CITRIX

XenApp and XenDesktop concepts and deployment on CloudPlatform

XenApp and XenDesktop concepts and deployment on CloudPlatform

Prepared by: Paul Howard

Commissioning Editor: Jan Penovich

Version: 1.0

Last Updated: July 18, 2014

Table of Contents

Part one – Overview of concepts and prerequisite infrastructure for Citrix CloudPlatform

This is an overview of concepts and prerequisite infrastructure you should know before setting up a XenApp or XenDesktop site on a CloudPlatform infrastructure. For information about deploying XenApp or XenDesktop with CloudPlatform, refer to *Part two – Deployment*.

Roles needed to create a cloud-based site

People or teams must fill three distinct roles before creating a cloud-based XenApp or XenDesktop site:

- **Infrastructure engineer** – This individual (or team/organization) hosts hardware and its presentation as a CloudPlatform cloud. The infrastructure engineer installs CloudPlatform and configures hosts, pods, clusters, zones, storage, and networking. The infrastructure engineer provides a cloud that can, in theory, be employed for any purpose (for example, Amazon can be considered to be the infrastructure engineers for their cloud).
- **Service designer** – This individual (or team) consumes the fixed cloud infrastructure defined by the infrastructure engineer. The service designer uses the raw cloud infrastructure to build a specific useful service for ultimate delivery to end users.
  Note: In this document, the specific service is XenDesktop.
- **Site administrator** – This individual (or team) consumes the cloud-based XenDesktop service and uses Studio to define the site and the facilities within it. XenDesktop's own Delegated Administration model can then be used to delegate parts of this management to further subadministrators. In this document, these are part of the same overall role.

At the simple end of the spectrum, a single individual acts in all three roles at the same time. At the complex end, each role might be filled by entire teams or organizations, with further subdivisions among them.

Assets and responsibilities

Infrastructure engineer is responsible for the CloudPlatform installation and its associated hosting hardware and overall organization. Because the infrastructure can be completely independent of XenDesktop, there are no XenDesktop-specific activities required of the **infrastructure engineer.**

Service designer and **site administrator** are responsible for all XenDesktop-specific activities. Of those two, the service designer is the person/team who transforms raw CloudPlatform infrastructure into a functioning XenDesktop workload.

The figure below shows the separation of responsibilities, and how various assets flow through the system. As a general rule, the service designer consumes the infrastructure, which is then consumed by the site administrator. This gives rise to the three-tier flow shown below. In

practice, it might sometimes be the same organization(s) or individuals who share responsibility for any or all of these steps.

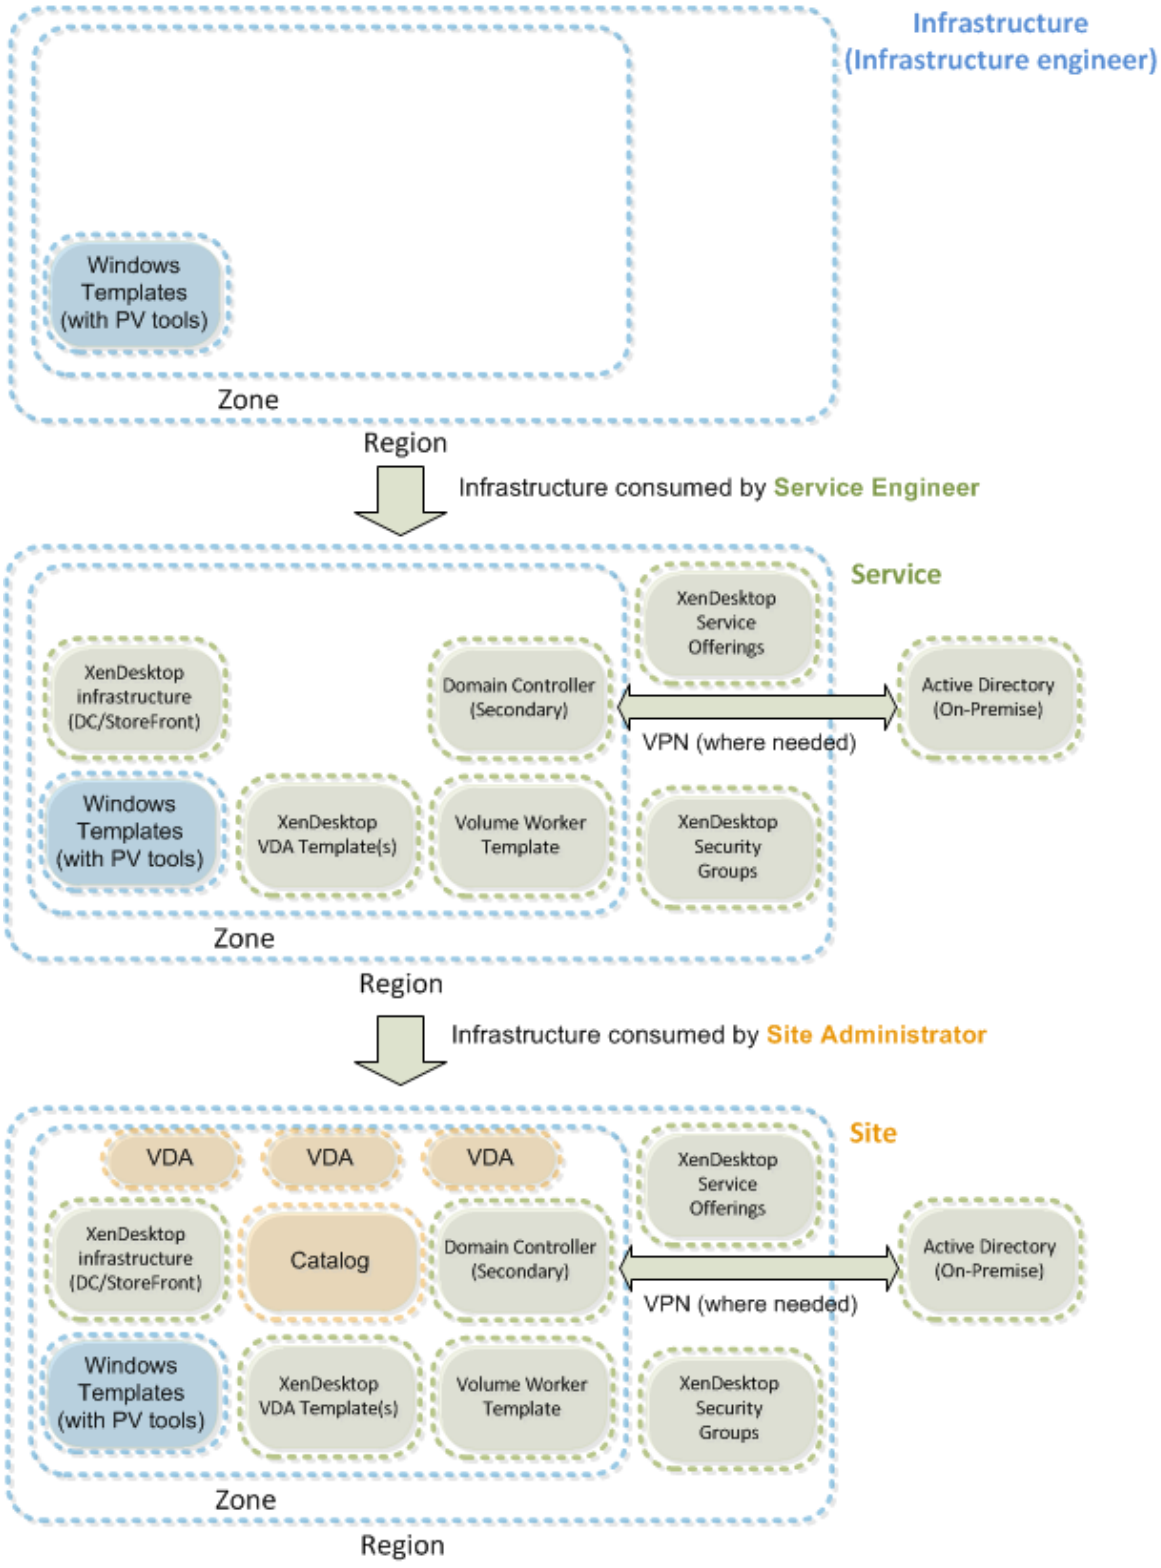

IaaS and DaaS

- Infrastructure as a Service (IaaS) refers to the delivery of cloud infrastructure from the infrastructure engineer.
- Desktops as a Service (DaaS) refers to the delivery of cloud-hosted desktops and applications to end users. The service designer consumes an IaaS cloud from the infrastructure engineer and delivers a DaaS cloud service to end users.

Public and private clouds

In a public cloud, the infrastructure engineer makes the cloud available for general use (IaaS), and typically bills it to consumers (or tenants), who then run their own custom services on infrastructure that they do not own. In a public cloud, there is a high degree of separation between the infrastructure engineer and the service designer.

In a private cloud, the same organization or entity providing the virtual desktop service owns the infrastructure. This implies a closer link between the infrastructure engineer and the service designer.

These terms reflect the exposure of the CloudPlatform infrastructure only, and not the exposure of the XenDesktop service. For example, an enterprise might use CloudPlatform internally in their datacenter to create a DaaS offering that is then exposed publicly. However, while the DaaS service is public, the infrastructure cloud itself is private, so this would qualify as a private cloud.

Regional isolation

Whether the cloud is private or public, its infrastructure might be composed of multiple regions. (Multi-region infrastructures are more likely to be the property of public clouds than private. Public cloud services generally operate on bigger scales and are more expansive geographically, so the regional isolation makes more sense.)

In CloudPlatform, regional isolations work in the same way as they do on Amazon Web Services (AWS): that is, they are 100% isolated services, with no crossover of resources. Cloud regions are, in effect, separate clouds that offer a similar looking service. Regions connect only in one simple way: each region provides a facility for discovering the names and connection details for the other regions; otherwise, they are completely segregated.

XenDesktop 7.5 has a one-region-per-site restriction. Any XenDesktop site must exist within one cloud region only. To use multiple regions, you must have multiple XenDesktop sites, potentially using StoreFront to aggregate them.

CloudPlatform API

If the cloud is public, it is assumed to be exposing the CloudPlatform REST API directly; that is, the user portion of the API documented here.

Some public IaaS clouds might use CloudPlatform internally but then expose its services using a proprietary API. Such public cloud offerings are not compatible with XenDesktop 7.5, because XenDesktop has been integrated only with the CloudPlatform API.

Time synchronization

It is important to have a common source of time throughout the system. This is especially critical when running a XenDesktop workload, because SSL certificates are used to secure some communications between the Delivery Controller and some virtualized services that run within the cloud zone. Poorly-synchronized time can cause the authentication of these channels to fail, which makes it impossible to provision desktops successfully.

Citrix recommends having a Network Time Protocol (NTP) server within the system. Configure all hosts in all cloud zones to obtain their time signal from this central server. Likewise, the Window Active Directory domain should use the same time source.

Check time synchronization carefully, because time disagreements are a very common source of errors within a XenDesktop workload.

Security groups

CloudPlatform supports the use of security groups in basic networking zones only.

Citrix recommends that you do not use security groups in advanced zones when using XenDesktop workloads within the cloud. (Security groups can be used in basic zones for isolation if needed.)

Prerequisite infrastructure

This section describes what assets the infrastructure engineer is required to make available to create a cloud to support a XenDesktop workload:

- At least one region, exposing an endpoint for the CloudStack User API and Web Console.
- At least one zone with zone-wide secondary storage.
- At least one cluster, with cluster-level primary storage.
- At least one host, with high-performance local storage. (The local storage requirement is not strictly necessary, but the best practice for XenDesktop is for pooled desktops to execute on host-attached storage. Any type of desktop can also execute on shared primary storage, which would probably be SAN-based).
- Physical networking.

In addition, the infrastructure engineer should make some operating system templates available. All of these must be extended where necessary with the appropriate paravirtual tools for the hypervisor platform that the infrastructure is using (for example, XenTools for XenServer-based hosts). Templates might also be extended with additional software deemed appropriate by the cloud provider - there could be some provider-specific tools or branding elements.

Operating systems

- **Windows Server 2008 R2 or Windows Server 2012 R2** – XenDesktop services such as the Delivery Controller, database server, domain controller, and StoreFront require these systems. They are also used for RDS workers and server VDI desktops.
- **Windows 7 or Windows 8.1** – A public cloud provider might be prevented from making these available due to license-reselling issues. (Amazon Web Services, for example, is currently not able to provide such templates.) Without these templates, XenDesktop services are limited to server VDI or RDS-type deployments.

Infrastructure providers can also make extended instances available, such as Windows Server with SQL Server pre-installed. XenDesktop does not require such templates, although the installation process can be optimized if they are available.

Templates must be suitably sysprepped. See chapter 13 of the *CloudPlatform Version 4.2 Administrator's Guide* for details on how to work with Windows templates.

Once the infrastructure is received by the service designer, it is the service designer who executes the remaining steps.

Recommended infrastructure configurations

CloudPlatform is heavily customizable to allow different workloads to make the best possible use of the infrastructure.

Note: The infrastructure administrator can make many of these customizations by editing the various global settings.

The following table shows recommendations for some specific global settings to achieve. For more information about global settings, see the *CloudPlatform Version 4.2 Administrator's Guide.*

Name	Description	Value	Notes
wait	Time in seconds to wait for control commands to return.	1800	When Storage server is responding slowly - Increasing this value helps to wait for more time
max.template.iso.size	The maximum size for a downloaded template or ISO (in GB).		This helps if the template is bigger than 50 GB
vmware.root.disk.controller	Set to **osdefault** for VMware-based clouds.		
vmware.create.full.clone	Setting this to FALSE allows you to use thin clones, which start faster and use less disk		If you have a hardware accelerated

Name	Description	Value	Notes
	space.		storage solution, or use other tools that are not compatible with thin clones, you probably do not want to disable this setting.

XenTools version setting

When using a zone with XenServer as the underlying hypervisor, you must check the **XenServer Tools 6.1+** property of all Windows templates that are created as part of the deployment. The CloudPlatform management console does not allow this to be done when the template is first created. When the template is ready, edit the properties of the template to turn on the required settings and select **XenServer Tools 6.1+** check box**.**

Disk, XenDesktop service, and network offerings and templates

- Disk offerings provide a choice of disk size and storage type.
- XenDesktop service offerings provide a choice of compute power and other choices.
- Network offerings provides a set of network sevices.
- Templates provide the base images.

Custom disk offerings

XenDesktop needs to create disks of arbitrary size, requiring a custom disk offering. If the infrastructure does not provide a custom disk offering, the service designer must create one with the web console using the procedure shown in the illustration below.

Note that you must create the custom disk offering with a **shared** storage type, but you can have additional custom offerings with other storage types. There must be at least one custom disk offering for shared storage; otherwise, XenDesktop is unable to provision machine catalogs.

XenDesktop service offerings

CloudPlatform comes with some basic compute offerings by default:

- Small Instance
- Medium Instance

For XenDesktop deployment, the service designer creates additional service offerings.

Service offerings should be given an amount of compute power and network performance that is appropriate to the different class of worker machine that will be used in the deployment. (RDS workers, for example, need more resources than single-user desktops).

Additionally, pooled desktops are intended to execute on host-attached (local) storage. This means you should create at least one service offering with the **Local** storage type by using the CloudPlatform web console.

Network offerings

For running XenDesktop workloads, networks need to support the CloudPlatform user data service. When creating networks for use in XenDesktop, always ensure they are derived from a network offering that supports this service.

Windows templates

CloudPlatform does not provide any Microsoft Windows OS templates out of the box. The infrastructure engineers must install these into the system and ensure that they contain the appropriate paravirtual tools for the hosting platform (for example, XenTools for XenServer-based clouds).

You can import templates as VHDs from, for example, XenServer. However, the best practice to get Windows running in CloudPlatform is to start by installing from the original Windows media as an ISO image file. ISOs must first be procured from Microsoft, or some suitable corporate ISO repository.

They must then be made available as downloads from a web service that is visible to CloudStack. A simple way to do this is to enable IIS on an existing Windows machine and publish the ISOs from a web directory within it.

A summary of the procedure is shown in this image.

1 *Make Windows ISO file(s) available on a local web service (eg. IIS).*

2 *Select Templates in CloudPlatform console, select ISO view, and click "+ Register ISO"*

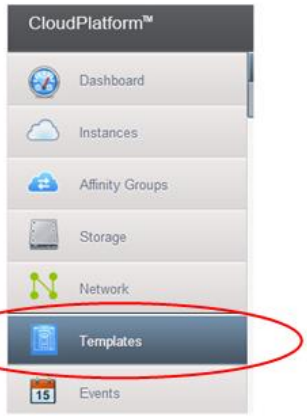

3 *Add details for the ISO, and use the URL from the local web service*

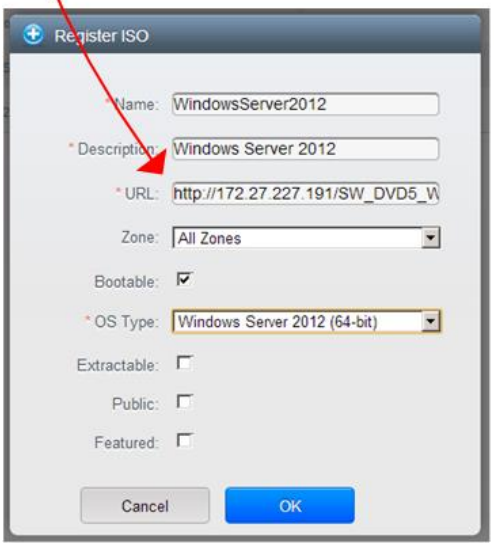

Be aware that ISO registration in CloudPlatform can appear to be a very fast process and your ISO is added to the list almost immediately. This does not mean that the template is ready. Windows ISOs take some minutes to download from your local web service into zone-level storage. To check on the true status of the template, select it to display its details. While it is downloading, the details show that it is not yet ready and the percentage of download progress. You can use the ISO only when it is fully downloaded and installed, at which point the status changes to show that the template is ready.

Volume worker template

To create machines in CloudPlatform using XenDesktop Machine Creation Services (MCS), the service designer must create a XenDesktop Volume Worker Template in CloudPlatform that performs some of the provisioning operations for the MCS on the XenDesktop controller.

When the template is created, it must be copied to all of the zones that the service designer is intending to use.

Create a volume worker template

Use the following procedures to create a volume worker template.

Part 1 – Install CentOS ISO

To begin the process, you need a CentOS 6.4 Linux installation. Do this by importing a CentOS ISO into the system in much the same way as Windows templates are imported.

1. Create an instance from the installed ISO.
2. Select the zone you are using for this installation.
3. Use any service offering with a small memory footprint. For example, the built-in **Small Instance** offering.
4. Select an appropriate network. Note: This can be any available network that provides internet access.
5. Complete the CentOS installation process.

Part 2- Enable networking

CentOS does not enable networking by default. Use **vi /etc/sysconfig/network-scripts/ifcfg-eth0** to enable networking.

1. Change **ONBOOT=no** to **ONBOOT=yes**.
2. Delete the **HWADDR** (MAC address) entry, if one is present in the file.
3. Delete the **UUID** entry, if one is present in the file.
4. Save the file.

Part 3 - Enable the interface and verify it obtained an IP address.

1. To bring the network interface online, run the **ifup eth0** command.
2. To verify that the instance has obtained an IP address, run the **ifconfig eth0** command.

Part 4 – SSL connections

1. To accept incoming SSL connections, ensure that port 443 is opened in the firewall (iptables) using **vi /etc/sysconfig/iptables** to make the following change, and then save the file
   Add the line **-A INPUT -p tcp --dport 443 -j ACCEPT** to the chain of input rules, above the line **-A INPUT -j REJECT --reject-with icmp-host-prohibited**.
2. Apply the new rules with **service iptables restart**.

Part 5 - Copy the Citrix software package onto the machine

1. If the instance requires an HTTP proxy to be defined to access external sites, set it with **export http_proxy=http://***<address>**:port.*
2. Ensure the operating system and installed packages are up to date with **yum update** and press **Y** to accept the download and installation.
3. Using WinSCP (http://winscp.net/eng/index.php) from the XenDesktop controller. Connect to the IP address listed by the ifconfig command above.
4. Navigate the local directory to **C:\Program Files\Citrix\MachineCreation\Service\VolumeWorker**.
5. Copy the **ctxvwd-*x.x*-1.i386.rpm** file to the remote centOS instance.

   Note: WinSCP is only one method of transferring the rpm package to the volume worker instance. Use any method you prefer such as mounting cifs/nfs shares or use curl to obtain it from an http server.
6. Install the Citrix Volume Worker package with **yum install ctxvwd-1.1-1.i386.rpm** and press **Y** to accept the downloads and installation.
7. Configure the Citrix service for CloudPlatform with **/etc/ctxvwd/select-platform CCP.**
8. Remove the **/etc/udev/rules.d/70-persistent-net.rules** file if present.
9. Remove the **/var/lib/dhclient/dhclient-eth0.leases** file if present.
10. Start the service with **service ctxvwd start**. This initializes the service data and then shuts down the machine.

Part 6 - Create a template from the instance's volume

When CloudPlatform indicates the instance has stopped, do the following:

1. On the **Storage** page of the CloudPlatform console, select the root volume for the instance created at the beginning of this procedure.
2. Choose **Create Templ**ate on the toolbar buttons.

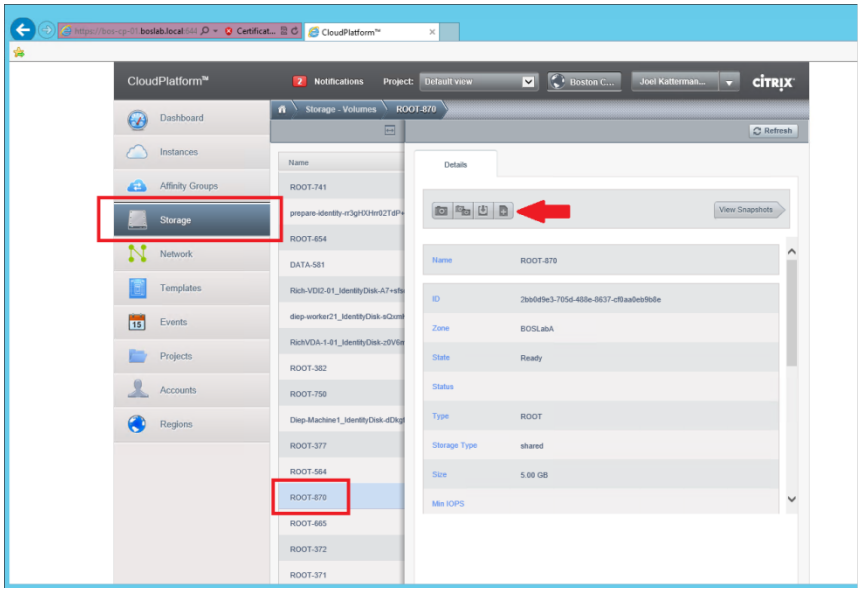

3. Provide a name and description and select CentOS 6.0 (32 bit) as the OS type, and click OK.
4. Tag the template as being a Citrix Volume Worker instance to enable XenDesktop to locate it.
5. Open the details page for the Template.
6. Scroll down to the "tags" section.
7. Add a tag using the key **Citrix.XenDesktop.Template.Role** and the value **VolumeServiceWorkerRole**.

Security updates

Once you have created a template, monitor for security patches. Also, as new versions of XenDesktop are released, it is equally important to update with the latest version of the volume worker (ctxvwd) package.

To apply a security update, create a new template containing relevant updates. If XenDesktop finds multiple templates with the same tag, it always uses the newest one, so it is not necessary to remove or untag the old templates.

Virtual Private Cloud (VPC)

If the infrastructure is being consumed as a public cloud service, the service designer needs to create a VPC within each availability zone in which the service will run. Establish VPN connections to that VPC.

Dedicated tenancy

If the cloud is private, and the zones are using basic (as opposed to advanced) networking, the service designer must create a suitable set of security groups, with appropriate ports open both for infrastructural instances and for worker (VDA) instances.

Customers can apply their own configuration to change the port numbers that are used by XenDesktop. However, in the absence of any such special requirements, service designers should create security groups with the following ingress rules:

- TCP port 80 (for HTTP/WCF)
- TCP port 443 (for SSL/HTTP(s))
- TCP port 2598 (for ICA)

Active Directory

Windows Active Directory is part of the service designer's deployment.

Supported patterns for Active Directory deployment fall within the following categories:

Active Directory fully on-premise – No part of Active Directory is running within the cloud. All cloud-hosted instances (for example, XenDesktop controllers and workers) can see the domain controller using a shared network. This means that the zone must either be a basic zone or an advanced zone with a shared network offering.

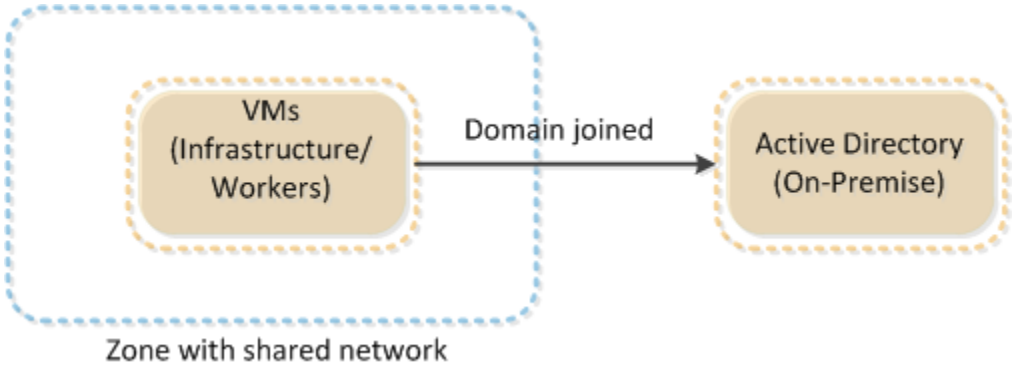

- **Active Directory on premise, with a secondary domain controller in the cloud** – The cloud is running within an isolated network but with a site-to-site VPN established.

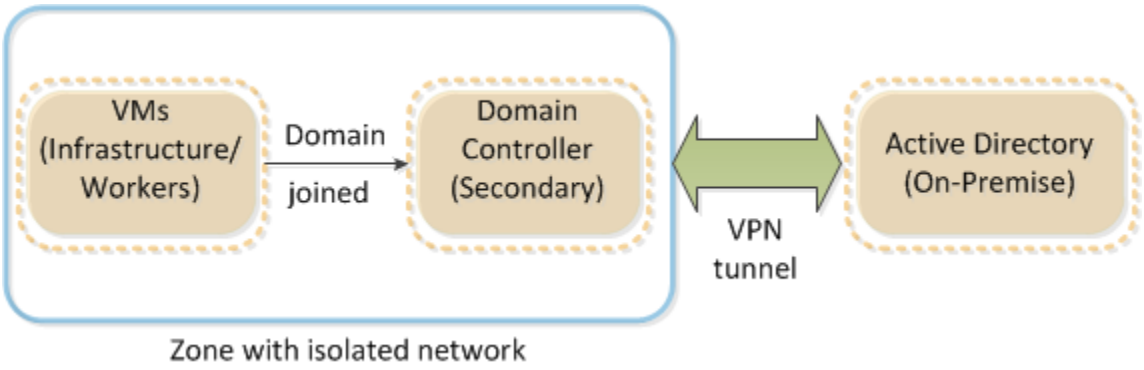

- **Active Directory fully within the cloud** – Typically, this scenario is restricted to small Proof of Concept or demonstration deployments.

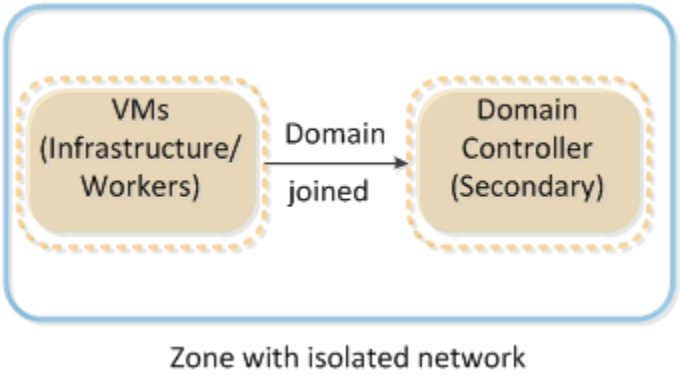

There are two key logistical issues surrounding the deployment of Active Directory:

- DHCP addresses are assigned by the cloud. Traditionally in a Windows domain, the domain controller is both a DHCP and DNS server. In CloudPlatform, DHCP cannot come from the Windows domain, because it is handled by CloudPlatform internally to ensure that IP address assignments follow the IP range restrictions established for the zone. This means that DHCP must be disabled on any Windows domain controller that is situated within the cloud.
- Domain controllers need fixed IP addresses. Any domain controller (whether primary or secondary) running within the cloud must have a fixed IP address. Note that fixed does not mean static. All cloud instances have DHCP addresses assigned by the local virtual router, so they are not static. However, it is possible to launch a cloud instance such that it is always assigned the same well-known address, making it effectively static. (Similar facilities are available in Amazon Web Services.) Be aware that in CloudPlatform the web console does not support the creation of instances with prescribed IP addresses (unlike the Amazon Web Services console). Such IP assignments are possible only through the REST API.

Service designers must make API calls to create domain controllers (primary or secondary) within the cloud.

Part two – Deployment

Part two describes what you need to run XenDesktop sites after the CloudPlatform installation is complete. Refer to the *Part one – Overview of concepts and prerequisite infrastructure* if you require basic CloudPlatform and conceptual information before starting the procedures.

A minimal XenDesktop site can run on an installation containing just a single host in a single availability zone. Zones can use either basic or advanced mode networking. Public clouds use advanced networking to allow tenants to use VPCs and isolated networks. Private clouds can use either mode.

Requirements

Ensure you are using the supported platforms and have the prerequisite infrastructure set up before deploying the XenDesktop site on the CloudPlatform. Refer to *Prerequisite infrastructure*.

CloudPlatform and hypervisor versions

For both public and private clouds, XenDesktop workloads can run only on infrastructure that has fully adopted a supported CloudPlatform version. Do not use older versions to host XenDesktop workloads. The CloudPlatform installation must use either the supported XenServer or VMWare hypervisor.

CloudPlatform	XenServer	VMWare
CloudPlatform 4.2.1-4	6.2, 6.2SP1	5.1
CloudPlatform 4.3.0.1	6.2SP1, 6.2SP1 + XS62ESP1004	5.1, 5.5

Launch a Delivery Controller

The infrastructure engineer has made the appropriate Windows Server (2008 or 2012) templates available, with the necessary hypervisor tools installed. To launch a Delivery Controller, create a new instance from that template and join it to the correct Active Directory domain.

If the zone network supports security groups, launch the controller with the security group that supports XenDesktop.

Install XenDesktop

When the necessary Windows instances are running and joined to the appropriate Active Directory domain, install XenDesktop from the media. The procedure is similar to Amazon Web Services and the complexity of it depends on whether the cloud is running in an isolated network or not.

If you use shared networking, you can use a simple Windows copy operation to import the XenDesktop.ISO image file onto the controller and Virtual Desktop Agent (VDA). In the case of Windows 8 or Windows Server 2012, you can then mount and install the ISO. For older operating systems, you need a third-party tool to mount the ISO as a CD-ROM. As on Amazon Web Services, a facility such as MagicISO achieves this.

When the XenDesktop media are mounted, you need no more special actions. Install as normal.

For the VDA installation, assume Machine Creation Services (MCS) style provisioning and choose the option for creating a master image. When asked how to locate the Delivery Controller, select the **Let Machine Creation Services do it automatically** option.

Example: Go from bare CloudPlatform to a XenDesktop site

The visual sequence in the three steps below indicates how you can transform raw infrastructure into a working XenDesktop site. Assume the site is a proof-of-concept site, so is quite minimal. From original ISO images for Windows and Linux, instantiate a single Delivery Controller instance, and a single VDA. After installing XenDesktop worker components on the VDA, shut it down and create a template. Machine provisioning then takes place from that template.

1. Prepare initial OS templates and Citrix Volume Worker

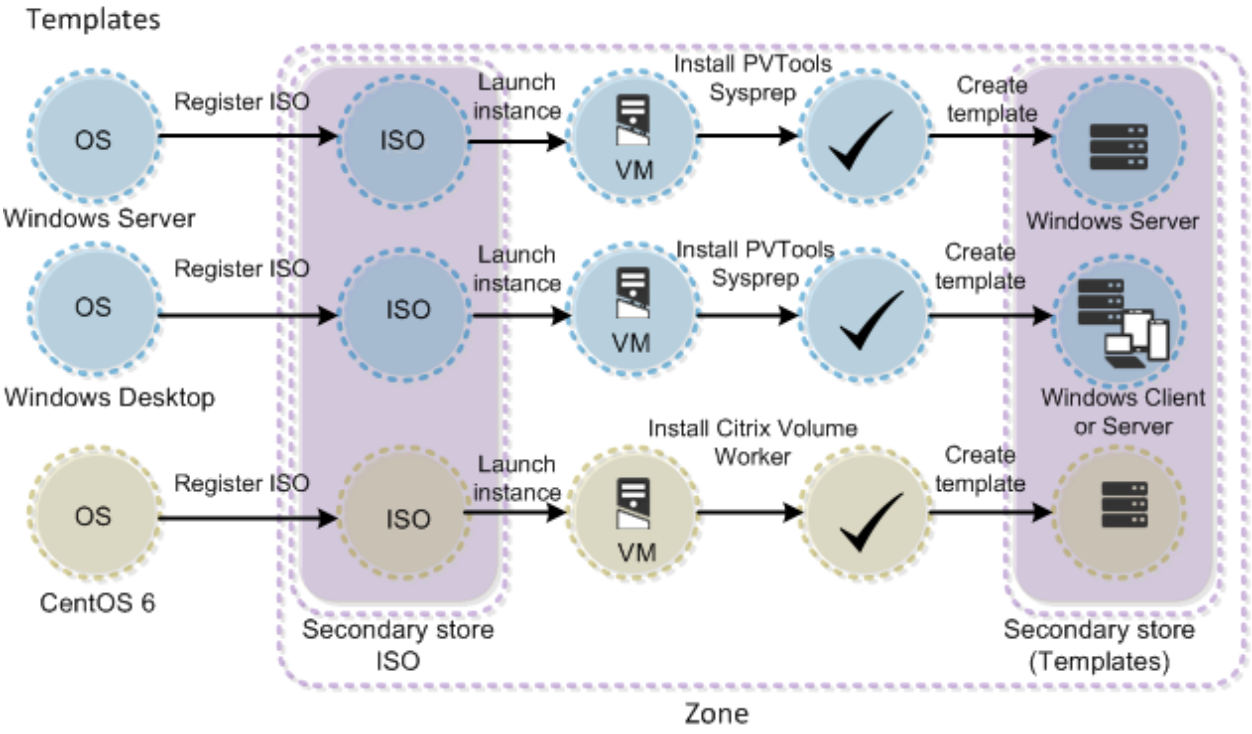

2. Install Controller and create VDA template

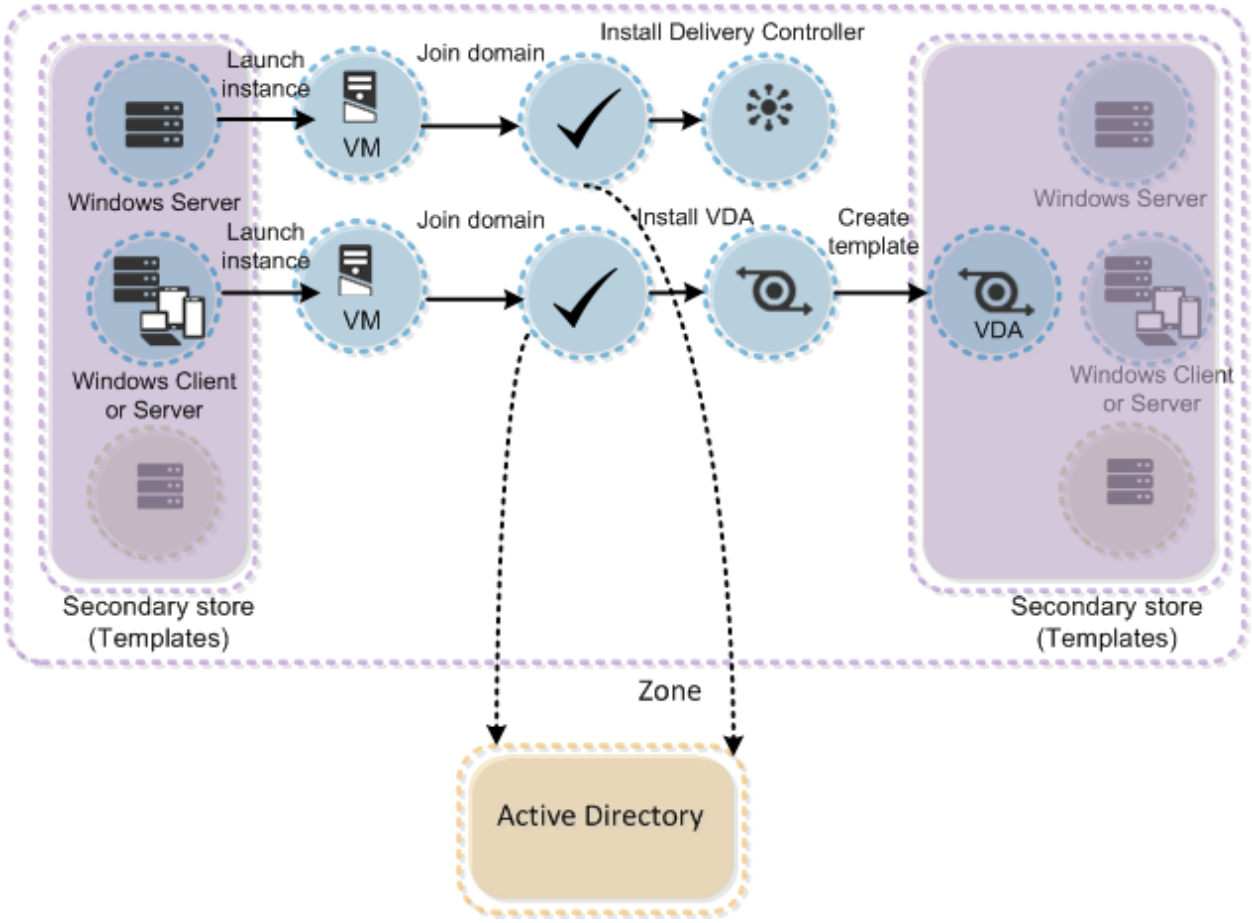

3. Create site and deploy desktops for users

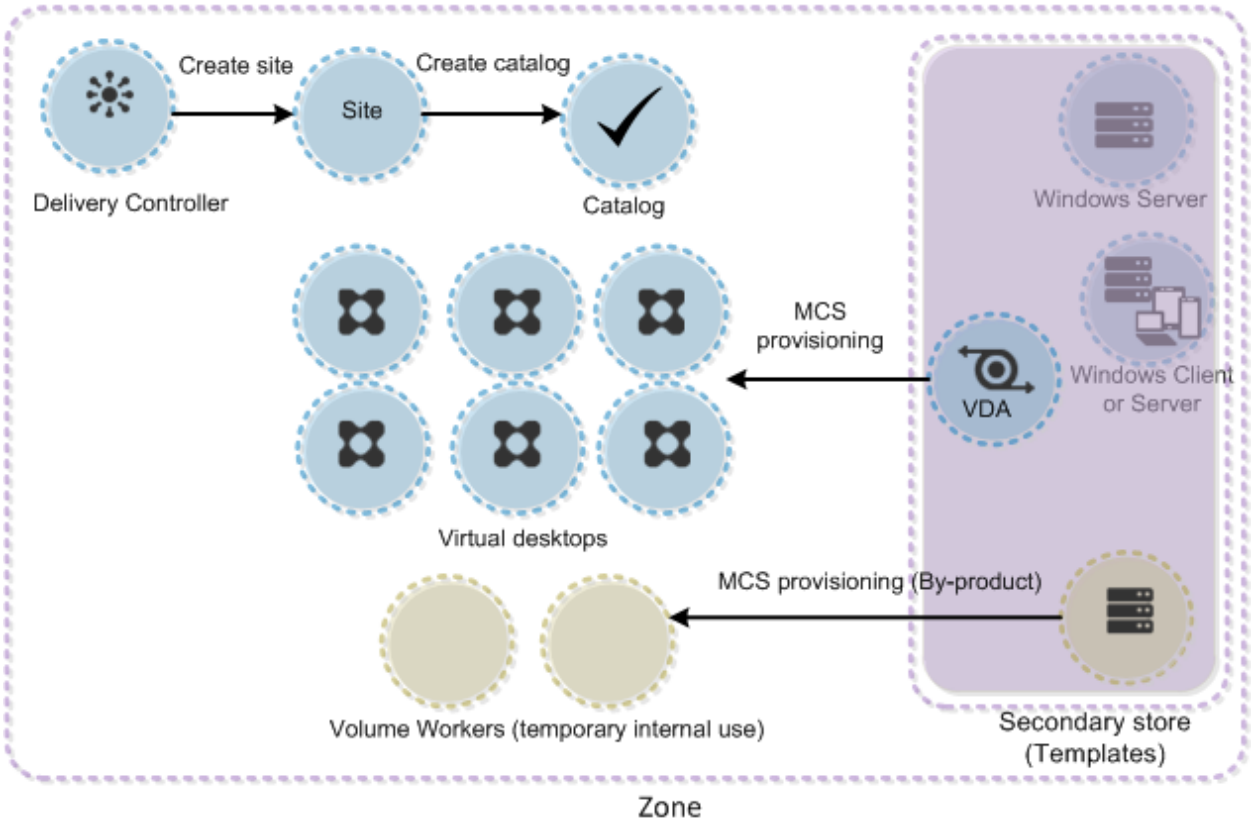

Image preparation isolation

In Amazon Web Services and CloudPlatform basic zones, the Machine Creation Services (MCS) catalog's image preparation VM is wrapped in a security group to isolate it from the network.

In CloudPlatform advanced zones, the image preparation VM is created on an isolated network, as CloudPlatform does not support security groups in advanced zones. The administrator must perform manual configuration to support MCS image preparation in advanced zones:

- Create an isolation network in the CloudPlatform account used by XenDesktop. Citrix recommends:
  - `DefaultIsolatedNetworkOfferingWithSourceNatService`
- Apply a tag to the network so XenDesktop can find it. Case sensitive tag values:
  - **Tag:** `Citrix.XenDesktop.Network.Role`
  - **Value:** `MachineIsolationRole`

When creating the preparation VM, XenDesktop identifies a network with these tags and supplies this information when creating the VM. If the tag is not found, no networks are supplied

when creating the VM and CloudPlatform attempts to find an existing, or create a new, isolated network using the same offering as above. This process can fail, for example, if two isolation networks exist because it is not clear which network CloudPlatform should select.

For debugging purposes, you can use one of the following HCL traces to determine whether the network was discovered:

```
[CloudStack]: PluginMachineManager.EstablishIsolationEnvironment: The network 'Tier0(192.168.30.0/24)' (2e32d848-f32a-4fae-8c02-baa5602c9c1e) will be used for isolation.
```

```
[CloudStack]: PluginMachineManager.EstablishIsolationEnvironment: No explicit isolation network found. Using an empty network array.
```

Example: Build a simple site with Studio

If you install the XenDesktop.ISO on the Delivery Controller and launch, the following sequence deploys a simple site. This example assumes:

- A very straightforward cloud setup with a single zone and a single region (and no VPC).
- The default security group has been edited to allow ingress on the ports that XenDesktop requires.

When you launch Studio before a site has been created, it automatically detects that no site exists. It then provides a 1-2-3 workflow, where the first step is to define the site and hosting resources, the second is to create a catalog (a set of virtual machines cloned from a template), and the third is to create a delivery group, which defines how users are allowed to access their desktops.

When you have created a delivery group, you can log on to a desktop using StoreFront.

1 - Define basic site parameters

1. Start Studio and select Get started! Create a Site.
2. Choose **A fully configured, production Site** and name the site **MyCloud** (The name can be anything you like.)
3. With the site name defined, XenDesktop needs to connect to the database that will store all persistent state information associated with that site. Unless you have Microsoft SQL Server pre-installed on the controller, SQLEXPRESS is installed by default, so the database server is **.\SQLEXPRESS**.
4. The default name of the database is your chosen site name, with **Citrix** prepended. In this example procedure, the default database name is **CitrixMyCloud**.
5. Because the database does not exist, Studio displays a message. Click **OK** and Studio creates the database automatically.

6. Add product licensing for XenDesktop. Because you do not have a license in this example, choose the 30-day trial. A license server runs by default on the local host on port 27000. XenDesktop detects this automatically and connects to it. Click **Next**.

Describe a connection to the hosting infrastructure

The hosting infrastructure is the hypervisor onto which MCS will deploy machines.

1. On the Studio Connection screen, select **Citrix CloudPlatform** as the host type, and type the URL to the client API endpoint in the **Address** text box. Note you must include the **/api** suffix on the URL, because XenDesktop talks to the REST API, not to the web console.
2. Copy and paste the API key and secret key from the CloudPlatform management console.
3. Give the connection a name. This can be anything you like and click **Next**.

Specify VM location and network

1. Use the Studio VM Location screen to specify the Machine Creation Services (MCS) to create machines within a chosen region and zone. This example is a single-region, single-zone cloud, so accept the defaults on this page.
2. Use the Studio Network screen to choose a network onto which the provisioned desktops will be linked. This example uses the zone's default guest network. There is no VPC.
3. This completes the definition of our hosting resources. Give the resources a name. The name can be anything you like.
4. In this example, skip the Studio App-V Publishing screen. There’s nothing CloudPlatform-specific on the screen.
5. To create the site, click **Finish**. This might take several minutes before the next screen displays.

2 - Create a catalog of virtual desktops

MCS creates a catalog of virtual desktops using the specifications in this example procedure. Under the Studio Full Deployment tab, choose **Create Catalog** to create VMs within the CloudPlatform zone.

For this procedure, assume you installed the XenDesktop VDA on a Windows 7 or Windows 8 machine (rather than Windows Server).

You will create a *pooled random* catalog, because it is the simplest type. A pooled random catalog comprises machines that reset to their original state after every reboot and do not save any user changes, neither on the root disk nor on any separate data disk attached to the

machine. Because there is no persistent storage on the VM, any user can attach to any desktop from the pool.

Every virtual desktop must join an Active Directory domain and needs a computer account within that domain. In this example procedure, only one Active Directory domain exists and the Delivery Controller is already joined to that domain.

1. On the Studio Operating System and Hardware screen, select **Windows Desktop OS** as the operating system.
2. On the Studio Machine Management screen, accept the defaults to create new virtual machines with MCS.
3. On the Studio Desktop Experience screen, select the random desktop option.
4. You created a template in the CloudPlatform secondary zone as part of the service design process. On the Studio Machine Template screen, select the template that you created.
5. Because we have a basic networking zone, we will have security groups. On the Studio Security screen, select the default security group to open the required ports for XenDesktop only if you edited it with the CloudPlatform web console.
6. Select Use any available hardware.
7. On the Studio Virtual Machines screen, for this example procedure, create three VMs and select **Medium Instance**. If you have a more specific service offering, use that instead.
8. On the Studio Network Interface Card Configuration screen, select **defaultGuestNetwork**.
9. On the Studio Active Directory Computer Accounts screen, accept the defaults, unless you want to use a specific OU for the computer accounts. Define a naming scheme. A scheme string of HarveyA## creates accounts with names HarveyA01, HarveyA02, HarveyA03, and so on.
10. Studio displays a summary of your choices on the Studio Summary screen. Define the name and description for the catalog.

MCS now creates the machines. This can take can take many minutes. The actual time depends on the performance of the storage within your CloudPlatform zone. Rather than watch the progress bar within Studio, you can look directly at the CloudPlatform web console and see machines gradually appearing. MCS creates some temporary VMs as part of the process. You should end up with the three virtual desktops that you specified in these steps. Their names are the same as the Active Directory account names (in our example, HarveyA01, HarveyA02, and HarveyA03).

3 - Create a Delivery Group

After the machines are created, create a delivery group, which makes it possible to log onto a desktop. You have only one catalog in the example, so that is selected by default. You can choose to put any number of machines from that catalog into the new delivery group. In the example below, add all three machines to the group. Because the catalog is pooled-random, three machines can support three users accessing their desktops concurrently. Any number of users can use the site, as long as it is at different times.

1. Under the Studio Full Deployment tab, choose **Create Delivery Group, Applications, and Assign Users** to get started creating a Delivery Group.
2. Your one catalog is chosen by default. Add three machines.
3. On the Delivery Type screen, choose the default (Desktops).
4. On the Users screen, specify which users are allowed to log on. Click **Add users...** to define a set of users within Active Directory.
5. On the Select Users or Groups dialog box, type **Domain Users** and click **Check Names** to ensure Windows recognizes it as a valid group. This allows all domain users to access the desktops.
6. To keep things simple in the example, do not define any StoreFront addresses for the desktops. Click **Next** to move to the Summary screen.
7. Specify the Delivery Group name and the display name. The display name displays within the StoreFront logon screen when users are choosing which desktop they want to use.

The delivery group is created more quickly than the catalog and should take just a few seconds. If you refer to the CloudPlatform web console at this point, you should find that one or more of the MCS virtual desktops gets powered on. This occurs because the XenDesktop broker is now aware that the machines exist and is getting ready for users to log on. It powers some machines on in advance so that users can connect quickly.

Now that the delivery group is created, you can connect to a desktop. The easiest way to do this is to use a web browser to log on through the StoreFront cluster, which is running on the Delivery Controller by default. Point your browser at **http://***x.x.x.x***/Citrix/StoreWeb**, where *x.x.x.x* is the IP address of your Delivery Controller. The StoreFront logon screen displays. Supply some appropriate credentials for your Active Directory domain and your desktop appears.