual desktops, but may be necessary for dedicated virtual desktops. When using pooled virtual desktops, a predefined number of idle virtual desktops can be configured to ensure that desktops are always online during normal business hours. If a server crashes, the users can immediately connect to a new virtual desktop without delay. Although this feature is not needed for the virtual desktops, it is advisable to set priority levels for other critical infrastructure components like: XenDesktop Controllers, Web Interface servers, etc. Best Practice recommendations can be found in the XenDesktop Design Handbook (CTX120760).

Of course, it is also recommended that high availability of the network and storage components be considered when designing and building a hosting platform for virtual desktops. While beyond the scope of this whitepaper, Citrix has developed best practices for planning and designing storage for XenDesktop. Network infrastructure should always be deployed with redundant configurations, from the server side (redundant NICs) through to the end user infrastructure to avoid single points of failure.

Operating System Delivery

With a foundation to host the virtual desktops, the delivery mechanism must be highly-available. In order for the operating system delivery to always function, the XenDesktop Machine Creation Services component, Active Directory and all related network services must be designed for high availability.

- DHCP: When a new virtual desktop is started, it requests an address from DHCP. The DHCP system used within the organization must be designed so the loss of one server will not prevent new DHCP requests from being fulfilled. Typically, enterprise DHCP solutions are built with high-availability options included. The most common methods for redundant DHCP are to utilize split scopes or to cluster DHCP servers. Specific pros and cons can be found in Microsoft TechNet under Design Options for DHCP Availability and Fault Tolerance. This should be analyzed and reviewed before production rollout.

- Active Directory and DNS: AD and DNS services are required to resolve FQDNs and authenticate both virtual desktop machine accounts and user access to the virtual desktops. As AD and DNS services are critical components of an enterprise IT architecture, they are usually designed with high availability in mind. Active Directory has built in availability features such as multi-master replication and Active Directory

integrated DNS. Generally, using these features will address availability needs for AD and DNS.

- Desktop Delivery Controllers: With Machine Creation Services, the Desktop Delivery Controller provides the mechanism to create and deliver the pooled virtual desktops for the users. The DDCs are designed with built in redundancy, and within a site, multiple DDCs automatically load balance site wide services and provide full redundancy in case of the failure of a controller. Additional considerations for Desktop Delivery Controllers and other delivery components are addressed in the next section.

Delivering the resources to the user, whether those resources are desktops or applications, is the responsibility of Web Interface and the Brokers. As discussed in the following sections, providing high-availability to these components is possible with the use of smart monitors.

Web Interface

In a full-scale XenDesktop infrastructure, the Web Interface servers are responsible for delivering the desktop and the applications to the users. Initially, a user makes a connection to Web Interface. Whether used through a browser or by a desktop appliance/thin client, Web Interface is a critical component for users. From the interface, users enter their credentials and select their desktops or applications. Based on the user interactions, Web Interface communicates with the XML Service for the XenDesktop sites to fulfill the user request. In situations where the sever hosting Web Interface fails, or the IIS service fails or Web Interface encounters issues, users would be unable to connect to the environment.

NetScaler provides an intelligent monitor for Web Interface. By launching a connection to the Citrix Web Interface, the monitor determines if the server is available, if the web service is running and if the Web Interface site is functioning and responding. If disruptions in the service are identified, NetScaler generates an alert.

The alert is then used as part of the NetScaler load balancing algorithm. If a Web Interface server is not responding correctly, the server is removed from the load balancing pool until the problem is corrected. New user requests are routed only to the available Web Interface servers.

Providing high-availability for Web Interface goes beyond providing an available server. Configuring high-availability for Web interface allows an administrator to determine the best way of balancing users across servers. For example, one Web Interface server could be the least loaded server but is currently busy with another process. NetScaler could direct new user requests to another Web Interface server that responds first instead of directing based on user connections. Intelligent monitoring incorporated into load balancing allows for users to get the fastest application delivery experience.

Desktop Controllers

The XenDesktop controllers are responsible for

- Maintaining the proper level of idle desktops to allow for instantaneous connections
- Monitoring state of online and connected virtual desktops
- Shutting down virtual desktops as needed

Without a XenDesktop controller, users would not be able to connect to new virtual desktop session. However, the XenDesktop controllers are designed with built-in redundancy. Communication between the Virtual Desktop Agent (VDA) and the XenDesktop controller through external means is neither required nor recommended. The VDA utilizes a list of controller addresses obtained from the registry (by default), and will attempt to connect to another controller if the first randomly selected from the list does not respond or registration fails. An external load balancer could also cause disruptions if it attempts to change the connection once it has been established as the VDA is dynamically registered to a single controller and is expecting to communicate with that controller.

XML Service

A critical component of any XenDesktop environment is the XML Service, which is a part of the XenDesktop Controller broker service. The broker service is responsible for user authentication, resource enumeration and resource launching processes. A failure in the broker will result in users being unable to start their virtual desktop. The following diagram shows how critical the XML Service is to users.

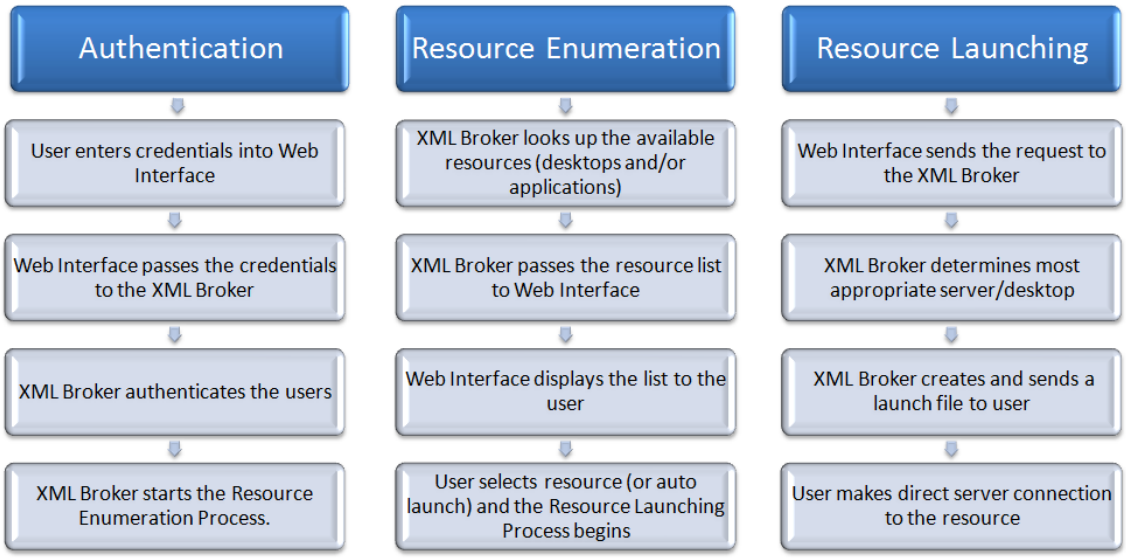

Figure 2: XML Broker Process Flow

The broker service is the link between the users and the XenDesktop infrastructure, which makes it critical. Monitoring the broker service and the DDCs is not a trivial task as the monitoring must go beyond simply identifying if the service is running to identifying if the service is responding appropriately. If the broker responds incorrectly, the Web Interface server could get stuck in a request/response loop resulting in users not gaining access to their resource.

NetScaler provides intelligent monitoring of the XML Service through the use of pre-configured monitor templates for the XenDesktop DDC. The monitoring determines if the XML Service is running and if the requested information from the DDC responds in a timely

manner and with expected information. If the monitor determines an unexpected result or complete failure to respond, NetScaler creates an alert, which is used in the load balancing algorithm. NetScaler dynamically adjusts the environment to bypass the failed DDC. If the XML Service functionality is restored, NetScaler automatically detects and incorporates the DDC back into the environment.

Providing high-availability to the XML Service is more than load balancing in the event of a failure. NetScaler also load balances connections across multiple DDCs to help spread the load and to provide a better and faster application initialization experience. In many organizations, a major shift change starts at 8 or 9 AM. This results in a huge load on the DDCs. Integrating NetScaler into the environment distributes the load across multiple Brokers while also monitoring and allocating requests based on availability.

SQL Server

The SQL database provides the foundation for all configuration information in the XenDesktop site, and acts as a message bus for communication between XenDesktop controllers. All desktop configuration and current utilization information is stored in the database. The server is crucial to the continuous operation of the XenDesktop site and if it fails, no new connections to virtual desktops will be possible. It is recommended that the SQL database be made highly available through SQL Mirroring and/or SQL Clustering technologies. For more information on SQL mirroring and clustering see the Microsoft whitepaper on High Availability with SQL Server 2008 R2.

Business Continuity/Disaster Recovery

As users continue to be spread across different regions, there is a requirement for many XenDesktop environments to provide business continuity and disaster recovery. Oftentimes, this results in one of the following actions by the administrators:

- Users are told to use a different address in the event of a failure at the main data center
- A manual change is made to the DNS table, which will direct user requests to the backup data center

Of course these items are not automatic based on monitors and they also idle roughly 50% of the hardware stored at the backup data center because these business continuity solutions are configured in an active-passive mode. In the active-passive mode, only one site is active at a time. When a failure occurs, the passive site becomes active, but until this occurs, the backup site is in standby mode. However, in an active-active model where there are multiple data centers, users require a solution that not only provides them access, but also provides them access to their desktop in their preferred data center, which is where the user's data is located. If the user cannot access the data required to perform their job roles, then failover makes little sense, whether automated or manual. Failover can be automated through a combination of these functions, depending upon the specific requirement:

- NetScaler Global Server Load Balancing
- XenDesktop Roaming User Support
- Disaster Recovery

Global Server Load Balancing

The first step towards providing a business continuity and disaster recovery solution for XenDesktop is to utilize the global server load balancing functionality of NetScaler.

Figure 3: High-Availability Architecture

External users make a request using a fully qualified domain name to the organization's Access Gateway login page, which is delivered from NetScaler. If the user is internal, the user would make a fully qualified domain name request to the Citrix Web Interface, which is load balanced by NetScaler. Providing the best site is the responsibility of the NetScaler. Each site communicates with the other sites using the Metric Exchange Protocol (MEP). This lets each NetScaler know the status of the other sites so each NetScaler can provide an accurate response to incoming user requests. The responses must take into account the following:

- Site Availability: Does the data center have all critical components available (Web Interface, Access Gateway, XML Broker, etc.)? NetScaler uses monitors to check the availability of the Access Gateway, Web Interface and XenDesktop DDC components within the sites in the GSLB configuration.

- Best Available Site: Based on the sites with critical component availability, which one is the best for this particular user? The best site is based on the current situation and configuration. With NetScaler, the best site could be the site which responds the fastest, the site that is closest in geographical proximity to the user, or it could be in a round robin flow. The choice should be based on the organizational architecture and needs.

Once these items are taken into account, the user is presented with an Access Gateway logon page at the best site.

Roaming User Support

When it comes to roaming user support, there are multiple considerations that need to be made to design and deliver an optimal solution for the users. The speed of the network and the location of user data relative to the virtual desktop must be considered. The global server load balancing functionality gets the user to the best responding site, but sometimes this is not the best site from the user perspective. The closest site from a proximity perspective is measured from the DNS server location, not the user location. For example, if a user's data is located in a data center in North America and they are currently travelling in Asia, global server load balancing would direct the user to the Asia data center. This would require the virtual desktop in the Asia data center to communicate with the North America data center to retrieve the user's data. If this network link is slow, the result would be an underperforming virtual desktop from the user perspective.

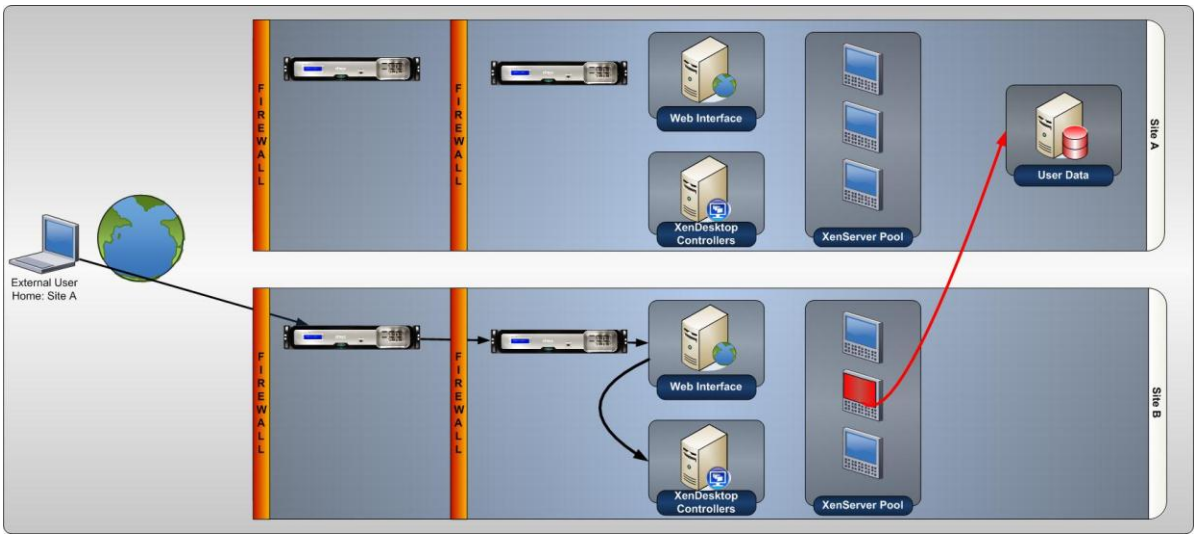

Figure 4: Roaming User Remote Site

If this is the case, the ideal solution would be to have the user would connect to the Asia site, based on GSLB, and then be re-routed to the North America data center for the virtual desktop, thus keeping the user's data local to the North America data center. The user would interact with the virtual desktop via the Citrix HDX, which is optimized for WAN environments.

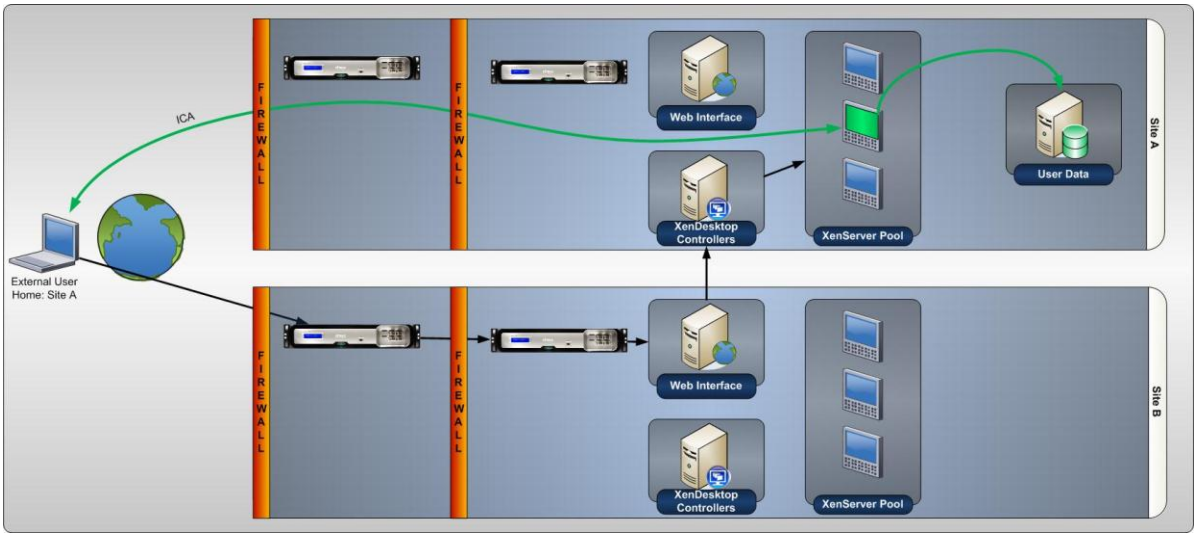

Figure 5: Roaming User Local Site

This approach provides global redundancy and also provides the user with the best virtual desktop experience as the only thing crossing the WAN link is HDX instead of application data. This can be accomplished by

1. Creating Active Directory groups corresponding to each geographic site.
2. Adding the appropriate users into the defined Active Directory groups. These groups are used to define the user's "Home" site.
3. Configuring all Web Interface servers with defined Active Directory groups and preferred "home" site.

When a user connects to the GSLB configuration, the NetScaler selects the best site based on the load balancing schemes configured. The user authenticates with the Access Gateway, and is connected to a load balanced Web Interface server. Based on the user's group membership, the Web Interface directs the user to their configured home site to access the virtual desktop and user data. The user's HDX traffic will traverse the corporate WAN structure behind the Access Gateway; no HDX traffic transits across the AG.

Disaster Recovery

If users are directed to their home site with component-level availability, there is still the potential of not enough capacity to support the user's virtual desktop. This could be due to a large percentage of servers being offline or communication failing between components. Regardless of the reason, if the Web Interface is not able to provide the user with an available virtual desktop, it can be configured to automatically fall back to a defined recovery site. The recovery site architecture could look like the following:

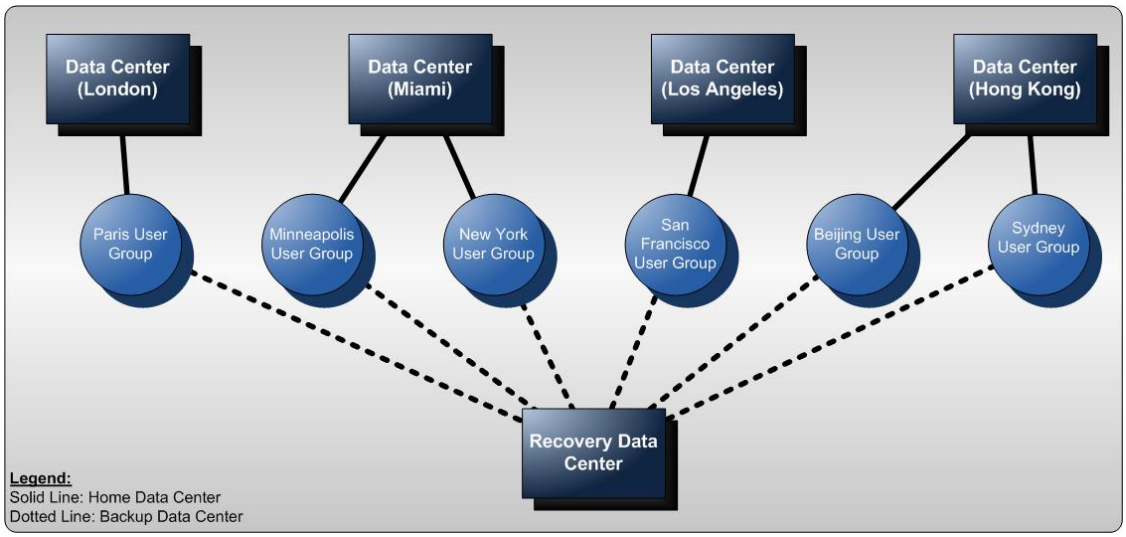

Figure 6: Fallback Data Center

The recovery site, defined by the administrator, allows the user to work in a virtual desktop, but if the user's data is not synchronized between the home and recovery sites, the user will experience a degraded experience. It is advisable for the administrator to configure the recovery site with a message informing the user of the potential degradation. By using automated recovery sites, the user will receive a virtual desktop, even in the event of a catastrophic failure. Use of automated recovery should be carefully considered within the context of an organization's IT operational model and change management process. It needs to be determined whether it is optimal for an organization to have an automated failover and deal with potential issues with performance, especially if the failover is not required (i.e. no actual disaster) or if a semi-automated process should be used by manually reconfiguring site preferences.

Conclusion

As the desktop operating environment is moved from the endpoint into the data center, fault tolerance, high availability and disaster recovery are even more critical. If the infrastructure suffers a failure, a large percentage of the user population is impacted. Being able to provide fault tolerance at each level mitigates the risk of centralized computing. When providing high-availability to a desktop virtualization solution, the architecture must include availability at each of the three layers:

- Desktop availability
- High Availability
- Business Continuity/Disaster Recovery

Product Versions

Product	Version
XenDesktop	4.0, 5.0, 5.5
NetScaler	9.3

Revision History

Revision	Change Description	Updated By	Date
1.0	Document released	Daniel Feller – Lead Architect	October 28, 2009
2.0	Document revised	Rich Meesters – Architect	October 14, 2011

About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. It's Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2010 was $1.9 billion.