Citrix XenApp and Desktop (Studio) Questions & Answers

What is FMA? In addition, mention the major differences between IMA and FMA Architecture. From Which Version, FMA Introduced?

Flex Management Architecture may have several definitions from both Marketing and Technical, My understanding FMA is , it is service oriented architecture where all services works independently with each other and XenApp and XenDesktop share a unified architecture called FlexCast Management Architecture (FMA).

IMA back then…	FMA as it is today…
IMA - Independent Management Architecture.	FMA - Flexcast Management Architecture
Farm.	Site
Worker Group.	Machine Catalog / Delivery Group.
Worker / Session Host / XenApp server.	Virtual Delivery Agent (VDA). There is a desktop OS VDA as well as a server OS VDA, including Linux.
Data Collector (one per zone).	Delivery Controller (multiple per Site).
Zones.	Zones (as of version7.7)
Local Host Cache (LHC).	Connection Leasing
Delivery Services Console / App center.	Citrix Studio (including StoreFront) and Director.
EdgeSight monitoring (optional).	Partly built into Director
IMA Data store.	Central Site database (SQL only).
Load evaluators.	Load management policies.
IMA protocol and service.	Virtual Delivery Agents / TCP.
Smart Auditor.	Session Recording.
Shadowing users.	Microsoft Remote Assistance, launched from Director.

Until XenApp 6.5, it is IMA and from XenDesktop 7.X, the architecture changed

Explain any few FMA Services.

FMA is a service-oriented architecture, all services are independent of each other and each Service runs complete separated from the other services, as a result, each service also has its own separate database connection string: if one service fails, it will not directly affect any of other services.

Although FMA services run completely isolated (independent) from each other, internal communications between the different services takes place using so-called WCF (Windows Communication Foundation) over port 80

Over the years, it has evolved from **6 services to 18 services in total.**

FMA Services	***Description***
Broker Service	*Brokers new session requests, handles disconnected sessions and resource enumeration, processes STA ticket verification and user validation. Additionally, it handles all communication to and from the VDA desktop.*
Machine Creation Service	*Handles the creation of new virtual machines (not physical machines).*
Configuration Service	*Handles all inter-service communication between FMA services.*
AD Identity Service	*Handles all Active Directory accounts related to any XenDesktop virtual or physical machines.*
Hosting Service	*Manages all connections between the physical hosts, the Delivery Controllers and the underlying Hypervisor (s)*
Delegated Administration Service	*It is one of critical service, All other FMA services will need to communicate with the Delegated Administration Service in order to validate if they have all the proper permissions and/or rights needed to make the necessary changes to the Central Site Database.Manages the creation, configuration and administration of all delegated administrative permissions.*
Monitoring Service	*Monitors the overall FMA architecture and produces alerts and warnings when it finds something is potentially wrong, such as a failing service.*

Environment Test Service	*Takes care of all Site-wide tests, initiated from Studio. You can run tests on your Delivery Groups, Machine Catalogs or even on your entire Site configuratio*
Configuration Logging Service	*Monitors and logs all configuration changes made within a XenDesktop site, to include all administrator activity.*
Analytics Service	*Collects analytical data from Citrix products.*
StoreFront Service	*Manages the StoreFront deployment from studio.*
High Availability Services -*New from 7.12*	Called as Secondary Broker Services which is responsible for connection Leasing when Primary Broker service is stopped
Config Synchronizer service(CSS) -*New from 7.12*	Every two minutes the Principal Broker Service will be checked for configuration changes,If a configuration change has been detected it will be copied over, or synchronised to the High Availability Service/Secondary Broker Service through CSS service
Configuration Service -*New from 7.12*	Located at the centre of the FMA, it holds and manages a list of all FMA service and all FMA services need to register with the configuration Service on start-up and all other FMA services need to communicate via this service
Analytics -*New from 7.12*	As the name implies: it collects analytical data used by Director/Studio (custom reports, to name one)

Addition of below services with Cloud-> 7.13 -total in 15

Citrix App Library
Citrix Orchestration Services
Citrix Remote Broker Provider
Citrix Storefront Privileged Administration Service
Citrix Telemetry Service

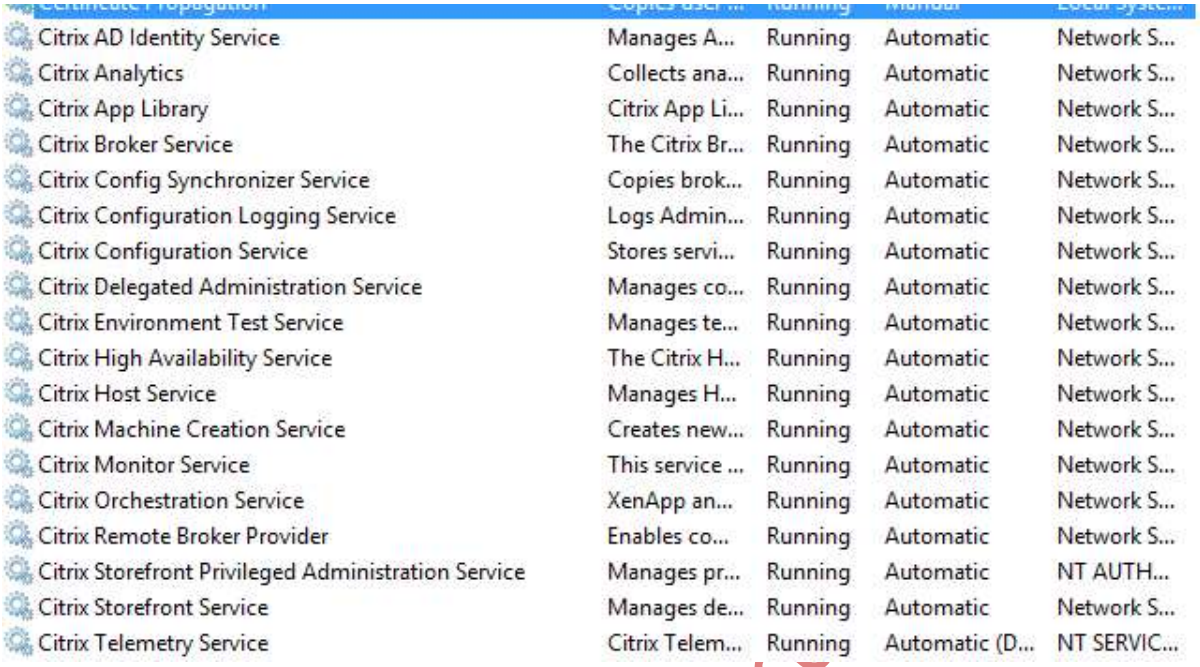

Citrix AD Identity Service	Manages A...	Running	Automatic	Network S...
Citrix Analytics	Collects ana...	Running	Automatic	Network S...
Citrix App Library	Citrix App Li...	Running	Automatic	Network S...
Citrix Broker Service	The Citrix Br...	Running	Automatic	Network S...
Citrix Config Synchronizer Service	Copies brok...	Running	Automatic	Network S...
Citrix Configuration Logging Service	Logs Admin...	Running	Automatic	Network S...
Citrix Configuration Service	Stores servi...	Running	Automatic	Network S...
Citrix Delegated Administration Service	Manages co...	Running	Automatic	Network S...
Citrix Environment Test Service	Manages te...	Running	Automatic	Network S...
Citrix High Availability Service	The Citrix H...	Running	Automatic	Network S...
Citrix Host Service	Manages H...	Running	Automatic	Network S...
Citrix Machine Creation Service	Creates new...	Running	Automatic	Network S...
Citrix Monitor Service	This service ...	Running	Automatic	Network S...
Citrix Orchestration Service	XenApp an...	Running	Automatic	Network S...
Citrix Remote Broker Provider	Enables co...	Running	Automatic	Network S...
Citrix Storefront Privileged Administration Service	Manages pr...	Running	Automatic	NT AUTH...
Citrix Storefront Service	Manages de...	Running	Automatic	Network S...
Citrix Telemetry Service	Citrix Telem...	Running	Automatic (D...	NT SERVIC...

Ref:

http://www.basvankaam.com/2016/12/15/the-citrix-xenapp-xendesktop-fma-services-complete-overview-new-7-12-services-included/

Is XenDesktop works without database connectivity? What is the immediate impact for existing sessions if database goes down?

Yes, Xendesktop work without database from 7.6 version by utilizing the feature connection leasing. No Immediate impact for existing sessions

What is connection leasing? Difference between LHC before 7.x and Connection leasing? From which version it is available.

The connection-leasing feature increases the SQL Server high availability by enabling users to connect and reconnect to their most recently used applications and desktops, even when the Site database is not available

Connection –leasing is enabled by default and introduced from the **version 7.6**, each Controller caches user connections in the folder C:\Program Data\Citrix\Broker\Cache when launching for first time and

periodically (10sec) synched to central site database. Connections are cached for a lease period of **two weeks**. So, if the database becomes unavailable, the desktops and applications that the user launched in the previous two weeks remain accessible.

When the Controller is in leased connection mode:

- Administrators cannot use Studio, Director, or the PowerShell console.
- Workspace Control is not available. When a user logs on to Receiver, sessions do not automatically reconnect; the user must relaunch the application.
- Static (assigned) desktops are not power-managed. VDAs that are powered off when the Controller enters leased connection mode remain unavailable until the database connection is restored, unless the administrator manually powers them on
- User cannot launch Pooled Desktops whereas static desktops are available
- Citrix Power Management is not available until the site database connection is restored.

Connection Leasing is not replacement of LHC and it is an alternative option provided in placement of LHC in FMA architecture

By default, a XenApp server (Data Collectors included) polls the central IMA store, through the local IMA service, every 30 minutes, the information obtained is stored into a local MDB database, which is referred to as the Local Host Cache (LHC). Also, when configuration changes are made within the Farm, the Zone Data Collectors will be notified so that they can update their LHC, next, the Data Collectors will notify their Zone member servers so they can do the same. The Local Host Cache stores the following information:

All this ensures that when the IMA Store for whatever reasons isn't reachable users can continue to work, logon, logoff etc. in fact, if needed the server can be rebooted while the IMA store is down and the local IMA service will start from the LHC without any issues

Ref:

https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-6/xad-connection-leasing.html

Lease Connection FAQ: https://support.citrix.com/article/CTX205169

http://www.basvankaam.com/2014/10/06/citrix-connection-leasing-vs-local-host-cache-cl-doesnt-stand-a-chance/

LHC reintroduced in XenDesktop 7.12? How it is different with Connection leasing? Any idea on this.

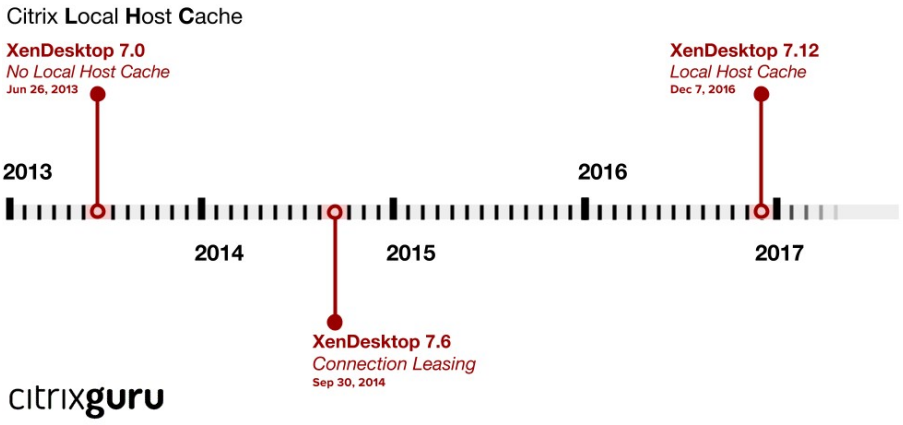

Citrix came up with a milestone achievement with its new idea as part of the XenDesktop 7.12. This time, they claimed to bring back all the Local Host Cache (LHC) features from XenApp 6.5, even adding few improvements to make it more reliable. LHC feature is offered for Cloud and On Premises implementations along Connection Leasing in 7.12, but is considered the primary mechanism to allow connection-brokering operations when database connectivity to the site database is disrupted. Surprisingly, Local Host Cache feature is disabled by default. We can expect Citrix to enable that feature by default in the next version.

The Local Host Cache (LHC) feature allows connection-brokering operations in a XenApp or XenDesktop Site to continue when an outage occurs. An outage occurs when:

- The connection between a Delivery Controller and the Site database fails in an on-premises Citrix environment.
- The WAN link between the Site and the Citrix control plane fails in a Citrix Cloud environment

Local Host Cache is the most comprehensive high availability feature in XenApp and XenDesktop. It is a more powerful alternative to the connection-leasing feature that was introduced in XenApp 7.6

When installing XenDesktop 7.12 and up, a SQL Express instance(LocalDB) will be installed locally (C:\Windows\ServiceProfiles\NetworkService\HaDatabaseName.mdf) on each Delivery Controller to store the Local Host Cache. For those familiar with XenDesktop 7, the Broker Service is still available and referenced as the Principal Broker and still manages VDA registrations and brokering during normal operations. This service also checks on the remote database to make sure that it is online.

The Config Synchronizer Service (CSS) and the Secondary Brokering Service (Citrix High Availability Service) are now part of the installation. CSS takes care of the synchronization between the remote database and the Local Host Cache (LocalDB). The Secondary Brokering Service takes over from the Principal Broker when an outage is detected and does all registration and brokering operations

How it works

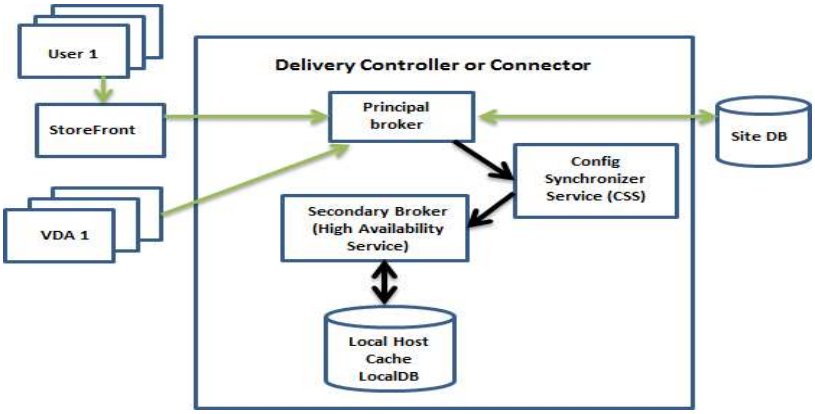

During normal operations:

- The *principal broker* (Citrix Broker Service) on a Controller accepts connection requests from StoreFront, and communicates with the Site database to connect users with VDAs that are registered with the Controller.
- A check is made every two minutes to determine whether changes have been made to the principal broker's configuration. Those changes could have been initiated by PowerShell/Studio actions (such as changing a Delivery Group property) or system actions (such as machine assignments).
- If a change has been made since the last check, the principal broker uses the Citrix Config Synchronizer Service (CSS) to synchronize (copy) information to a *secondary broker* (Citrix High Availability Service) on the Controller. All broker configuration data is copied, not just items that have changed since the previous check. The secondary broker imports the data into a Microsoft SQL Server Express LocalDB database on the Controller. The CSS ensures that the information in the secondary broker's LocalDB database matches the information in the Site database. The LocalDB database is re-created each time synchronization occurs.
- If no changes have occurred since the last check, no data is copied.

When an outage begins:

- The principal broker can no longer communicate with the Site database, and stops listening for StoreFront and VDA information .The principal broker then instructs the secondary broker (High Availability Service) to start listening for and processing connection requests.
- When the outage begins, the secondary broker has no current VDA registration data, but as soon as a VDA communicates with it, a re-registration process is triggered. During that process, the secondary broker also gets current session information about that VDA.
- While the secondary broker is handling connections, the principal broker continues to monitor the connection to the Site database. When the connection is restored, the principal broker

instructs the secondary broker to stop listening for connection information, and the principal broker resumes brokering operations. The next time a VDA communicates with the principal broker, a re-registration process is triggered. The secondary broker removes any remaining VDA registrations from the previous outage, and resumes updating the LocalDB database with configuration changes received from the CSS.

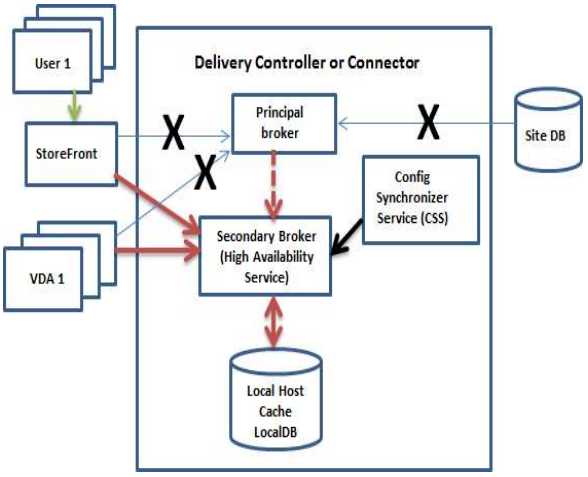

What is unavailable or changes during an outage

- You cannot use Studio or run PowerShell cmdlets. No Site Management
- Hypervisor credentials cannot be obtained from the Host Service. All machines are in the unknown power state, and no power operations can be issued. However, VMs on the host that are powered-on can be used for connection requests.
- Machines with VDAs in pooled Delivery Groups that are configured with "Shut down after use" are placed into maintenance mode.
- An assigned machine can be used only if the assignment occurred during normal operations. New assignments cannot be made during an outage.
- Automatic enrollment and configuration of Remote PC Access machines is not possible. However, machines that were enrolled and configured during normal operation are usable.
- Server-hosted applications and desktop users may use more sessions than their configured session limits, if the resources are in different zones

Highlights

- Local Host Cache supports more use cases than connection leasing.
- When operational, Local Host Cache requires more resources (CPU and memory) than connection leasing.

- During outage mode, only a single broker per zone will handle VDA registrations and broker sessions.
- An election process decides which broker will be active during outage, but does not take into account broker resources.
- If any single broker in a zone would not be capable of handling all logons during normal operation, it won't work well in outage mode.
- No site management is available during outage mode.
- A highly available SQL Server is still the recommended design.
- For intermittent database connectivity scenarios, it is still better to isolate the SQL Server and leave the site in outage mode until all underlying issues are fixed.
- There is a limit of 5,000 VDAs per zone (not enforced).
- There is **no 14-day limit.**
- **Pooled desktops are not supported** in outage mode, in the default configuration.

PowerShell Commands

Get-BrokerSite **– To know status of Cache**

Set-BrokerSite -LocalHostCacheEnabled $true -ConnectionLeasingEnabled $false **– To enable Local Host Cache**

Ref:

Local Host Cache sizing and scaling - http://docs.citrix.com/en-us/categories/solution_content/implementation_guides/local-host-cache-sizing-scaling.html

https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-12/manage-deployment/local-host-cache.html

http://www.citrixguru.com/2017/05/02/never-ending-citrix-local-host-cache-story/

http://www.basvankaam.com/2016/12/05/the-long-awaited-xenapp-and-xendesktop-7-12-local-host-cache/

http://www.citrixguru.com/2017/05/02/never-ending-citrix-local-host-cache-story/ -Good

Define FMA key component's in short.

Delivery Controller

The Delivery Controller is the central management component of a XenApp or XenDesktop Site. Each Site has one or more Delivery Controllers. It is installed on at least one server in the data center. For Site reliability and availability, Controllers should be installed on more than one server. If your deployment includes virtual machines hosted on a hypervisor or cloud service, the Controller services communicate with the hypervisor to distribute applications and desktops, authenticate and manage user access,

broker connections between users and their virtual desktops and applications, optimize use connections, and load-balance these connections.

The Controller's Broker Service tracks which users are logged on and where, what session resources the users have, and if users need to reconnect to existing applications. The Broker Service executes PowerShell cmdlets and communicates with a broker agent on the VDAs over TCP port 80. It does not have the option to use TCP port 443.

The Monitor Service collects historical data and places it in the Monitor database. This service uses TCP port 80 or 443.

Data from the Controller services is stored in the Site database.

The Controller manages the state of desktops, starting and stopping them based on demand and administrative configuration. In some editions, the Controller allows you to install Profile management to manage user personalization settings in virtualized or physical Windows environments

Site Database

A Microsoft SQL database that stores data for the Delivery Controller, such as site policies, machine catalogs, and delivery groups.

Virtual Delivery Agent (VDA)

The VDA is installed on each physical or virtual machine in your Site that you make available to users; those machines can deliver applications or desktops. The VDA enables the machine to register with the Controller, which in turn allows the machine and the resources it is hosting to be made available to users. VDAs establish and manage the connection between the machine and the user device, verify that a Citrix license is available for the user or session, and apply whatever policies have been configured for the session.

The VDA communicates session information to the Broker Service in the Controller through the broker agent included in the VDA. The broker agent hosts multiple plugins and collects real-time data. It communicates with the Controller over TCP port 80. It does not have the option to use TCP port 443.

The word "VDA" is often used to refer to the agent as well as the machine on which it is installed.

VDAs are available for Windows server and desktop operating systems. VDAs for Windows server operating systems allow multiple users to connect to the server at one time. VDAs for Windows desktop operating systems allow only one user to connect to the desktop at a time. A Linux VDA is also available

StoreFront

The interface that authenticates users, manages applications and desktops, and hosts the application store. StoreFront communicates with the Delivery Controller using XM

It also keeps track of users' application subscriptions, shortcut names, and other data to ensure they have a consistent experience across multiple devices

Citrix Receiver

A software client that is installed on the user device, supplies the connection to the virtual machine via TCP port 80 or 443, and communicates with StoreFront using the StoreFront Service API

Studio

A management console that allows administers to configure and manage Sites, and gives access to real-time data from the Broker agent. Studio communicates with the Controller on TCP port 80.

Director

A web-based tool that allows administers access to real-time data from the Broker agent, historical data from the Site database, and HDX data from NetScaler for troubleshooting and support. Director communicates with the Controller on TCP port 80 or 443

License server

License server manages your product licenses. It communicates with the Controller to manage licensing for each user's session and with Studio to allocate license files. You must create at least one license server to store and manage your license files.

These additional components, not shown in the illustration above, may also be included in typical XenApp or XenDesktop deployments:

- **Provisioning Services** — Provisioning Services is an optional component of XenApp and XenDesktop available with some editions. It provides an alternative to MCS for provisioning virtual machines. Whereas MCS creates copies of a master image, Provisioning Services streams the master image to user device. When Provisioning Services is included in a Site, it communicates with the Controller to provide users with resources.
- **NetScaler Gateway** — A data-access solution that provides secure access inside or outside the LAN's firewall with additional credentials. When users connect from outside the corporate firewall, this release can use Citrix NetScaler Gateway (formerly Access Gateway) technology to secure these connections with SSL. NetScaler Gateway or NetScaler VPX virtual appliance is an

SSL VPN appliance that is deployed in the demilitarized zone (DMZ) to provide a single secure point of access through the corporate firewall.

- **Citrix Cloud Bridge** — In deployments where virtual desktops are delivered to users at remote locations such as branch offices, Citrix CloudBridge (formerly Citrix Branch Repeater or WANScaler) technology can be employed to optimize performance. Repeaters accelerate performance across wide-area networks.
- **NetScaler SD-WAN** - In deployments where virtual desktops are delivered to users at remote locations such as branch offices, Citrix NetScaler SD-WAN (formerly Citrix CloudBridge, Branch Repeater, or WANScaler) technology can be employed to optimize performance.

With XenApp and XenDesktop, you set up the resources you want to provide to users with machine catalogs, but you designate which users have access to these resources with Delivery Groups.

Machine Catalogs

Machine Catalogs are collections of virtual or physical machines that you manage as a single entity. These machines, and the application or virtual desktops on them, are the resources you provide to your users. All the machines in a catalog have the same operating system and the same VDA installed. They also have the same applications or virtual desktops.

Typically, you create a master image and use it to create identical VMs in the catalog. For VMs you can specify the provisioning method for the machines in that catalog: Citrix tools (PVS or MCS) or other tools. Alternatively, you can use your own existing images. In that case, you must manage target devices on an individual basis or collectively using third-party electronic software distribution (ESD) tools.

Valid machine types are:

- **Server OS machines**: Virtual or physical machines based on a server operating system used for delivering XenApp published apps, also known as server-based hosted applications, and XenApp published desktops, also known as server-hosted desktops. These machines allow multiple users to connect to them at one time.
- **Desktop OS machines**: Virtual or physical machines based on a desktop operating system used for delivering VDI desktops (desktops running desktop operating systems that can be fully personalized, depending on the options you choose), and VM-hosted apps (applications from desktop operating systems) and hosted physical desktops. Only one user at a time can connect each of these desktops.

- **Remote PC Access**: Enables remote users to access their physical office PCs from any device running Citrix Receiver. The office PCs are managed through the XenDesktop deployment, and require user devices to be specified in a whitelist.

Delivery Groups

Delivery Groups specify which users can access which applications and/or desktops on which machines. Delivery Groups contain machines from your Machine Catalogs, and Active Directory users who have access to your Site.

Each Delivery Group can contain machines from more than one Machine Catalog, and each catalog can contribute machines to more than one Delivery Group, but each individual machine can only belong to one Delivery Group at a time.

Application Groups

Application Groups provide application management and resource control advantages over using more Delivery Groups. Using the tag restriction feature, you can use your existing machines for more than one publishing task, saving the costs associated with deployment and managing additional machines. A tag restriction can be thought of as subdividing (or partitioning) the machines in a Delivery Group. Application Groups can also be helpful when isolating and troubleshooting a subset of machines in a Delivery Group.

.

What is Machine Catalog? Can we have single Machine in two Machine Catalog's ? If possible? Explain.

No

What is Delivery Group? Can we have single machine in two delivery groups? If possible? Explain.

Yes

How the Policies are implement, is it through Citrix or AD? Which is best method and Effective Priority?

Policies are a collection of settings that define how sessions, bandwidth, and security are managed for a group of users, devices, or connection type

You can apply policy settings to physical and virtual machines or to users. You can apply settings to individual users at the local level or in security groups in Active Directory. The configurations define

specific criteria and rules, and if you do not specifically assign the policies, the settings are applied to all connections.

Policy settings placed at the Organizational Unit GPO level take the highest precedence on the network.

All Citrix Local Policies are created and managed in the Citrix Studio console and stored in the Site Database; whereas, Group Policies are created and managed with the Microsoft Group Policy Management Console (GPMC) and stored in Active Directory. Microsoft Local Policies are created in the Windows Operating System and are stored in the registry

Studio uses a Modeling Wizard to help administrators compare configuration settings within templates and policies to help eliminate conflicting and redundant setting.

All policies are processed in the following order:

1. The end user logs on to a machine using domain credentials.
2. Credentials are sent to the domain controller.
3. Active Directory applies all policies (end user, endpoint, organizational unit, and domain).
4. The end user logs on to Receiver and accesses an application or desktop.
5. Citrix and Microsoft policies are processed for the end user and machine hosting the resource.
6. Active Directory determines precedence for policy settings and applies them to the registries of the endpoint device and to the machine hosting the resource.
7. The end user logs off from the resource. Citrix policies for the end user and endpoint device are no longer active.
8. The end user logs off the user device, which releases the GPO user policies.
9. The end user turns off the device, which releases the GPO machine policies.

Note: Mixing Windows and Citrix policies in the same GPO is not supported.

Policy processing order and precedence

Group policy settings are processed in the following order:

1. Local GPO
2. XenApp or XenDesktop Site GPO (stored in the Site database)
3. Site-level GPOs
4. Domain-level GPOs
5. Organizational Units

However, if a conflict occurs, policy settings that are processed last can overwrite those that are processed earlier. This means that policy settings take **precedence** in the following order:

1. **Organizational Units**
2. Domain-level GPOs
3. Site-level GPOs
4. XenApp or XenDesktop Site GPO (stored in the Site database)

5. Local GPO

For example, a Citrix administrator uses Studio to create a policy (Policy A) that enables client file redirection for the company's sales employees. Meanwhile, another administrator uses the Group Policy Editor to create a policy (Policy B) that disables client file redirection for sales employees. When the sales employees log on to the virtual desktops, Policy B is applied and Policy A is ignored because Policy B was processed at the domain level and Policy A was processed at the XenApp or XenDesktop Site GPO level

Ref:
http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-12/policies.html

How VDI will register to controller? Explain how VDI registration issues is resolved in your infrastructure.

Before a VDA(Virtual Delivery Agent) can be used, it must register (establish communication) with one or more Controllers on the site. The VDA finds a Controller by checking a list of Controllers called the ListofDDCs. The ListOfDDCs on a VDA contains DNS entries that point that VDA to Controllers on the site. For load balancing, the VDA automatically distributes connections across all Controllers in the list.

Why is VDA registration so important?

- From a security perspective, you're establishing a connection between the Delivery Controller and the VDA. For such a sensitive operation, the expected behavior is to reject the connection if not everything is in perfect shape. You are effectively establishing two separate communication channels: *VDA to Controller, and Controller to VDA.* The connection uses Kerberos, so time synchronization and domain membership issues are challenging. Kerberos uses Service Principal Names (SPNs), *so you cannot use load balanced IP\hostname.*
- If a VDA does not have accurate and current Controller information as you add and remove Controllers in a site, the VDA might reject session launches that were brokered by an unlisted Controller. Invalid entries can delay the startup of the virtual desktop system software. A VDA won't accept a connection from an unknown and untrusted Controller.

If a ListOfDDCs specifies more than one Controller, the VDA attempts to connect to them in random order through load balancing . Failover and load balancing functionality is built into the Citrix Brokering Protocol (CBP). For security reasons, you cannot use a network load balancer, such as NetScaler. VDA registration uses Kerberos mutual authentication,

The ListOfDDCs can also contain Controller groups. The VDA attempts to connect to each Controller in a group before moving to other entries in the ListOfDDCs

There are two communications channels**: VDA -> Controller and Controller -> VDA**.
A component in this process is called Service Principal Name (SPN), which stored as a property in an Active Directory computer object. When your VDA connects to a Controller, it must specify "who" it wants to communicate with; this address is an SPN. If you use a load-balanced IP, mutual Kerberos authentication correctly recognizes that the *IP does not belong to the expected Controller*

There are several methods for configuring Controller addresses on a VDA.

- Policy-based (LGPO or GPO)
- Registry-based (manual, GPP, specified during VDA installation)
- Active Directory OU-based (legacy OU discovery)
- MCS-based (personality.ini)

Regardless of which method you use to specify Controllers, Citrix recommends using an FQDN address. An IP address is not considered a trusted configuration, because it's easier to compromise an IP than a DNS record

Virtual Delivery Agent (VDA) Registration Troubleshooting

- Check that VDA software is installed
- Check Event Log
- Confirm Controller configuration
- VDA machine is member of domain –SPN Issues
- VDA Citrix Desktop Service "BrokerAgent.exe" is running
- Windows Firewall is open for communication
- Check Functional level
- If power managed, make sure machine is powered on
- Check DHCP configuration
-Domain Name Services (DNS) not Properly Configured
-Time Synchronization not Properly Configured

Tools:

XDPing:
The XDPing tool is a legacy command-line based application, which automates the process of checking for the causes of common configuration issues in a XenDesktop environment

Citrix Health Assistant:

The Citrix Health Assistant tool has replaced XDPing. Citrix Health Assistant is the recommended tool; automating a series of health checks to identify possible root causes for common VDA registration issues. The tool is graphical UI based but also supports command line commands

Ref:
VDA Troubleshooting
https://support.citrix.com/article/CTX136668
https://support.citrix.com/article/CTX126992

XDPing
https://support.citrix.com/article/CTX123278
Citrix Health Assistant
https://support.citrix.com/article/CTX207624

How to resolve power state unknown issue? Will there be any impact if power state is unknown in accessing the VDI

Symptoms or Error

Machines in Desktop Studio or Desktop Director display a **Power State of Unknown**

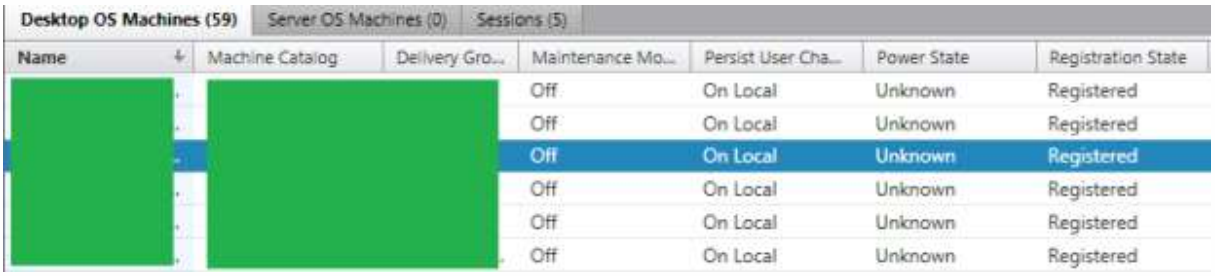

And even though the machines are Registered, they CANNOT be brokered:

Background

The Desktop Delivery Controller (DDC) broker log contains Citrix Pool Management exceptions relating to the machine in question, such as:

23/08/11 15:18:41.7081: HostingManagement:PollForPowerStateUpdates: problem while fetching state for 02e0d308-c9cf-4fbc-9fb2-aaaaaaaaaaaa:
Citrix.ManagedMachineAPI.NoSuchManagedMachineException: Exception of type 'Citrix.ManagedMachineAPI.NoSuchManagedMachineException' was thrown. at Citrix.PoolManagement.VMManager.VmmImplementation.MicrosoftScvmmProvider.Connection.Execute(Runspace runspace, Action`1 action) at Citrix.PoolManagement.VMManager.VmmImplementation.MicrosoftScvmmProvider.Connection.GetVM

(String id, Runspace runspace) at Citrix.PoolManagement.VMManager.VmmImplementation.MicrosoftScvmmProvider.SCVMMConnector.<>c__DisplayClass44`1.<CallAndRetry>b__42(Runspace runspace) at Citrix.PoolManagement.VMManager.VmmImplementation.MicrosoftScvmmProvider.SCVMMConnector.Call[T](Func`2 worker) at Citrix.PoolManagement.VMManager.VmmImplementation.MicrosoftScvmmProvider.SCVMMConnector.CallAndRetry[T](Func`3 worker)

Solution

To update the correct host machine ID on the DDC, complete one of the following solutions:

Follow KB to check different solutions https://support.citrix.com/article/CTX131267

Problem Cause

The DDC must communicate with the hypervisor using the virtual machine ID. If the DDC has an incorrect machine ID for the virtual machine, it is unable to read the machine power state and throws an exception in the broker log. If the power status is **Unknown**, the DDC will not be able to manage any power functions on the virtual machine

Cmd :Get-BrokerHypervisorConnection -> Shows DDC connection status

What is LTSR(Long Term Service Release) and Current Release (CR)? By opting this, what is the benefit to customer? What is LTS Assistant?

What is LTSR?

As a benefit of Software Maintenance, Long Term Service Releases (LTSR) of XenApp ,XenDesktop,XenServer enable enterprises to retain a particular release for an extended period of time while receiving minor updates that provide fixes**,** typically void of new functionality. Long Term Service Releases (LTSR) is ideal for large enterprise production environments where you would prefer to retain the same base version for an extended period

A Long Term Service Release guarantees 5 years of mainstream support and an optional 5 years of extended support (needs to purchased separately). This includes cumulative updates every 4 to 6 months, a new LTSR version of XenApp / XenDesktop every 12 to 24 months and any potential (hot) fixes

A valid Software Maintenance (SM) contract is needed to make use of the LTSR or CR servicing option.

Ideal customer environment for a LTSR is for the **c**ustomers who typically follow a 3-5 year version upgrade cycle

Long Term Service Releases will have a regular cadence of Cumulative Updates that will typically contain only fixes

What is Current Release?

Any new release of XenApp/XenDesktop/XenServer will be labeled a Current Release. With the CR servicing option you can always make use of (install) the most recent XenApp and/or XenDesktop versions including all the latest enhancements and additions that come with it.

Its release cycles are much shorter with a new version release being announced every three to nine months in general.

Citrix recommends that large enterprise customers have a combination of Current Release and Long Term Service Release environments.

Switching from a LTSR to a CR servicing, and vice versa, is always optional as well

All initial releases of XenApp/XenDesktop/XenServer will be a Current Release. There will likely be multiple Current Releases of a major XenApp/XenDesktop/XenServer version (i.e. 7.6, 7.6 FP1, 7.6 FP2, 7.6 FP3, 7.7, 7.8 ,7.9,7.11,7.13,7.14); however, there will likely only be one LTSR release of that version after that release is considered customer-tested and industry-proven (i.e. 7.6 FP3).

How will the customer know if their environment is Long Term Service Release compliant?

Citrix support and engineering have developed the LTSR Assistant tool which will scan your environment and compare your environment with the necessary LTSR components to determine if you are compliant. The tool provides a report that will outline the necessary updates to achieve compliance. The LTSR Assistant tool is available for download at http://support.citrix.com/article/CTX209577.

Explain the difference in functionality Data Collector vs Delivery Controller

Data Collector	Delivery Controller
LHC(Local Host Cache)	No LHC
No Connection Leasing	Connection Leasing
Has static as well dynamic(run-time) information cached locally	Pulls all information, static as well as dynamic from the central Site database
Communicates with the IMA store, Peer Data Collectors and its session Hosts(within its own zone) on a scheduled interval, or when a Farm configuration change has been made.	There is no direct communication between delivery controllers. No scheduled communication between the VDA's and/or Site databases, only when needed.

Often hosts user session, but can be configured as a dedicated data collector as well.	Is responsible for brokering and maintaining new and existing user session only.
Need to have the same operating system as all other session hosts and DC's within same Farm.	Can have a different operating system installed then the server and desktop VDA's
Has all the XenApp 6.5 or earlier bits and bytes fully installed.	Core services installed only. The HDX stack is part of VDA software
Each Zone has one Data Collector. Having multiple data collector means multiple zones.	Zones are optional. When configured they do need at least one Delivery controller present.
Can, and sometimes need to be elected. Configure at least one other Session Host per zone that can be elected as a Data Collector when needed.	Election does not apply. Deploy multiple, at least 2 Delivery controllers per site /zone (again one per zone is the minimum)
When IMA Db is down, no Farm wide configuration changes are possible. Everything else continues to work as expected due to the LHC present on the Data Collectors and Session Hosts in each Zone.	When Central Site DB is down, no site wide configuration changes are possible. By default, Connection Leasing will kick in, enabling users to launch sessions which are assigned at least once during last 2 weeks prior to DB going offline.
Does not have any direct connection (API) with a Hypervisor or cloud platform management capabilities.	A Delivery controller can have a direct connection (API) with a Hypervisor or cloud platform of choice.
Session Hosts as well as Data Collectors directly communicate with IMA database.	Almost all the communication directly flows through a Delivery Controller to Central Site DB.
When a XenApp server boots it needs a IMA service but it will not register itself anywhere.	VDA's needs to successfully register themselves with a Delivery controller.

What databases will create during installation of Citrix App/Xendesktop and mention supported database High Availability methods?

A XenApp or XenDesktop Site uses three SQL Server databases:

- **Site** – (also known as Site Configuration) stores the running Site configuration, plus the current session state and connection information.
- **Configuration Logging** – (also known as Logging) stores information about Site configuration changes and administrative activities. This database is used when the Configuring Logging feature is enabled (default = enabled).
- **Monitoring** – stores data used by Director, such as session and connection information.

There are several high availability solutions to consider for ensuring automatic failover

- **AlwaysOn Availability Groups**
- **SQL Server database mirroring**
- **SQL clustering:**
- **Using the hypervisor's high availability features**

Note: Installing a Controller on a node in an SQL clustering or SQL mirroring installation is not supported.

Have you done any Hotfix Rollup installation in XenApp or Hotfixes/Cumulative updates installations in XenDesktop 7.x? If yes, what is the best procedure to follow in both?

Old versions

This article introduces and discusses some of the best practices for installing and deploying hotfix rollup packs for XenApp and Presentation Server. Citrix periodically releases hotfix rollup packs that might include bug fixes, security fixes, and enhancements for all currently supported XenApp versions.

Order of Deployment
The order of hotfix rollup pack deployment is very important, especially in large or complex farm configurations. Before installing a hotfix rollup pack:

- Ensure that there are no existing connections on the servers where the update is being run
- Ensure that the Independent Management Service is running
- Back up the data store database

After a deployment, always ensure that all servers in the farm are updated to the same hotfix rollup pack level.

Citrix recommends the following order of deployment:

- Zone data collector

- Backup zone data collectors
- Database connection server (Applies only to Resource Manager for XenApp 5 for Microsoft Windows Server 2003)
- Primary farm metric server (Applies only to Resource Manager for XenApp 5 for Microsoft Windows Server 2003)
- Backup farm metric server (Applies only to Resource Manager for XenApp 5 for Microsoft Windows Server 2003)
- Member servers

The hotfix rollup pack installs and updates only servers that have core XenApp (mps.msi) and/or XenApp Advanced Configuration (cmc.msi) installed

How to Install XenDesktop/XenApp 7.x Controller Hotfixes

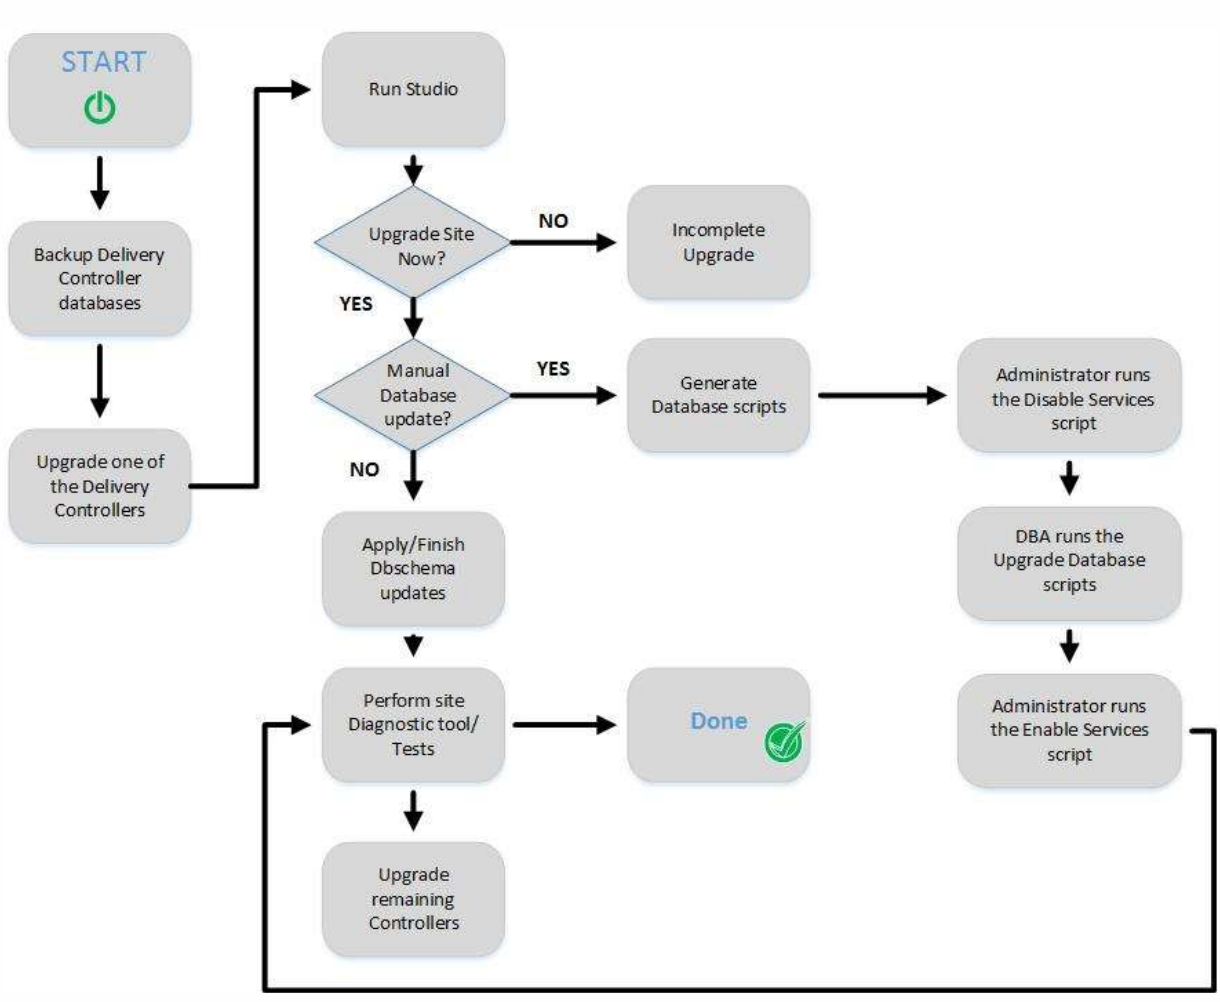

Ref:
https://support.citrix.com/article/CTX120842
https://support.citrix.com/article/CTX201988

Explain step-by-step communication process flow when user launches VDI or Application through Netscaler Access gateway.

The following list explains the basic Logon Process and the flow of communication for XenDesktop 5.x and XenDesktop/XenApp 7.x with out Netscaler

Below is general diagram only for reference…

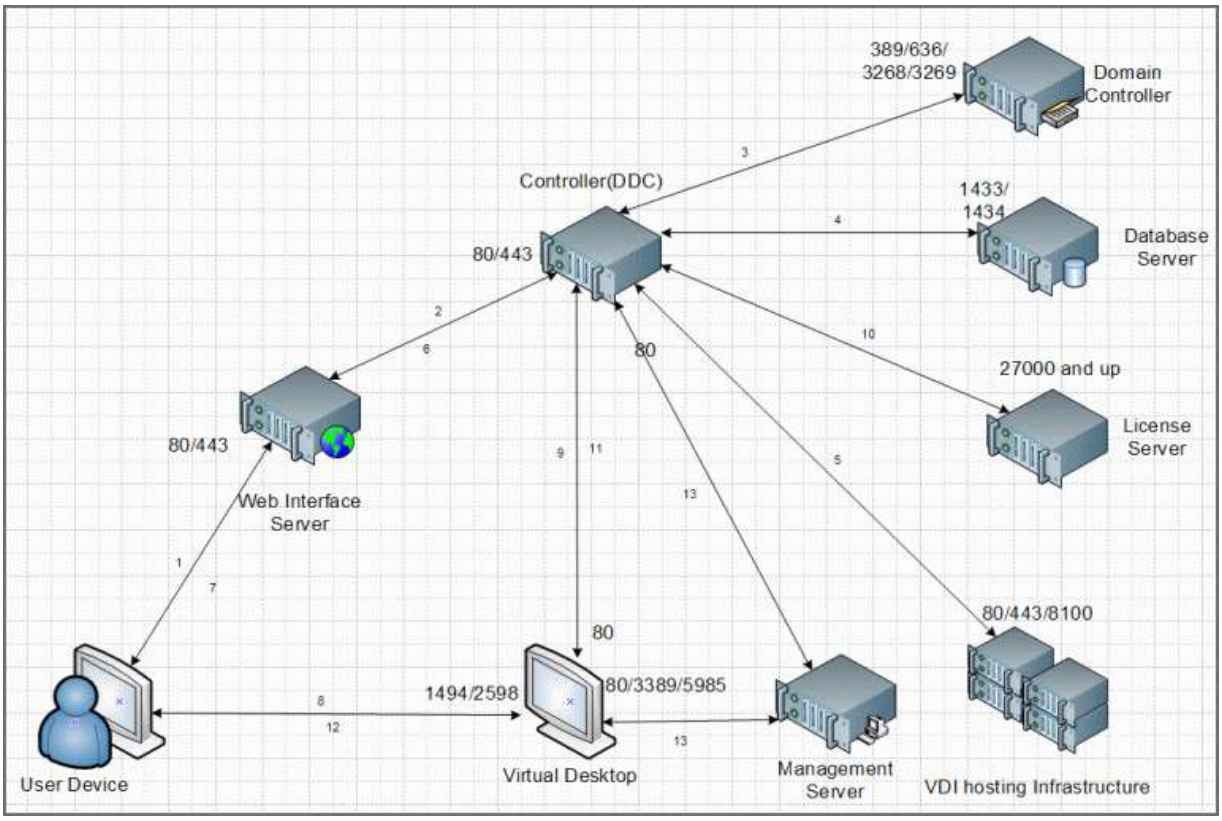

1. The user device submits credentials to the Citrix web site hosted on either a Web Interface (WI) or StoreFront(SF) web server.
2. a For StoreFront, the credentials are verified directly to a Domain Controller (Similar to Step 3). SF then passes the validation over to the DDC to begin resource enumeration (Step 4).
2.b For Web Interface, the username and password is passed to the Desktop Delivery Controller (DDC).
3. The DDC then queries a Domain Controller with the end user's credentials to verify user authorization.
4. DDC then queries the site SQL database for the end user's assigned Delivery Groups (SQL named instance uses default ports 1434 and 1433).
The resources defined by the Delivery Groups are sent to the WI or SF server and presented to the user (enumeration).

5. User clicks on one of the resource icons to start a desktop or application session. This request is sent to the Delivery controller through the storefront. Using the Delivery Group information obtained from the database for the resource the user requested, DDC queries the hypervisor about the status of resources within that group.
6. DDC identifies to Web Interface/StoreFront the virtual machine it assigned for this particular session.
7. Web Interface/StoreFront creates and sends an ICA file to the Citrix Receiver pointing to the virtual machine that hypervisor identified.
8. Citrix Receiver establishes an ICA connection to the specific virtual machine that the DDC allocated for this session and passes the Ticket embedded in the ICA file.
Note: In the case of Single Sign-on or Smartcard authentication, instead of the ticket, Launch Ref is exchanged with the DDC.
9. Virtual Delivery Agent (VDA) verifies the license file with the DDC and exchanges the ticket for user credentials with the DDC.
10. DDC queries Citrix License server to verify that the end user has a valid ticket.
11. DDC passes session policies to the VDA, which then applies session policies to the virtual machine.
12. Citrix Receiver displays the selected resource to the end user.
13. Administrator and help desk personnel use Desktop Director and Desktop Studio tools to manage the desktops from the management server.

To discuss in deeper

External user authentication through NetScaler (Below steps given until resource enumeration)

Let's assume that your NetScaler Gateway is set up and configured to integrate your StoreFront server(s), you have Receiver installed, SSL certificates are present, and that a STA / XML / Broker service address (Delivery Controller) and a domain controller for authentication purposes are also configured.

1. A user opens up a web browser and connects to the external URL of the NetScaler Gateway (using SSL over port Nr. 443) here he or she will fill in his or her username and password. A locally installed Citrix Receiver can also be used to establish a direct connection to the NetScaler Gateway. Citrix Receiver uses so called Beacons to determine if a connection is internal or external and handles it accordingly.
2. The NetScaler will take the user credentials and authenticate them (*session ticket*) against Active Directory over TCP port Nr. 389. The NetScaler has its own authentication service just like StoreFront mentioned earlier.

3. Once authenticated the user session gets redirected to StoreFront where it will first perform a callback to the NetScaler that handled authentication to validate the user. The authentication details will then be send to the StoreFront.
4. From here, the user credentials will be forwarded, as part of the earlier mentioned XML query, to the configured Broker (XML) service on one of the Delivery Controllers. Both these transactions will use port Nr. 443 / SSL.
5. In between, StoreFront will check its local data store for any existing recourse subscriptions and stores these in memory.
6. The Broker (XML) service will again contact a domain controller (port Nr. 389) to validate the user credentials, *note that this is different to the user authentication process, as we've established earlier. During this process it will find out to which security groups the user belongs*.
7. With this information the Delivery Controller, or Broker (XML) service, will contact the central Site database to find out which resources have been assigned to the user. It does this over port Nr. 1433 / 1434
8. The Broker (XML) service will return an XML file to the StoreFront server including all assigned resources.
9. StoreFront will generate a web page containing all the assigned resources, which will be routed through the NetScaler Gateway and presented to the user. The users home screen will be populated with any pre-subscribed resources (Keywords). Depending on how you connected, your resources will be displayed either directly using a Receiver for web sites or you'll find them within the locally installed Citrix Receiver instead. The user will be able to browse its own personal app store for any assigned resources to which he or she can subscribe and then launch.

Note:

If you don't enable authentication on the NetScalers login page the NetScaler will contact StoreFront and the user will be presented (through the NetScaler) with the StoreFront login page (Receiver for web sites). The user fills in his or her credentials and authentication will be handled by StoreFront.

Internal user authentication through StoreFront (Below steps given until resource enumeration)

So what happens when a user authenticates internally, directly to StoreFront? Let's have a look. Same rules apply here, use port Nr. 443 where you can.

1. A user opens up a web browser and connects to the internal StoreFront URL where he or she will fill in his or her username and password. This method is also referred to as Receiver for web sites as mentioned above (*don't confuse this with the HTML 5 based Receiver for web, they're not the same*). A locally installed Citrix Receiver can also be used to establish a direct connection to StoreFront, which is probably the preferred method whenever possible. The earlier mentioned (NetScaler) Beacon functionality applies here as well.
2. Next, the StoreFront authentication service will pick up the user credentials and contact a domain controller to authenticate the user in Active Directory over TCP port Nr. 389. Here I'd like to note that if domain pass-through authentication is enabled on the StoreFront server, this step would automatically be skipped.
3. Once authenticated the user credentials, as part of the XML query, will be send to a Delivery Controller.
4. In between, StoreFront will check its local data store for any existing recourse subscriptions and stores these in memory.
5. During the next phase the Broker (XML) service will again contact a domain controller (port Nr. 389) to validate the user credentials, this is different to the user authentication process, as we've established earlier. During this process it will find out to which security groups the user belongs.
6. With this information the Delivery Controller, or Broker (XML) service, will contact the central Site Database to find out which resources have been assigned to the user. It does this over port Nr. 1433 / 1434.
7. The Broker (XML) service will return an XML file to the StoreFront server over port Nr. 443 / SSL.
8. StoreFront will generate a web page containing all the assigned resources, which will be presented to the user. The users home screen will be populated with any pre-subscribed resources (Keywords). Depending on how you connected, your resources will be displayed either directly using a Receiver for web sites or you'll find them within the locally installed Citrix Receiver instead. The user will be able to browse its own personal app store for any assigned resources to which he or she can subscribe and then launch.

Launch process for HSD or Published application post enumeration

Above process explained till resource enumeration process, resource displayed on home screen once enumerated. Let see how the launch process. Just as with the authentication process, there are some differences between the internal and external resource launch process.

Also, when launching a HSD or a published application there is an extra load balance step involved as well. Let's start with an external HSD launch trough NetScaler. Note that the below process is basically XenApp as we knew it earlier

Points to remember:

- *"The STA is only used when traffic traverses a NetScaler, so you don't have to worry about the STA service and its tickets when authentication takes place internally."*
- *Note that STA (service) is also part of the Broker service, and it has been since Presentation Server 4.0*
- *STA ticket gets generated and send back after a user launches an application/desktop, not during the resource enumeration process.*

Steps:

1. Assuming that the login and enumeration process finished without any issues (see above) the user is free to subscribe to and *launch any applications and/or desktops that have been assigned* to him or her. As an example, let's say that the user tries to launch a (XenApp) Hosted Shared Desktop session.
2. **After the user clicks the icon the launch request** is send to the NetScaler Gateway from where it will be forwarded to the StoreFront server.
3. The StoreFront server will contact the Broker (XML/STA) service, or Delivery Controller, to find out if and where the resource is available and where it can be best started. This is where the well-known XenApp load balancing mechanism comes into play. *Which as of FMA needs to be configured through policies.*
4. During this time, the StoreFront server will also request an STA ticket from the Broker (XML/STA) service. It will include the user, domain and resource name it wants to start. It will also request a 'least loaded' server as part of the load balancing process.

5. The Broker (XML/STA) service will query the central Site database (ports Nr. 1433 and 1434) to find out which server is able to offer the requested resource. The Delivery Controller will use this information together with its load balance algorithm to decide which server to connect to.
6. At this time the Broker (XML/STA) service will create the STA ticket mentioned earlier. *This will include information on the server and resource to connect to, as discovered in the previous steps mentioned.*
7. Next the Broker (XML/STA) service will send this information back to the StoreFront server in the form of an XML file.
8. Based on this information the StoreFront server will then generate a **launch.ica** file (uses the default.ica file as a template) *containing the STA ticket and a whole bunch of other connection properties that are, or might, be needed. This will also include the FQDN/DNS name of the NetScaler Gateway itself.*
9. StoreFront passes on this information down through the NetScaler Gateway onto the locally installed Receiver, which initiated the connection to begin with.
10. *The locally installed Receiver will read and autolaunch the launch.ica file to set up a connection to the NetScaler Gateway (443 / SSL).
11. **From here the NetScaler Gateway will first contact the Broker (XML/STA) service** (this address is configured on the NetScaler as well) to verify if the earlier generated STA ticket, as part of the launch.ica file, is still valid.
12. The Broker (STA) service will validate the STA ticket from memory. Once verified it will send back the IP address, port Nr. Resource name etc. of the machine and the resource it needs to connect to. Once done the STA ticket will be deleted.
13. The NetScaler Gateway will set up a new ICA connection using port 1494 (ICA) or 2598 (CGP – Common Gateway Protocol) depending on config.
14. ***The installed VDA will verify its license file with the Delivery Controller***.
15. **The Delivery Controller checks with the Citrix License server to verify that the end user has a valid ticket**. This will also be done for any Microsoft (CAL) licenses, with regards to HSD and published applications, that might be involved.
16. At this time any applicable session policies will be passed onto the VDA applying them to the session.
17. Finally the HSD is launched and the NetScaler Gateway acts as a proxy between the user and the XenDesktop resource in the data center.

18. Somewhere in between the session/connection information will be passed on and registered in the central Site Database where it will be used for future load balance purposes.

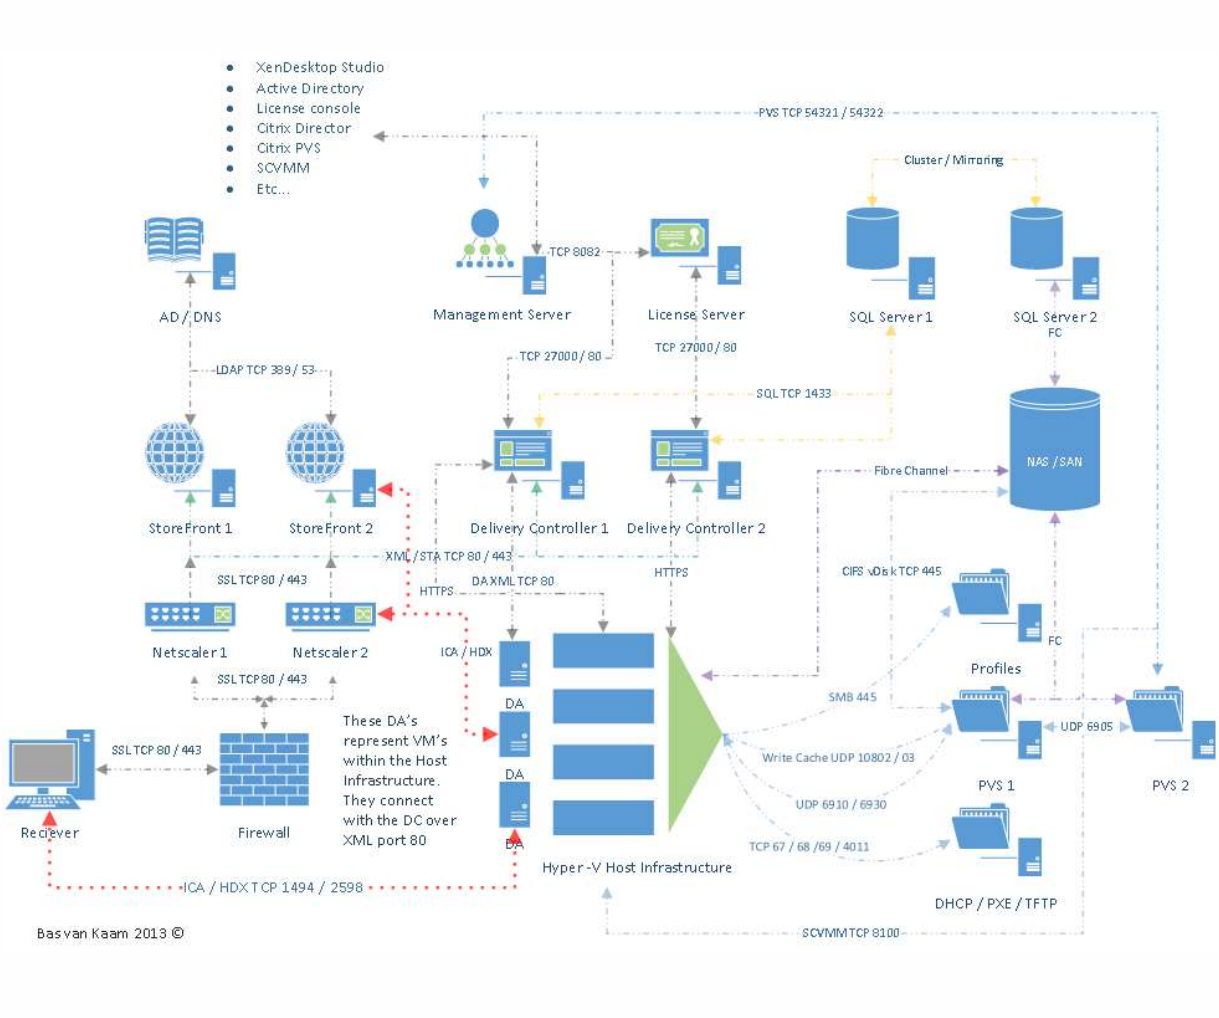

Launch process for a pooled VDI post enumeration

Now that we have seen which steps are involved when launching a resource externally, a Hosted Shared Desktop in this case, let's have a look and see what happens when we launch a pooled VDI virtual machine internally.

After this we will have looked at an external and internal resource launch, a HSD, which is comparable to a published application, and a VDI virtual machine. Again, user authentication and resource enumeration has successfully completed, here we go (again)

Steps

1. *As mentioned we will launch a pooled VDI virtual machine this time*. Let's assume that the VM is pre-subscribed and *already present on the users home screen*, never mind how we connected: locally installed Receiver or using the Receiver for web sites.
2. After the user clicks the icon the StoreFront server will contact the Broker (XML/STA) service, or Delivery Controller, to check if any registered VDA's are available. It does this by communicating with underlying Hypervisor platform through the Host service on the Delivery Controller.
3. *If needed it will first start / boot a VM*. It's not uncommon to pre-boot a few VM's, since, as you can probably imagine, this will positively influence the overall user experience.
4. Next the Delivery Controller, or Broker (XML/STA) service, will contact one of the VDA's and sends a startlistening request. *By default the VDA isn't listening for any new connections on port Nr. 1494 or 2595 until it gets notified that a user wants to connect*.
5. As soon as the VDA is listening, the Broker (XML/STA) service will send this information back to the StoreFront server in the form of an XML file.
6. Based on this information the StoreFront server will then generate a launch.ica file (it uses the default.ica file as a template) containing the IP address of the VDA and a whole bunch of other connection properties that are, or might, be needed. This is send down to the user.
7. The locally installed Receiver (or HTML 5 based Receiver) will read and autolaunch the launch.ica file initiating a direct connection from the users end-point to the VDA.
8. The installed VDA will verify its license file with the Delivery Controller.
9. The Delivery Controller checks with the Citrix License server to verify that the end user has a valid ticket.
10. At this time any applicable session policies will be passed on to the VDA and the session is launched.

Note: In above steps, STA process will not in place.

Points to remember:

- *"The STA is only used when traffic traverses a NetScaler, so you don't have to worry about the STA service and its tickets when authentication takes place internally."*
- *Note that STA (service) is also part of the Broker service, and it has been since Presentation Server 4.0*

- *STA ticket gets generated and send back after a user launches an application/desktop, not during the resource enumeration process.*

Ref:

https://support.citrix.com/article/CTX128909

http://www.basvankaam.com/2014/11/24/the-ultimate-xendesktop-7-x-internals-cheat-sheet/

What happens inside the VDA during launch time? Explain the services communication flow for which responsible to launch VDI

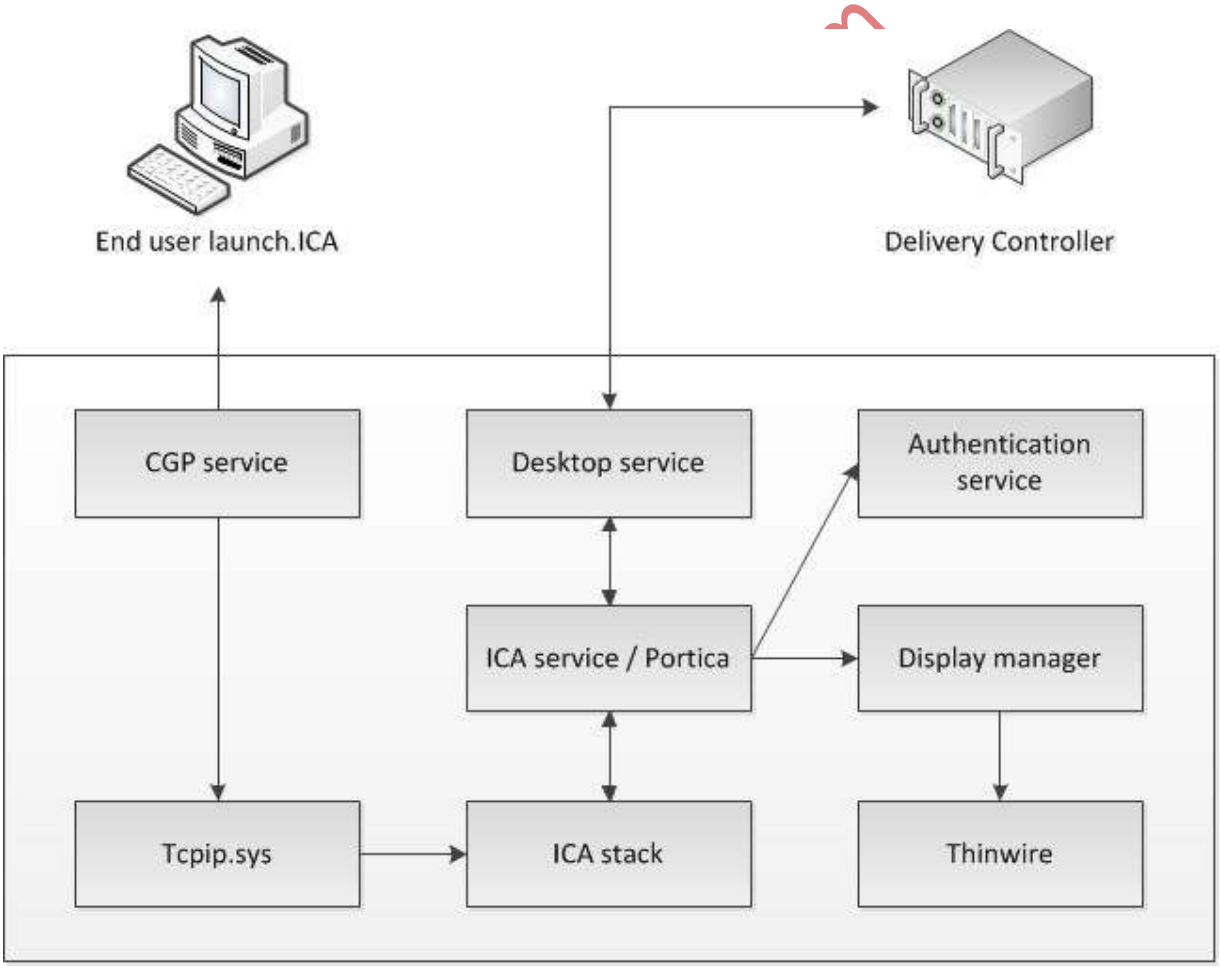

The process below assumes that session reliability is enabled and that a desktop OS VDA gets launched as we've seen in the previous section, the below takes place somewhere around step 7. Remember that as of XenDesktop 7 there is also a server OS based VDA.

1. The CGP service will receive the connection and sends this information on to the tcpip.sys, which will forward it to the ICA stack.
2. The ICA stack will notify the ICA service a.k.a. the Portica service (Picasvc) that a connection has been made after which the Picasvc will accept the connection.
3. Then the ICA service will lock the workstation because the user needs to be authenticated to ensure that the user is allowed access to that particular machine.
4. As soon as the user logs on to the workstation the Portica service will communicate with the display manager to change the display mode to remote ICA, this request will be forwarded to the Thinwire driver.
5. In the meantime the Portica service will hand over the 'pre logon' ticket data, which it received from the ICA stack, up to the Desktop service and from there back to the Delivery Controller in exchange for 'real' credentials.
6. The Desktop service receives the users credentials, which are send back to the Portica service.
7. The Portica service contacts the authentication service to logon the user and this is sort of where the process ends

It's kind of hard to find information on these kinds of traffic flows. Do note that with the newer XenDesktop 7.x versions things might have slightly changed, but I still think this should give you a good idea on what's going on internally

What is Secure Ticket Authority (STA)? Explain importance and functionality.

STA was first introduced with one of the earlier Secure Gateway editions over ten years ago. It (the STA) runs as a service and is part of the Broker service on the Delivery Controller just like the XML service. During the resource launch sequence the StoreFront server as well as the NetScaler will both need to be able to communicate with the STA. As such, you will need to configure the NetScaler as well as the StoreFront server(s) or Web-Interface server(s) to point to the same XML/STA service(s)/Deliver Controller(s).

Once a user launches a resource, externally through NetScaler Gateway, at one point a secure ticket will be requested. As we will see shortly the STA ticket will eventually end up in the launch.ica file generated by StoreFront and/or Web-Interface. Once generated, the Delivery Controller hosting the STA service will hold the STA ticket information in memory for a configurable amount of time. As soon as a secure session is established, the NetScaler Gateway responsible for handling the session only has to check the STA ticket (as part of the .ica launch file) with the STA service that originally generated the ticket. It (the STA service) does this from memory where the ticket was stored after it was created and send back to

the StoreFront server as part of the XML file mentioned earlier. More (detail) on this in the overview below.

"The STA is only used when traffic traverses a NetScaler, so you don't have to worry about the STA service and its tickets when authentication takes place internally."

Broker, XML and STA.

Be aware that the STA (service) is also part of the Broker service, and has been as of Presentation Server 4.0. Before that, it was written *as an ISAPI extension for Microsoft Internet Information Services*, or IIS. I also highlighted the so-called XML service multiple times. I put the XML and STA services between brackets because as of XenDesktop 4.x the XML service (**ctxxmlss.exe**) has been rewritten in *.NET and became part of the Broker service*. So the Broker service is actually build up out of three separate services, all handling different tasks, it brokers connections, it enumerates resources and it acts as the Secure Ticket Authority, generating and validating STA tickets

"Make sure that the Broker (XML/STA) service on the NetScaler and the Storefront server is configured identically. The same applies to the load balance/fail over order in which you configure them."

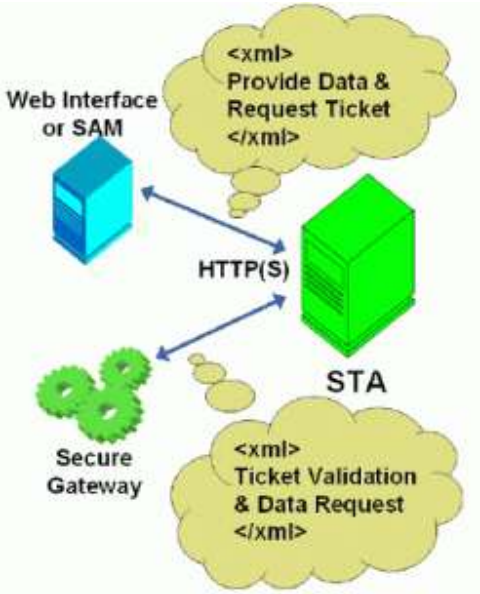

Ref:

FAQ: Citrix Secure Gateway/NetScaler Gateway Secure Ticket Authority -
https://support.citrix.com/article/CTX101997
https://support.citrix.com/article/CTX132334/
http://www.basvankaam.com/2014/11/24/the-ultimate-xendesktop-7-x-internals-cheat-sheet/

What is Machine Catalogue and options while creations? Requirement to have two different OS in one catalog? Is it possible, if yes explain?

Collections of physical or virtual machines are managed as a single entity called a machine catalog. **All the machines in a catalog have the same type of operating system: server or desktop**. A catalog containing Server OS machines can contain either Windows or Linux machines, not both.

MCS catalog creation summary

Here's a brief overview of default MCS actions after you provide information in the catalog creation wizard.

- If you selected a master image (rather than a snapshot), MCS creates a snapshot.
- MCS creates a full copy of the snapshot and places the copy on each storage location defined in the host connection.
- MCS adds the machines to Active Directory, which creates unique identities.
- MCS creates the number of VMs specified in the wizard, with two disks defined for each VM. In addition to the two disks per VM, a master is also stored in the same storage location. If you have multiple storage locations defined, each gets the following disk types:
    - The full copy of the snapshot (noted above), which is read-only and shared across the just-created VMs.
    - A unique 16 MB identity disk that gives each VM a unique identity. Each VM gets an identity disk.
    - A unique difference disk to store writes made to the VM. This disk is thin provisioned (if supported by the host storage) and increases to the maximum size of the master image, if necessary. Each VM gets a difference disk. The difference disk holds changes made during sessions. It is permanent for dedicated desktops. For pooled desktops, it is deleted and a new one created after each restart.

Alternatively, when creating VMs to deliver static desktops, you can specify (on the **Machines** page of the catalog creation wizard) thick (full copy) VM clones. Full clones do not require retention of the master image on every data store. Each VM has its own file

Operating system

Each catalog contains machines of only one type:

- **Server OS:** A Server OS catalog provides hosted shared desktops and applications. The machines can be running supported versions of the Windows or Linux operating systems, but the catalog cannot contain both. (See the Linux VDA documentation for details about that OS.)
- **Desktop OS:** A Desktop OS catalog provides VDI desktops and applications that can be assigned to various different users.
- **Remote PC Access:** A Remote PC Access catalog provides users with remote access to their physical office desktop machines. Remote PC Access does not require a VPN to provide security.

Ref:

https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-13/install-configure/machine-catalogs-create.html

Difference between Server OS VDA & Desktop OS VDA

During Server VDA installation one of the things it will do is register the Broker Agent Service (used for direct communication with the Delivery Controller) which is similar to the Desktop VDA process. Next it will install the multi-user ICA stack, as it does with earlier XenApp versions, which will then become part of Termsrv creating the so called ICA stack listener waiting for new ICA connections (kernel mode). The ICA stack itself has changed very little with the introduction of the FMA, one its biggest changes is to be found in its communication interface, which is now better known as the earlier mentioned Broker Agent. Last but not least it will install and configure the Citrix Stack Control Service, this is its display name within the Windows services overview.

"Each Terminal Server protocol (like Citrix's ICA) will have a protocol stack instance loaded (a listener stack awaiting a connection request). When installed, the Server VDA basically extends Microsoft's RDS protocol with the HDX feature set / protocol. I guess some things never change."

The biggest difference between the two Delivery Agents is the ICA protocol stack. ***For desktop machines, Citrix ships a single-user ICA stack (internally known as Portica) which allows only one ICA session at a time***. This version connects users to the machine's console session, similar GoToMyPC or other Remote Access products for a Desktop OS. It also includes additional HDX features such as USB and Aero redirection, which are only available on a single-user machine. ***For server machines, Citrix includes a multi-user ICA stack, which extends Windows Remote Desktop Services with the HDX protocol.*** This is the same ICA protocol stack developed for Citrix XenApp, just with a different management interface to make it compatible with Excalibur controllers.

Ref:

http://www.basvankaam.com/2014/12/15/xendesktop-7-x-fma-internals-continued-the-server-vda-in-more-detail/

http://www.basvankaam.com/2014/11/24/the-ultimate-xendesktop-7-x-internals-cheat-sheet/

Network Ports used by Citrix

The following table lists the default network ports used by XenApp and XenDesktop Delivery Controllers, Windows VDAs, Director, and Citrix License Server. When Citrix components are installed, the operating system's host firewall is also updated by default, to match these default network ports.

Component	Usage	Protocol	Default incoming ports
VDA	ICA/HDX	TCP, UDP	1494
VDA	ICA/HDX with Session Reliability	TCP, UDP	2598
VDA	ICA/HDX over TLS	TCP	443
VDA	ICA/HDX over WebSocket	TCP	8008
VDA	ICA/HDX Audio over UDP Real-time Transport	UDP	16500..16509
VDA	ICA/HDX Framehawk	UDP	3224-3324
VDA	ICA/Universal Print Server	TCP	7229
VDA	ICA/Universal Print Server	TCP	8080
VDA	Wake On LAN	UDP	9
VDA	Wake Up Proxy	TCP	135
VDA	Delivery Controller	TCP	80
Delivery Controller	VDA, StoreFront, Director, Studio	TCP	80
Delivery Controller	StoreFront, Director, Studio over TLS	TCP	443
Director	Delivery Controller	TCP	80, 443
License Server	License Server	TCP	27000
License Server	License Server for Citrix (vendor daemon)	TCP	7279
License Server	License Administration Console	TCP	8082
License Server	Web Services for Licensing	TCP	8023

Ref:

http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-12/technical-overview/default-network-ports.html

Provisioning Methods – (MCS & PVS)

Major Differences between MCS & PVS with pros & cons? Which design do you suggest to customer with pros & Cons?

Machine Creation Services:

- MCS is considered to be easy. It is managed and configured directly from studio and you do not need any additional infrastructural components as you do with PVS.
- MCS is based on differencing disk technology.
- Your base(Golden Image) will be copied to overall data stores, which are part of hosting connection(storage connection).
- When using MCS, rollbacks are treated the same way as a new or updated base image : they will again need to be copied over to all data stores involved. Note that in some cases the previous image might still in use by some machines. If so, then no full copy will be needed.

Provisioning Services:

Provisioning services is based on software streaming technology. Simply put ,a single read-only vDisk(Virtual Disk) will be streamed over the network to multiple so-called target devices, which can be XenApp servers or XenDesktop VDI-based VM's. You will always have at least 2 provisioning servers for HA purposes, or more depending on the size of your deployment and the number of target devices that need to be serviced

- Provisioning Services streams a base image over the network down to either virtual or physical machines.
- It works for both desktop as well as server Operating Systems.
- A device using a vDisk is also referred to as Target Device.
- The Machine used to create and maintain the vDisk is referred to as the Master Target Device.
- Target Devices are managed using Device Collections.
- The life cycle of a vDisk consists of Creation,Deployment,Maintainance and finally Retirement. For this, we can leverage the built-in PVS versioning mechanism.
- While Personal vDisks (PvDs) have their use, apply them wisely: it is not for everyone. And while this may be somewhat off topic, in many cases where VDI ie being considered , RDSG might make more sense.
- Check CTX117372,CTX124185 for some PVS best practices & PVS vDisk's.

While past , it was always considered to use physical machines for your PVS ,today virtual machines are almost always recommended by Citrix. The same applies to isolating PVS traffic but still it depends on the requirement whether to isolate PVS traffic or not as multiple blogs advised not required of network isolation ,however isolation sometimes might make sense is because of security considerations

Citrix recommends, Provisioning Services is the best option for managing enterprise-scale XenDesktop deployments, but Machine Creation Services might be easier for smaller businesses to use because it's simpler to set up

Citrix PVS has been around since 2006, when the company acquired the technology from Ardence. MCS came out in 2010 with XenDesktop 5, but Citrix has stated that MCS is not a replacement for PVS; it is another virtual desktop provisioning option for IT

Support for Azure or AWS.

Azure and AWS do not support PVS, so if your company's long-term plan is to move in that direction, MCS is a better option. And if you plan to keep your XenApp or XenDesktop environments on premise, while moving all of your other servers to Azure or AWS, you're going to potentially introduce application latency issues

PVS	MCS
Citrix PVS has been around since 2006, when the company acquired the technology from Ardence	MCS came out in 2010 with XenDesktop 5, but Citrix has stated that MCS is not a replacement for PVS.it is another virtual desktop provisioning option for IT
PVS requires administrators to set up an additional infrastructure of Provisioning Servers before they begin virtual desktop production.	Citrix MCS is a component of XenDesktop that IT can control through the Studio management console
PVS captures a XenApp/XenDesktop image as a .vhd/.vhdx file then streams a read-only copy to a target device	MCS takes a snapshot of a virtual machine, copies it to a storage location and clones are made that read this copy
PVS is network-based	MCS is a Storage based
PVS less IOPS	MCS uses more IOPS
Azure and AWS do not support PVS	Supports Azure & AWS ->
Citrix recommends Provisioning Services is the best option for managing enterprise-scale XenDesktop deployments	Citrix recommends Machine Creation Services might be easier for smaller businesses to use because it's simpler to set up
PVS we have more scalability - Though technically there is no limit in provisioning VM's	It can take a long time to define multiple storage locations in MCS because the read only copy of the image needs to be copied out to each storage location (each volumes) defined in Host connection
PVS has better versioning	No versioning concept
PVS allows for instant rollback by simply assigning the previous working version	Need to copy the previous snapshot in order to rollback which is not easy compare to PVS
PVS allows support of physical machines - Stream Physical	No Streaming

Ref:

- MCS or PVS, what should I be using? - https://www.citrix.com/blogs/2011/02/17/mcs-or-pvs-what-should-i-be-using/?_ga=2.257472075.1391456941.1500362834-1668697678.1498637089
- https://support.citrix.com/article/CTX218082
- https://www.citrix.com/blogs/2016/06/28/provisioning-services-or-machine-creation-services-2016-edition/?_ga=2.21544315.1391456941.1500362834-1668697678.1498637089

What is the impact if License Server not available to PVS Farm?

You need one license server per Provisioning Services farm. Provisioning Services servers must be connected to the license server to operate successfully:

When Provisioning Services is in a grace period, administrators are notified through warning messages in the Provisioning Services console. When a grace period expires, all target devices are shut down

Provisioning Services does not work out-of-the-box. You must use the most recent version of the Citrix License server to get the latest features. When you upgrade an existing environment to the newest version of Provisioning Services, you must also upgrade to the latest version of the licensing server or the product license will enter a **30-day grace period** and new product features will be unavailable.

Licensing grace periods

There are two types of grace period:

- **Out-of-box grace period** is **30 days (720 hours).** Initial installation of the licensing server provides startup licenses for all Citrix products. Startup licenses expire after 30 days. The 30-day countdown begins when the product prompts you for the startup license for the first time. Provisioning Services product licenses must be installed during this period. A startup license for a Citrix product is voided if a license for that product is installed, regardless of whether it is valid or invalid.
- **License server connectivity outage grace period** is 30 days (720 hours). If connectivity to the Citrix License Server is lost, Provisioning Services continues to provision systems for 30 days.

P.S: Applicable from 7.8
Ref:
https://docs.citrix.com/en-us/provisioning/7-8/install/pvs-license.html

General License Questions

Q: *How do Provisioning Services clients communicate with the Citrix license server?*

A: Clients do not communicate with the Citrix License Server.

Q: *What happens to the booting clients if there is no License Server installed or if there are not enough licenses available in the pool?*

A: On startup, clients receive a five-minute shutdown-warning message.

Q: *What happens to the booting client if the Citrix License Server is installed but there is no license key registered?*

A: The clients use the out-of-box grace period.

Q: *What is the Out-Of-Box Grace (OOBG) period?*

A: The out-of-box grace period takes place when the Citrix License Server is installed but no licenses are registered. It lasts for 96 hours from installation.

Q: *What happens if the OOBG grace period ends and the clients restart?*

A: Upon restarting, clients receive a five-minute shutdown warning message.

Q*: If the Citrix License Server loses its connection to the Provisioning Server, what happens to the active clients?*

A: If Provisioning Services 4.5.x with a hotfix and Provisioning Services 5.x are being used, clients keep working for an unlimited period of time until a restart occurs.

Q: *If the Citrix License Server loses its connection to the Provisioning Server, what happens to the clients that start up after this event?*

A: It goes to grace period but still works without getting 5-minute shutdown message.

Ref:
https://support.citrix.com/article/CTX117378

How many disk is used by VM when it created by MCS? Explain file significance.

Each VM created by MCS is given at minimum two disks upon creation.

- Disk0 = Delta or Differencing (Diff) Disk (contains the OS as copied from the Master Base Image)
- Disk1 = Identity Disk (16MB - contains Active Directory data for each VM)

As the product has evolved, additional disks may be added to satisfy certain use cases and feature consumption. For example:

- **Personal vDisk** (Feature which basically provides end users with the ability to install applications without admin intervention on a separate disk attached to the VM)

- **AppDisk** (Feature which provides the ability to attach application only disks to VMs primarily for Server OS Catalogs)
- New **MCS Storage Optimization** feature which creates a write cache style disk for each VM
- MCS added the ability to use **full clones** as opposed to the Delta disk scenario described above.

Hypervisor features may also enter into the equation. For example:

- **XenServer intelli-cache Feature** (creates a Read Disk on local storage for each XenServer to save on IOPS against the master image which may be help on the shared storage location).

Ref:
Machine Creation Services (MCS) Storage Considerations - https://support.citrix.com/article/CTX218082

What is the impact if PVS database is not reachable? How to enable Offline database

The Offline Database Support option allows Provisioning Servers to use a snapshot of the Provisioning Services database in the event that the connection to the database is lost.

Note: This option is disabled by default and is only recommended for use with a stable farm running in production. It is not recommended when running an evaluation environment or when reconfiguring farm components 'on the fly'. Only a farm administrator can set this option.

When offline database support is enabled on the farm, a snapshot of the database is created and initialized at server startup. It is then continually updated by the Stream Process. If the database becomes unavailable, the Stream Process uses the snapshot to get information about the Provisioning Server and the target devices available to the server; this allows Provisioning Servers and target devices to remain operational. *However, when the database is offline, Provisioning Services management functions and the Console become unavailable*.

When the database connection becomes available, the Stream Process synchronizes any Provisioning Server or target device status changes made to the snapshot, back to the database.

Considerations

The following features, options, and processes remain unavailable when the database connection is lost, regardless if the Offline Database Support option is enabled:

- Auto Add target devices
- vDisk updates
- vDisk creation
- Active Directory password changes
- Stream Process startup

- Image Update service
- Management functions; PowerShell, MCLI, SoapServer and the Console

Enabling Offline Database Support

To enable the Offline Database Support option

1. In the Console tree, right-click on the Farm, then select Properties. The Farm Properties dialog appears.
2. On the Options tab, check the checkbox next to Offline Database Support.
3. Restart Stream services.

Ref: http://docs.citrix.com/en-us/provisioning/7-1/pvs-ha-wrapper/pvs-ha-offline-db-config.html

What are the different methods to create Target devices? What is Auto-Add feature in PVS? In which scenario it will be useful.

Auto-Add is one of the method to create new target device entries in the Provisioning Services database, below are the different methods to create a target device:

- Using the Console to Manually Create Target Device Entries
- Using Auto-add to Create Target Device Entries
- Importing Target Device Entries

Using the Console to Manually Create Target Device Entries

1. In the Console, right-click on the Device Collection where this target device is to become a member, then select the Create Device menu option. The Create Device dialog appears.
2. Type a name, description, and the MAC address for this target device in the appropriate text boxes.
   Note: If the target device is a domain member, use the same name as in the Windows domain. When the target device boots from the vDisk, the machine name of the device becomes the name entered. For more information about target devices and Active Directory or NT 4.0 domains, refer to “Enabling Automatic Password Management”
3. Optionally, if a collection template exists for this collection, you have the option to enable the checkbox next to Apply the collection template to this new device.
4. Click the Add device button. The target device inherits all the template properties except for the target device name and MAC address.
5. Click OK to close the dialog box. The target device is created and assigned to a vDisk

Importing Target Devices Entries

Target device entries can be imported into any device collection from a .csv file. The imported target devices can then inherit the properties of the template target device that is associated with that collection. For more details, refer to Importing Target Devices into Collections.

Using the Auto-Add Wizard

The Auto-Add Wizard automates the configuration of rules for automatically adding new target devices to the Provisioning Services database using the Auto-Add feature.

The Auto-Add Wizard can be started at the Farm, Site, Collection or Device level. When started at a level lower than Farm, the wizard uses that choice as the default choice. For example, if it is started on a particular target device, it will:

- Select the Site for that Device as the Default Site choice in the combo-box.
- Select the Collection for that Device as the Default Collection choice in the combo-box.
- Select that Device as the Template Device choice in the combo-box.

The wizard displays each page with choices pre-selected based on the location that the Auto-Add Wizard was started from.

A Farm Admininistrator has the ability to turn Auto-Add on or off and to select the default Site.

A Site Admininistrator only has the ability to select the default site if the current default site is a site in which that administrator is the Site Administrator. If the Site Administrator is not the Administrator of the currently selected default Site, then that administrator can only configure the sites they has access to.

To configure Auto-Add settings (the default collection of a site, template device for the default collection and target device naming rules):

Steps: http://docs.citrix.com/en-us/provisioning/7-1/pvs-target-device-wrapper/pvs-target-database-add/pvs-target-auto-add-wizard.html

How can you enable Auditing for on PVS Farm?

Provisioning Services provides an auditing tool that records configuration actions on components within the Provisioning Services farm, to the Provisioning Services database. This provides administrators with a way to troubleshoot and monitor recent changes that might impact system performance and behavior.

Auditing is off by default. To enable it:

1. In the Console tree, right-click on the farm, then select the farm Properties menu option.
2. On the Options tab, under Auditing, check the Enable auditing checkbox.

Ref:

https://docs.citrix.com/en-us/provisioning/7-6/pvs-audit-wrapper.html

Difference between Private & Standard vdisk Access modes

Use the Console to select from the following vDisk access modes:

- Private Image – Select this mode if a vDisk is only used by a single target device (read/write access is enabled).
- Standard Image – Select this mode if a vDisk is shared by multiple target devices (write-cache options enabled).

Write Cache option enabled if Standard disk access mode enabled cache required during read access mode.

Types of vdisk load balances.

To achieve optimum server and target device performance within a highly available network configuration, enable load balancing for each vDisk.

1. Right-click on the vDisk in the Console, then select the Load balancing... menu option. ThevDisk Load Balancing dialog appears.
2. After enabling load balancing for the vDisk, the following additional load balancing algorithm customizations can be set:
   - **Subnet Affinity** – When assigning the server and NIC combination to use to provide this vDisk to target devices, select from the following subnet settings:
     - **None** – ignore subnets; uses least busy server. This is the *default setting*.
     - **Best Effort** – use 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.
     - **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.
   - **Rebalance Enabled using Trigger Percent** – Enable to rebalance the number of target devices on each server in the event that the trigger percent is exceeded. When enabled, Provisioning Services checks the trigger percent on each server approximately every ten

minutes. For example: If the trigger percent on this vDisk is set to 25%, rebalancing occurs within ten minutes if this server has 25% more load in comparison to other servers that can provide this vDisk.

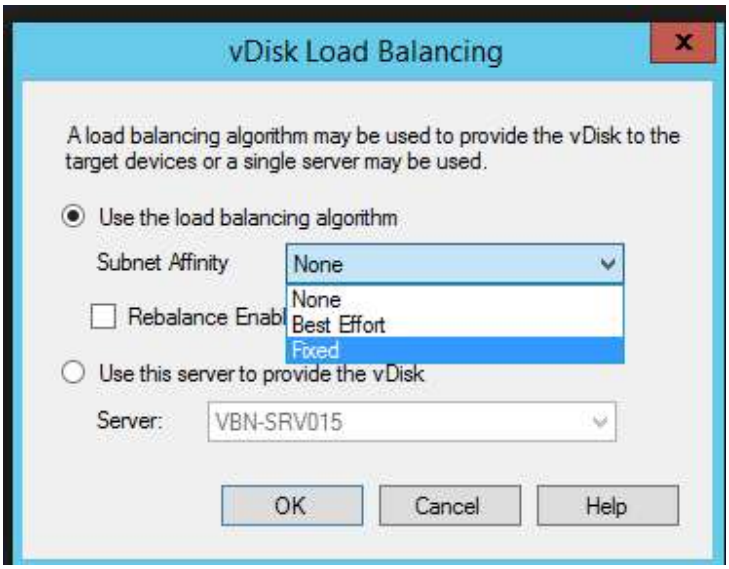

Ref:
http://docs.citrix.com/en-us/provisioning/7-1/pvs-vdisks-lifecycle-wrapper/pvs-vdisks-load-balancing.html
https://docs.citrix.com/en-us/provisioning/7-6/pvs-server-wrapper/pvs-server-devices-balance.html

How your PVS servers connected with storage? How to connect shared storage to PVS Server and procedure? What is the impact for existing user sessions if any one of PVS Server goes down? Is connections moved automatically? If yes explain.

Multiple Provisioning Servers can access the same physical files located on shared storage, which allows a target device to establish a connection on an alternate Provisioning Server if the connection to the active Provisioning Server is interrupted for any reason.

Or Multiple Provisioning Servers can access the same physical files located on local storage but still can provide HA if both joined in same farm

Provisioning Services supports various shared-storage configurations. The steps for configuring for highly available storage in the network varies depending on shared-storage configuration

Using Provisioning Services 5.1 you can use a SAN without a shared file system in some instances (you can have "read-only" vDisk storage). The desired boot modes for PVS target devices are important when using this feature since Provisioning Services ***only allows read-only shared access to the SAN LUN(s***)

Procedure to connect shared storage to PVS Server

- You need to install Microsoft iSCSI initiator on all PVS servers that access the SAN.
- It is only for Standard Image mode, Private Image mode is not supported.
- If cache is located on server disk, a separate shared storage location that has read-write access is needed for write cache file
- Post receiving LUN from storage, present LUN to PVS server either iscsi or HBA or VHBA.
- Format the drive and assign the drive letter,
- Copy all the vDisk image files (.vhd) and associated properties files (.pvp) to the volume, no need to copy the Lock files. Before you copy the files, make sure all properties for the vDisks that will reside on the volume are set correctly (including High Availability in PVS console).
- Next step is to make the volume read-only. You can use diskpart.exe , select disk and run command "*attributes volume set readonly*"
- Now mount the same volume to other provisioning servers

A target device does not experience any disruption in service or loss of data when failover occurs.

Note: Provisioning Services does not support the high availability of vDisks on local storage that are in Private Image mode or that are currently in maintenance (read/write enabled)

Ref:
http://docs.citrix.com/en-us/provisioning/7-8/managing-high-availability/pvs-ha-option-intro.html
Provisioning Services Read-only vDisk Storage
https://www.citrix.com/blogs/2009/09/17/provisioning-services-read-only-vdisk-storage/
Placing your vDisks on a Storage Area Network (SAN)
https://www.citrix.com/blogs/2009/08/03/provisioning-services-high-availability-considerations-part-4/

What is the difference between XenDesktop Setup Wizard and Streamed VM Setup wizard? Name the scenarios where these methods will use.

In XenDesktop setup wizard, it guides same as creating catalog, adding number of vms, create ad accounts or not,what will be the new vm name, what is the new vms boot order, PXE or BDM etc etc.. same process as creating catalog. It creates catalog in xendesktop and manages those vms.

XenDesktop Setup Wizard is the newer machine deployment option offered by PVS and should in most cases be your method of choice when deploying extra target VMs to your Citrix environment.

The XenDesktop Setup Wizard creates the Delivery Group and imports the created VDA's in to Citrix for you, automating part of the process

In Streamed VM Setup wizard, it just creates multiple vms, and adds them to pvs so that they can be monitored/managed only from pvs console. If needed, you have to manually add them to xendesktop by creating new catalog

When would you use Streamed VM Setup Wizard instead? When deploying machines that are not going to be part of your Citrix environment. PVS does not just have to stream Citrix related VM's. It can also stream non-Citrix related VM's

Ref:
http://www.kylewise.net/citrix/provisioning-services/pvs-7-8-streamed-vm-setup-wizard/
http://www.jgspiers.com/provisioning-services-xendesktop-setup-wizard/
https://docs.citrix.com/en-us/provisioning/7-6/pvs-vm-wizard-using.html
https://docs.citrix.com/en-us/provisioning/7-9/xendesktop-setup-wizard.html

Is PVS will be used only for Citrix Xen Desktop or any other environments?

Yes, you can use for non citrix environment where streaming of vdisks or physical servers required

Is PVS can be streamed to Physical Servers? In which Scenario it will used.

Yes, it can be used to stream physical servers too.. There may be scenarios where applications require high CPU, RAM & Multimedia applications (Graphic Intensive) or License

Do PVS Target devices support Network teaming? if yes what are the options available? How many Network cards required for PVS target devices

Provisioning Services provides the ability to run redundant networks between the servers and the target devices

Multiple NICs on the target device may configured into a virtual team by using Manufacturer's NIC teaming drivers, or into a failover group while installing Provisioning Target device software

Tip

When a machine powers up, the BIOS goes through the list of available boot devices and the boot order of those devices. Boot devices can include multiple PXE-enabled NICs. Provisioning Services uses the first NIC in the list as the primary boot NIC. The primary boot NIC's MAC address is used as the lookup key for the target device record in the database. If the primary boot NIC is not available at boot time, Provisioning Services will not be able to locate the target device record in the database (a non-primary NIC may be able to just process the PXE boot phase). Although a workaround would be to add a separate target device entry for each NIC on each system, and then maintain synchronization for all

entries, it is not recommended (unless the successful startup of a system is considered as critical as the continued operation of the system that is already running).

Target Device should have
NIC1 – Legacy NIC for PXE boot
NIC2 ->Synthetic NIC in streaming subnet
NIC3 ->NIC for Production (Need to check)

HA on NIC level on my target devices have below options

1) a virtual team by using Manufacturer's NIC teaming drivers (Windows NIC Teaming inBox)
2) failover group using the Provisioning Services NIC failover feature

Requirements and considerations for Provisioning Services NIC failover

Provisioning Services supports NIC failover for vDisks in either Standard and Private Image Mode.

- ***The PXE boot NIC is considered the primary target device MAC address***, which is stored in the Provisioning Services database.
- ***You define the failover group of NICs when you run the Provisioning Services target device installer on the Master Target Device***. If the machine has more than one NIC, the user is prompted to select the NICs in which to bind. Select all the NICs that participate in NIC failover. Alternatively, in Provisioning Services 5.1 or later, run bindcfg.exe, to selectively bind NICs post installation. bindcfg.exe is located in the product installation directory, .
- A target device will only failover to NICs that are in the same subnet as the PXE boot NIC.
- Teaming of multi-port network interfaces is not supported with Provisioning Services whereas multi network interface cards is supported
- In the event that the physical layer fails, such as when a network cable is disconnected, the target device fails over to the next available NIC. The failover timing is essentially instantaneous.
- The NIC failover feature and Provisioning Services HA feature compliment each other providing network layer failover support. If a failure occurs in the higher network layer, the target device fails over to the next Provisioning Server subject to HA rules.
- The next available NIC from the failover group is used should the NIC fail and the target device reboots. NICs must be PXE capable and PXE enabled.
- If a virtual NIC (teamed NICs) is inserted into the failover group, the vDisk becomes limited to Private Image Mode. This is a limitation imposed by NIC teaming drivers.
- **By default, Provisioning Services automatically switches from legacy Hyper-V NICs to synthetic NICs if both exist in the same subnet**. To disable the default behavior (allowing for the use of legacy HyperV NICS even if synthetic NICs exist), edit the target device's registry settings:[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\BNIStack\Parameters] DisableHyperVLegacyNic"=dword:00000000
- Load balancing is not supported in the NIC failover implementation.

Requirements and considerations for manufacturer's NIC teaming

When configuring NIC teaming, consider the following requirements:

- Provisioning Services supports Broadcom and Intel NIC teaming drivers. A vDisk that is built after configuring NIC teaming can run Standard or Private Image Mode. Broadcom NIC Teaming Drivers v9.52 and 10.24b are not compatible with Provisioning Services target device drivers.
- Teaming of multi-port network interfaces is not supported with Provisioning Services whereas multi network interface cards is supported
- Multi-NIC is supported for XenDesktop Private virtual machine desktops. Using the wizard, Provisioning Services allows you to select the network to associate with the Provisioning Services NIC (NIC 0). The Delivery Controller provides the list of associated network resources for host connections.
- The target device operating system must be a server-class operating system, such as Microsoft Windows 2008.
- The new virtual team NIC MAC address has to match the physical NIC that performs the PXE boot.
- Microsoft Windows Server 2012 built-in NIC teaming or OEM NIC teaming software should be installed and configured **prior** to the Target Device software.
- Configure NIC teaming and verify that the selected teaming mode is expected by the application and the network topology. It should expose at least one virtual team NIC to the operating system.
- When provisioning machines to a SCVMM server, the XenDesktop Setup wizard automatically changes the network configuration of both the first legacy NIC and the second synthetic NIC.
- During the Master target device installation process, Provisioning Services target device client drivers need to bind to the new virtual team NIC MAC address. If all physical NICs have been teamed up to a single virtual NIC, the Provisioning Services installer automatically chooses the virtual NIC silently, without prompting.
- If changes are required, Provisioning Services Target Device software must be uninstalled before making changes to the teaming configuration, then reinstalled after changes are complete. Changes to teaming configurations on a Master target device that has target device software installed, may result in unpredictable behavior.
- When installing Provisioning Services target device software on NT6.x systems within a multi-NIC environment, all available NICs can be used. **Therefore bindcfg.exe is no longer required and no longer installed with target device software**.

Ref:
https://docs.citrix.com/en-us/provisioning/7-1/pvs-network-components-wrapper/pvs-nics/pvs-nics-requirements-failover.html
http://docs.citrix.com/en-us/provisioning/7-1/pvs-network-components-wrapper/pvs-nics.html

What is the significance of Threads per port on PVS Server Advanced Properties? Additionally, ask any one of advance properties (Buffers per Thread, Server Cache Timeout,MTU, pacing etc..)

Provisioning Server Properties dialog allows you to modify Provisioning Server configuration settings.

Threads per port — Number of threads in the thread pool that service UDP packets received on a given UDP port. Between four and eight are reasonable settings. Larger numbers of threads allow more target device requests to be processed simultaneously, but is consumes more system resources

PVS is actually not using a standard threading model where each client gets its own port/thread much like a TFTP server does. Instead, PVS has a listener for each port that receives a request and dumps it on a port specific thread pool. The threads in the pool process each request, one per thread. If there are more threads than cores, the leftover threads simply block. Adding more threads than CPU logical core is not going to help on performance

The best Stream Process performance is attained when the threads per port is not greater than the number of cores available on the Provisioning Server.

For best performance, use the following formula: **# of ports x # of threads/port = max clients**

Example:
For a typical 1000 target devices per PVS server situation, the following configuration is recommended to be a baseline to start with:
3-6 vCPUs (I will explain in more details below), 32 GB RAM (depends on the vDisk numbers)
Streaming Port re-configured from 6910 to 6968 (default 6910 – 6930).
Threads per port set to match the vCPU number.
Leave the rest advanced options to be unchanged.

Note:
Increase the threads per port. The number of threads per port should match the number of vCPUs assigned to the server.

The Server Properties dialog includes the following tabs:

- General
- Network
- Stores
- Options
- Logging

General

Power Rating

A power rating is assigned to each server, which is then used when determining which server is least busy. The scale to use is defined by the administrator. For example in 1-10 scale , 2 is considered twice as powerful as a server with a rating of 1; therefore it would be assigned twice as many target devices.

Log events to the server's Window Event Log

Select this option if you want this Provisioning Server's events to be logged in the Windows Event log.

Advanced Server Properties
Server tab

Threads per port — Number of threads in the thread pool that service UDP packets received on a given UDP port. Between four and eight are reasonable settings. Larger numbers of threads allow more target device requests to be processed simultaneously, but is consumes more system resources.

Buffers per thread — Number of packet buffers allocated for every thread in a thread pool. The number of buffers per thread should be large enough to enable a single thread to read one IO transaction from a target device.

Server cache timeout — Every server writes status information periodically to the Provisioning Services database. This status information is time-stamped on every write. A server is considered 'Up' by other servers in the farm, if the status information in the database is newer than the Server cache timeout seconds(2sec)

Local and Remote Concurrent I/O limits — Controls the number of concurrent outstanding I/O transactions that can be sent to a given storage device. A storage device is defined as either a local drive letter (C: or D: for example) or as the base of a UNC path, for example \\ServerName

Network tab

Maximum transmission unit — Number of bytes that fit in a single UDP packet. For standard Ethernet, the default value is correct. If you are attempting to operate over a WAN, then a smaller value may be needed to prevent IP fragmentation. Provisioning Services(v7.13) currently does not support IP fragmentation and reassembly.

I/O burst size — The number of bytes that will be transmitted in a single read/write transaction before an ACK is sent from the server or device

Pacing tab

Boot pause seconds — The amount of time that the device will be told to pause if the Maximum devices booting limit has been reached

Maximum boot time — The amount of time a device will be considered in the booting state. Once a device starts to boot, the device will be considered booting until the Maximum boot time has elapsed for that device. After this period, it will no longer be considered booting (as far as boot pacing is concerned) even if the device has not actually finished booting. Maximum boot time can be thought of as a time limit per device for the booting state for boot pacing.

Maximum devices booting — The maximum number of devices a server allows to boot at one time before pausing new booting devices. The number of booting devices must drop below this limit before the server will allow more devices to boot.

vDisk creation pacing — Amount of pacing delay to introduce when creating a vDisk on this Provisioning Server. Larger values increase the vDisk creation time, but reduce Provisioning Server overhead to allow target devices that are running, to continue to run efficientl

Device tab

License timeout — Amount of time since last hearing from a target device to hold a license before releasing it for use by another target device. If a target device shuts down abnormally (loses power for example) its license is held for this long.

Network

IP Address

The IP addresses that the Stream Service should use for a target device to communicate with this Provisioning Server. When adding a new Provisioning Server, enter the valid IP address for the new server.

Ports
Enter the First and Last UDP port numbers to indicate a range of ports to be used by the Stream Service for target device communications.
Note: The minimum is five ports in a range. The default first port number is 6910 and the last port number is 6930 .

FYI

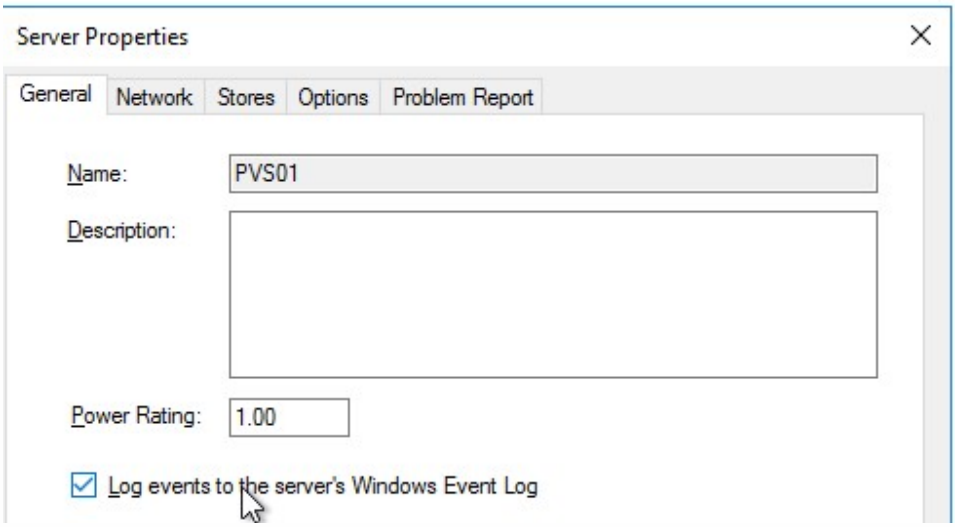

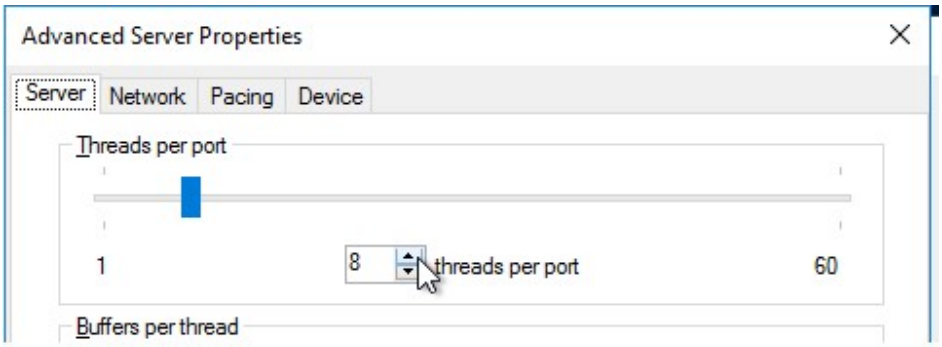

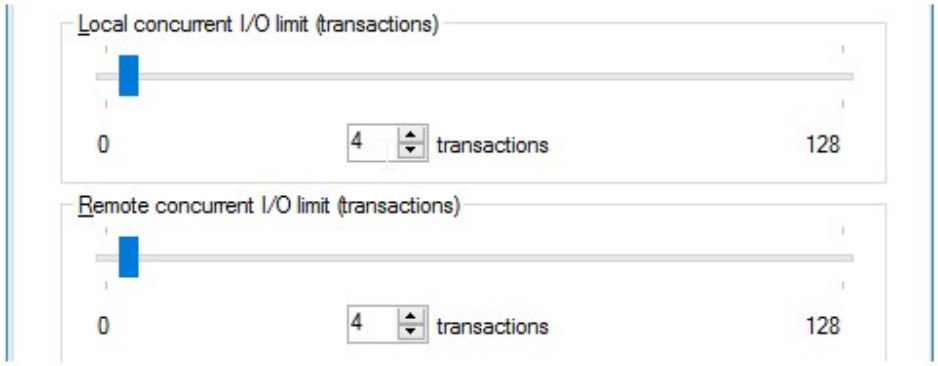

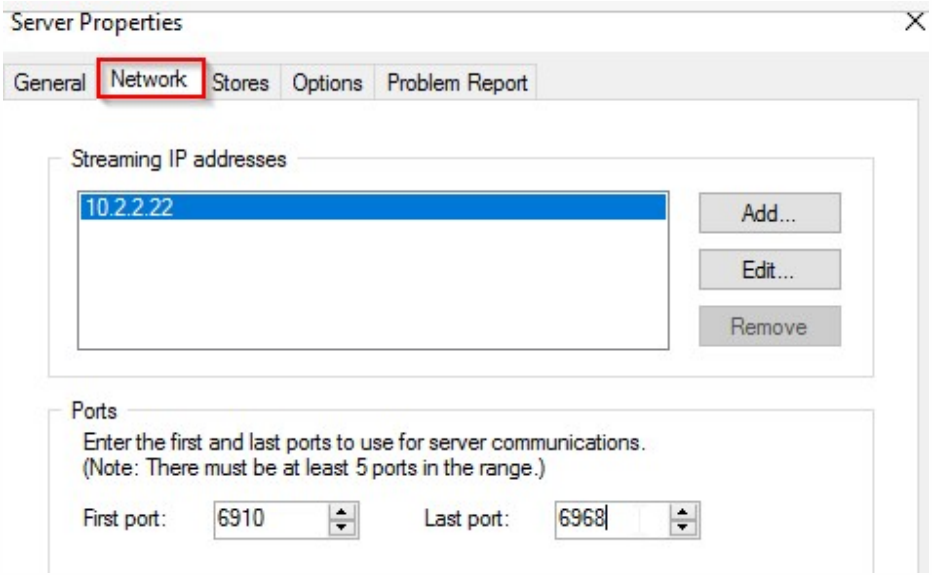

Ref:
https://www.citrix.com/blogs/2016/03/30/updated-guidance-on-pvs-ports-and-threads/

https://www.citrix.com/blogs/2011/07/11/pvs-secrets-part-3-ports-threads/
http://www.carlstalhood.com/provisioning-services-console-config/
https://docs.citrix.com/en-us/provisioning/7-13/managing-servers/server-properties.html

What are the logs available for Provisioning Servers and how can you enable logging for Target devices for troubleshooting? Name the Logs tools available for PVS.

Provisioning Services has been providing the ability to capture logs since the time of PVS version 4.5. Using the Log4cxx framework, PVS logging has been able to provide information about the goings on of the PVS console, the Soap Service, and the Stream Service. These logs have provided valuable historical data when trying to troubleshoot issues.

There have been many improvements to the logging capability and as the feature set for PVS grew Citrix introduced new logs as well as refined some existing logs. With the introduction of PVS 7.0 the logging format has been changed from Log4cxx to CDF. Citrix has long used CDF as a mainstay of our logging mechanism across many products

Log Tools available are CDFControl ,CDFmonitor(For continuous trace), CDF Analyzer, PVS Datacollector, ->All logs which are generated by this tool can be uploaded in to TAAS site (https://taas.citrix.com)

Enable Logging for PVS Server & Target Devices:

- Open the PVS Console > Sites > Site > Servers node->Right-click on the PVS server and select Properties.
- Select the Logging tab and set the Logging level to Debug.

For Target Device

- For target side logging right click any target in your device collection and click on properties.

In Provisioning Services all traces and logs are kept on the Provisioning Services Server, even for Target Devices (because Target devices are Readonly)

Default Log file location in all PVS servers are C:\ProgramData\Citrix\Provisioning Services\Log

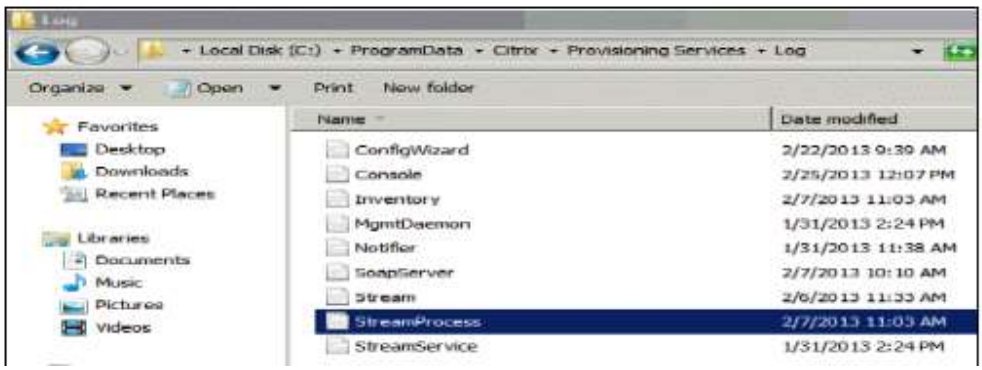

Stream_log.config- This file should not be edited manually. Logging levels should be set through the Console.
Stream.log -StreamProcess.exe, Manager.dll and Streamdb.dll all write to the Stream.log file.
MCLI.log - Writes MCLI logging information to MCLI.log.
SoapServer.log - Writes SoapServer logging information to Soapserver.log.
Console.log - Writes Console logging information to Console.log.
ConfigWizard.log - Writes Provisioning Server configuration logging information to ConfigWizard.log.

Ref:
CDF Control - https://support.citrix.com/article/CTX111961
CDF Monitor (For Continuous trace) - https://support.citrix.com/article/CTX138698
PVS Data Tools -> https://support.citrix.com/article/CTX136079

Why Write Cache required for vdisks? Can vdisks run without write cache? Type of cache available in Standard Mode? Explain Difference between Write caches.

The streamed image is read-only and changes, writes, are cached to a write cache location configured per streamed image. When using Provisioning Services, one of the most important things to ensure optimal performance is the write cache type.

The write cache includes data written by the target device. If data is written to the PVS server vDisk in a 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 following locations:

- Cache on device hard disk
- Cache in device RAM
- Cache on device RAM with overflow on hard disk
- Cache on server

When the target device is booted, write cache information is checked to determine the presence of the cache file. If the cache file is not present, the data is then read from the original vDisk file.

All current versions of PVS have the option for distributing write cache. It is called **Multiple Write Cache Paths**. The multiple write cache paths (for a store) option provides the capability of distributing the write cache files across multiple physical media. This feature helps to improve I/O throughput for heavily loaded servers.

Cache on server:

This write-cache type places the write-cache on the PVS server. By default, it is placed in the same location as the vDisk, but a different path can be specified.

This write-cache type provides poor performance when compared to other write cache types, limits high availability configurations, and should almost never be used in a virtual environment. This option was typically only used when streaming to physical end points or thin clients that are diskless.

Cache on device's hard drive:

This write-cache type creates a write-cache file (.vdiskcache) on the target devices' hard drive. This cache type has been our leading Citrix best practice for environments to date and most of our deployments use this write-cache type as it provides the best balance between cost and performance. To achieve the highest throughput to the write-cache drive, Intermediate Buffering should almost always be used (caution should be used with target devices hosted on Hyper-V where we have occasionally seen adverse effects). By default this feature is disabled (CTX126042.)

Cache in device RAM:

This write-cache type reserves a portion of the target device's memory(in non-paged pool memory) for the write cache, meaning that whatever portion of RAM is used for write-cache is not available to the operating system. The amount of memory reserved for write-cache is specified in the vDisk properties. This option provides better throughput, better response times, and higher IOPS for write-cache than the previous types because it writes to memory rather than disk.

There are some challenges like, firstly if RAM is not sufficient then there are high chances of blue screen errors hence plenty if RAM to purchase which is costly, Second, if there is a need to store persistent settings or data such as event logs, a hard drive will still be required on each target. On the flip side, this hard disk will not be as large or use as many IOPS as when using "Cache on device's hard drive"

Cache on device RAM with overflow on hard disk:

This is a new write-cache type and is basically a combination of the previous two, but with a different underlying architecture. It provides a write-cache buffer in memory and the overflow is written to disk. However, the way that memory and disk are used is different than with "Cache in device RAM" and "Cache in device's hard drive" respectively. This is how it works:

- Just as before, the buffer size is specified in the vDisk properties. By default, the buffer is set to 64 MB but can be set to any size.
- Rather than reserving a portion of the device's memory, the cache is mapped to Non-paged pool memory and used as needed, and the memory is given back to the system if the system needs it.
- On the hard drive, instead of using the old ".vdiskcache" file, a VHDX (vdiskdif.vhdx) file is used.
- On startup, the VHDX file is created and is 4 MB due to the VHDX header.
- Data is written to the buffer in memory first. Once the buffer is full, "stale" data is flushed to disk.
- Data is written to the VHDX in 2 MB blocks, instead of 4 KB blocks as before. This will cause the write-cache file to grow faster in the beginning than the old ".vdiskcache" cache file. However, over time, the total space consumed by this new format will not be significantly larger as data will eventually back fill into the 2 MB blocks that are reserved.

Note:

- RAM size recommended in cache is for desktop operating systems start with 25**6-512MB** and for server operating systems start with **2-4GB**
- This write-cache type is only available for Windows 7/2008 R2 and later.
- If the pagefile is too big, it fails back to server caching. Set a fixed size for the pagefile that is smaller than the cache disk.

Write Cache Size

Write Cache File Name

PVS has had three different cache names:

- **.vdiskCache** is Legacy Ardence format (5 .x and before not supported anymore, you can delete this if your target software is running latest, this cache was optimized for size)
- **.vdiskdif.vhd** is legacy hard drive cache (6.0 and above local hard drive cache, used standard 1mb sector size and is larger than the legacy cache but worked better with storage and was incrementally faster than Legacy Ardence format)
- **vdiskdiff.vhdx** is Ram cache with overflow (7.1.4 and above RAM cache with overflow, 2 mb sectors larger than vhd but much faster and more compatible with storage)

Deleting cache on a difference disk

The Delete Cache from Selected Device(s)... context menu option allows you to manually delete cache on a difference disk. The option is only available if the vDisk cache mode is set to Server Persistent Cache.

Note: Write cache on a Difference Disk is not automatically deleted if that file becomes invalid. Files marked as invalid should periodically be deleted manually.

To delete a cache on a Difference Disk:

In the Console, right-click on the vDisk that is associated with difference disk files to delete. Select the Delete Cache from Selected Device(s) menu option. The Delete Cache for Devices dialog appears.

Check each target device box for which the cache should be deleted, or click Select all to delete all cache files associated with this vDisk.

Click Delete to delete the cache files from the server

Ref:

https://www.citrix.com/blogs/2011/10/06/pvs-write-cache-sizing-considerations/
https://www.citrix.com/blogs/2014/04/18/turbo-charging-your-iops-with-the-new-pvs-cache-in-ram-with-disk-overflow-feature-part-one/?_ga=2.105753024.234697361.1500524189-1668697678.1498637089
http://docs.citrix.com/en-us/provisioning/7-8/managing-vdisks/pvs-write-cache.html
https://www.citrix.com/blogs/2015/01/19/size-matters-pvs-ram-cache-overflow-sizing/

What is Personal vdisk (PvD) & Why it required? Is it possible to connect pvdisk to Server OS version, if yes explain?

The personal vDisk feature holds the single image management of pooled and streamed desktops while allowing users to install applications and change their desktop setting

Users lose their customization and personal applications when the administrator changes the master image, deployments using personal vDisks retain their changes. This means administrators can easily and centrally manage their master images while providing users with a customized and personalized desktop experience.

Personal vDisks provide this separation by redirecting all changes made on the user's VM to a separate disk (the personal vDisk), which is attached to the user's VM.

Personal vDisks have two parts, which use different drive letters and are by **default equally sized**:

- **User profile** - This contains user data, documents, and the user profile. By default this uses drive P: but you can choose a different drive letter when you create a catalog with machines using personal vDisks. The drive used also depends on the EnableUserProfileRedirection setting.
- **Virtual Hard Disk (.vhd) file** - This contains all other items, for example applications installed in C:\Program Files. This part *is not displayed in Windows Explorer* and, since Version 5.6.7, does not require a drive letter.

Physically, you store personal vDisks on the hypervisor but they do not have to be in the same location as other disks attached to the virtual desktop. This can lower the cost of personal vDisk storage.

Back up personal vDisks regularly using any preferred method. The vDisks are standard volumes in a hypervisor's storage tier, so you can back them up, just like any other volume

FAQ's

- PvD is a feature available in all editions of XenDesktop 5.6 and 7.x.
- PvD compatible with XenDesktop and applicable to XenAPP
- PvD works in both MCS & PVS. PvD can only be configured within a virtual target device. A physical target device is not support.
- There can only be one PvD per Virtual Machine. The PvD is assigned to a Virtual Machine when building the catalog of desktops. The pool type for a PvD catalog is a pooled static, which the desktop is assigned to the user on first use.
- Actually, PvD is a 1:1 mapping to a Virtual Machine in a catalog, which is then assigned to the user on first use. A PvD is attached to a Virtual Machine assigned to the user. The administrator can move a PvD to a new virtual machine in a recovery situation
- If you create a catalog for pooled with PvD, it does not mean that the user is always required to be assigned to that Virtual Machine.The base image is still shared and updated across the pool. However, once the user makes an initial connection to a Virtual Machine, the Virtual Machine is kept assigned to the use in pooled catalog
- *PvD only supported on Desktop Operating Systems*

Ref

FAQ: Personal vDisk in XenDesktop **:** **https://support.citrix.com/article/CTX131553**
https://docs.citrix.com/en-us/provisioning/7-8/configure-personal-vdisk.html
https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-6/cds-pvd-intro.html

Why Vdisk version required (or how vdisks updated) and how many versions are recommended? What files are created during vdisk version?

Versioning simplifies vDisk update and management tasks, providing a more flexible and robust approach to managing vDisks.

Differencing disks(AVHD) 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. Each time a vDisk is put into Maintenance Mode a new version of the VHD differencing disk is created and the file name is numerically incremented

Numerous differencing disk versions can exist for a vDisk. Device access to a particular version, or the ability to make updates to that version, depends on that versions Access mode setting and the device Type

A version's Access mode is managed on the vDisk Versioning Dialog. New versions of a vDisk are generally promoted from Maintenance to Test and then into Production. Access mode options include:

Maintenance – new read/write difference disk version that is only available to the first Maintenance device that selects to boots from it in order to make updates.

Test – read-only version used for test purposes and only available to Test or Maintenance devices.

Default – read-only version that is bootable by all device types. If the Boot production devices from version is set to Newest released, then the latest released production version is marked with a green checkmark and the status is set to Default.

Override – read-only version that is bootable by all device types. If a specific version is selected from the Boot production devices from version drop-down list, then that version is marked with a green checkmark and the access changes to Override.

Newest released – read-only version that is bootable by all devices. If a specific version is selected from the Boot production devices from version drop-down list, then that version is marked with a green checkmark and the access changes to Override.

Merging – a merge is occurring to this new version. This version is unavailable to all device types until the merge completes. After the merge completes, the status of the new version depends on the Access mode selected on the Mode to set the vDisk to after automatic merge drop-down list (Production, Maintenance, or Test). This Farm Properties setting is available on the vDisk Versions tab.

vDisk Versioning dialog

vDisk versioning is managed from the vDisk Versions dialog. To open the dialog, right-click on a vDisk in the Console, then select the Versions... menu option. The table that follows provides a general description of the vDisk Versions dialog

It is often necessary to update an existing vDisk so that the image contains the most current software and patches. Each time the vDisk is to be updated, a new version of that vDisk is created (VHD file) to capture the changes without changing the base vDisk image.

Updating a vDisk involves the following:

- Create a new version of the vDisk, manually or automatically.
- Boot the newly created version from a device (Maintenance device or Update device), make and save any changes to the vDisk, then shut-down the device.
- Promote the new version to Production.

The following illustrates the general promotion of a vDisk update:

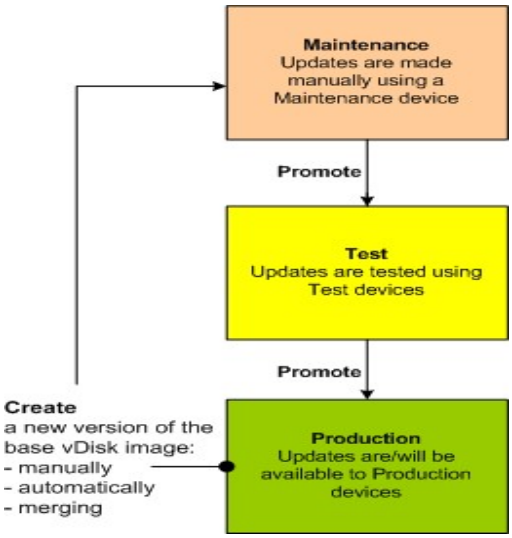

The availability of the updated version depends on the current promotion of that version (Maintenance, Test, or Production), and the type of device attempting to access it (Maintenance Device, Update Device, Test Device, or Production Device).

If updating a device that uses a personal vDisk image, ensure compatibility in your production environment using this procedure:

Note: Updating images for devices that use a personal vDisk, must be done on a virtual machine that does not have a personal vDisk attached. Otherwise, updates are saved to the personal vDisk image rather than the virtual machine image.

1. Create a new maintenance version of the vDisk.
2. Make any necessary updates to the maintenance version.
3. Promote the new maintenance version to test.
4. Boot the PvD test device, and then verify updates were made.
5. Promote the test version to production.

Update Scenarios

The following vDisk update scenarios are supported:

- **Manual Update** – An administrator may choose to update a vDisk manually by creating a new version of that vDisk, and then using a Maintenance device to capture updates to that version. Manual updates are initiated by selecting the New button on the vDisk Versions dialog.
- **Automated Update** – Creating automated updates saves administration time and physical resources. Updates are initiated on-demand or from a schedule and are configured using vDisk Update Management. If updating automatically, the Access column on the vDisk Versioning dialog displays that the newly created version is currently under maintenance.
    - **Note:** vDisk Update Management is intended for use with Standard Image Mode vDisks only. Private Image Mode vDisks can be updated using normal software distribution tool

procedures. Attempting to register a Private Image Mode vDisk for vDisk Update Management, or switching a vdisk that is already registered, will cause errors to occur.

- **Merge** – Merging VHD differencing disk files can save disk space and increase performance, depending on the merge option selected. A merge update is initiated manually by selecting the Merge button on the vDisk Versions dialog, or automatically when the maximum vDisk versions count is reached.

Merging VHD Differencing Disks

Merging VHD differencing disk files can save disk space and increase performance, depending on the merge method selected.

Citrix recommends merging vDisk versions to either a new base image or to a consolidated differencing disk. Each time the vDisk is versioned five times.

Merge methods include:

- Merging to a new base image
- Merging to a consolidated differencing disk

Note: A merge can only occur when no Maintenance version exists for this vDisk or when the vDisk is in Private Image mode. A merge starts from the top of the chain down to a base disk. A starting disk cannot be specified for the merge

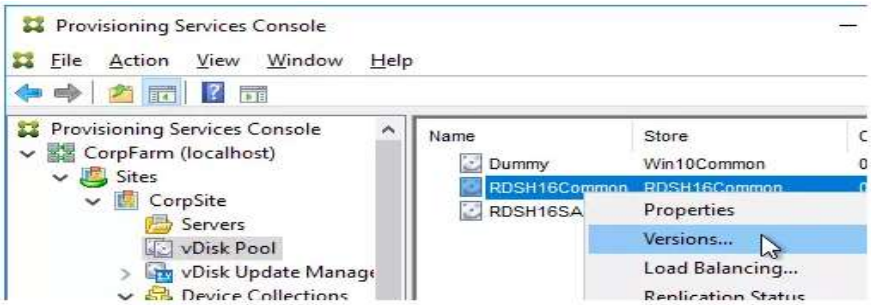

. In the *vDisk Versions* window, click **New**.

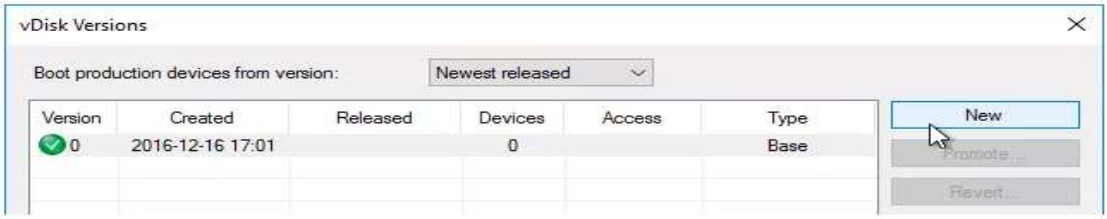

Notice that the *Access* is set to *Maintenance*. Click **Done**.

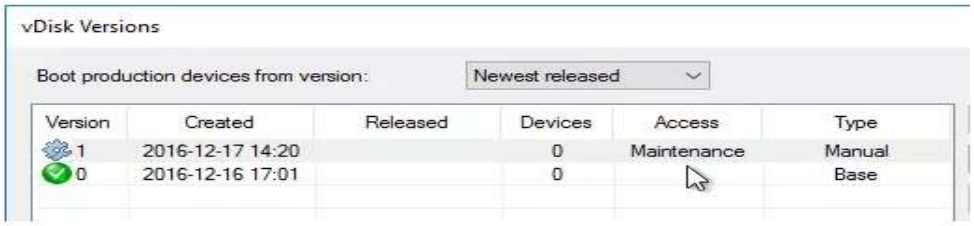

If you look at the physical location where the vDisks are stored, you'll see a new **.avhdx** file.

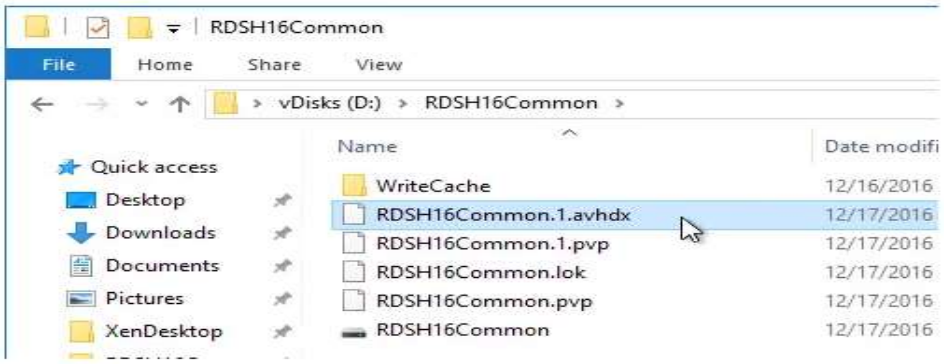

3. Go to the properties of an Updater Target Device, and change the *Type* to **Maintenance**. You'll use this Target Device to update the vDisk. Make sure this Target Device you are using for vDisk Updating is not in any Delivery Group so that users don't accidentally connect to it when it is powered on.

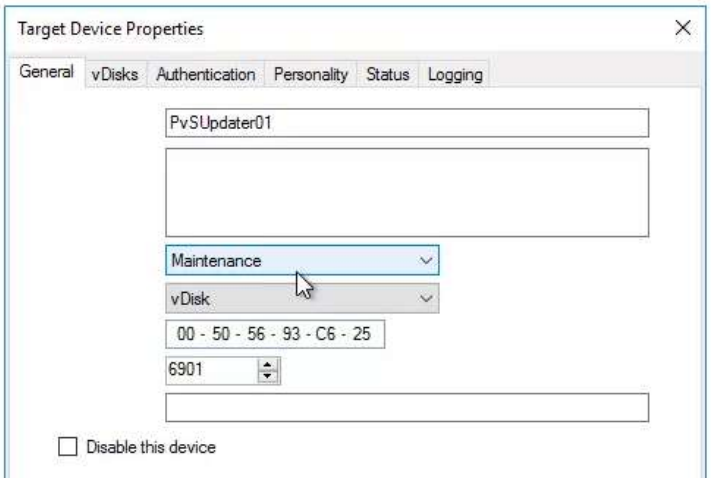

Ref:

http://docs.citrix.com/en-us/provisioning/7-12/managing-vdisks/pvs-vdisks-update-disk-types.html
http://docs.citrix.com/en-us/provisioning/7-6/pvs-vdisks-lifecycle-wrapper/pvs-vdisks-update-wrapper.html

http://docs.citrix.com/en-us/provisioning/7-6/pvs-vdisks-lifecycle-wrapper/pvs-vdisks-update-wrapper/pvs-vdisks-vhd-merge.html

http://www.carlstalhood.com/pvs-update-vdisk/#updater

How many files will be created for a vdisk and explain importance? How to Import & export vdisks and what files are required to replicate between PVS servers and replication process.

vDisks act as a hard disk for a target device and exist as disk image files on a Provisioning Server or on a shared storage device. A vDisk consists of a VHDX base image file, any associated properties files (.pvp), and if applicable, a chain of referenced VHDX differencing disks (.avhdx) for versions

.VHDX -> Contains Image File
.PVP -> It is properties file i.e. it contains vdisk properties (Version,Access mode, load balance etc..)
.AVHDX -> This file created when vdisk version is created
.LOK -> This is the LOCK file which contains device connected information -> when any device connects, that lock information is stored in this file

vDisks act as a hard disk for a target device and exist as disk image files on a Provisioning Server or on a shared storage device. A vDisk consists of a VHDX base image file, any associated properties files (.pvp), and if applicable, a chain of referenced VHDX differencing disks (.avhdx).

When creating a vDisk image file, keep the following facts in mind:

- You can create as many vDisk image files as needed, as long as you have enough space available on the Provisioning Server, or on the storage device containing the vDisk image files.
- vDisk files use FAT or NTFS file systems for Microsoft operating systems.
- Depending upon the file system used to store the vDisk, the maximum size of a VHDX file (vDisk) is 2 terabytes (NTFS) or 4096MB (FAT).
- A vDisk may be shared (Standard Image) by one or more target devices, or it can exist for only one target device to access (Private Image).

Provisioning Services allows for the exporting and importing of both versioned (having Difference disks with versions) and unversioned vDisks, from an existing store to a store in a different farm.
Note: If importing VHDs that were not exported using Provisioning Services, all differencing disks must first be merged to a base disk using third party tools, then the new VHD base disk can be imported.

Exporting vDisks

To export a vDisk:
Note: When deleting a vDisk that will be exported, be sure to export the vDisk first, then copy the resulting XML file to the new location before deleting it from the original location.

- Right-click on the vDisk in the Console, then select the Export menu option. The Export dialog appears.
- Select the version to export from the drop-down menu, then click OK. The manifest file is created in the Store.

Importing vDisks

A vDisk or vDisk chain of differencing VHD files can be imported into a store if:
The VHD being imported does not already exist in the store and both the highest version number of the VHD and associated manifest files match, and if the VHD chain includes a base image, and that base image version number matches the base image version in the manifest file.
The VHD does exist in the store but the imported version number in the associated manifest file is greater than the existing VHD version number.

Back in the days of Citrix Provisioning Services 5.x it was common practice to copy an existing vDisk mostly due to maintenance reasons in order to apply updates, install new software and so on and so forth

Aside from the fact that Citrix introduced with PVS 6.0 a more flexible and robust vDisk updating approach based on VHD chaining, the good old copy approach still works. In the case of an unversioned vDisk the procedure is as straightforward as in PVS 5.x (copy/label the according *.vhd and *.pvp files, then choose 'Add or Import Existing vDisk' in the context menu of the vDisk Store).

If you go the same path with a versioned vDisk (copy the complete VHD cain) you'll get an error message. Is it impossible to copy a versioned vDisk after all? No, you have two options: merge and export/import

Merging the VHD chain of the source vDisk seems to be the most obvious preparation step in order to copy the vDisk since it results in a single set of VHD/PVP files, meaning that afterwards you're able to copy that vDisk as simple as ever. But to merge means to commit oneself to a single source vDisk version. You'll lose flexibility

If you need to keep the source vDisk with all its versions as it is and want to get a copy anyway you should choose the export/import option. This approach is actually supposed to export and import both versioned and unversioned vDisks from an existing store to a store in a different PVS farm. However you can leverage it to get a copy in one and the same vDisk store as follows

- Copy the versioned vDisk including all *.vhd, *.ahvd, and *.pvp files. (Don't copy *.lok files).
- Rename or label the copied files as required.
- Open the PVS console, right click the source vDisk, and choose 'Export vDisk…'. A dialog appears.
- In the export dialog, choose the the latest version in the drop-down menu named 'Export versions starting at', then click OK. After a short delay the dialog closes. In the source vDisk's store you'll find a manifest file containing the entire information about the versions of that vDisk. The manifest file name matches with the name of the vDisk and it has a **.xml suffix**
- Rename the manifest file using the name of the copied vDisk.
- Open the manifest file in a text editor and accurately replace all references to the name of the source vDisk with the name of the copied vDisk. Double-check the changes, then save the file.
- Now, you should have a set of VHD/AVHD/PVP file and an XML file with the same base file name. And the manifest/XML refers to the new VHD and AVHD files.
- In the PVS console, right click the vDisk store, and choose 'Add or Import Existing vDisk…'. A dialog appears.
- In the import dialog, check the settings for Site, Store, and Server, then click Search. After a short delay the copied vDisk will be displayed.

- Check the vDisk, check/uncheck the load balancing option as wanted, then click Add. After a short delay a popup appears saying that the import was successful. Click OK, then click Close in the import dialog.

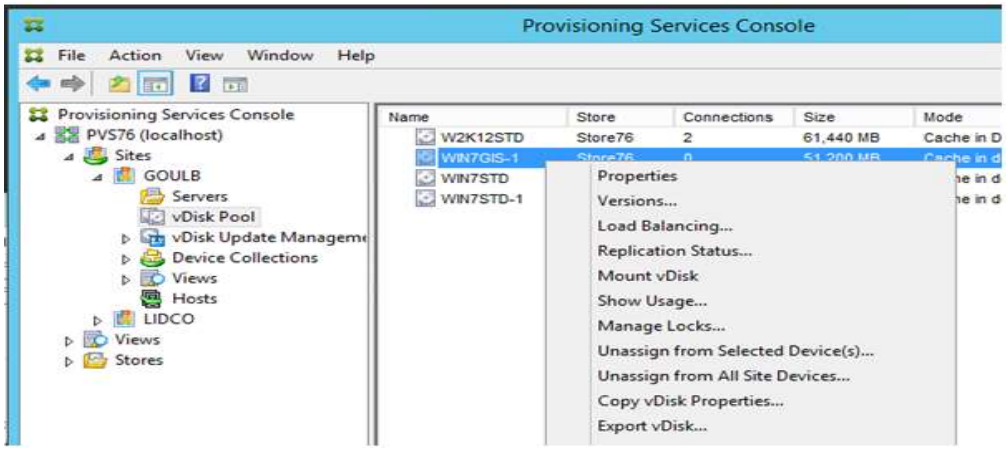

Export

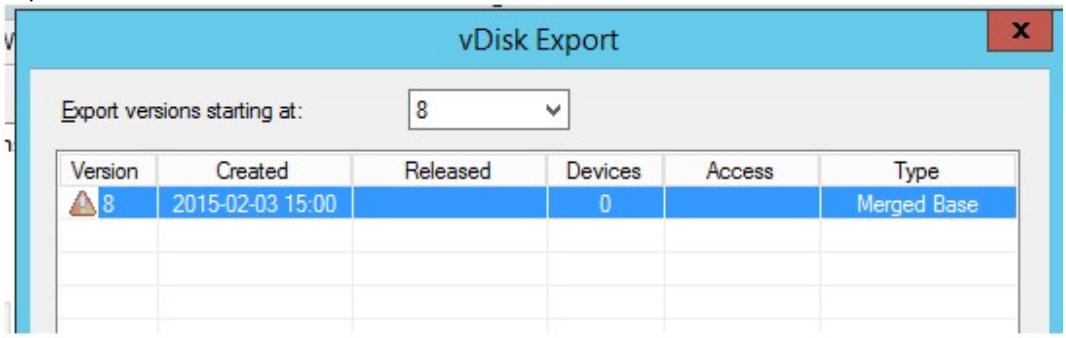

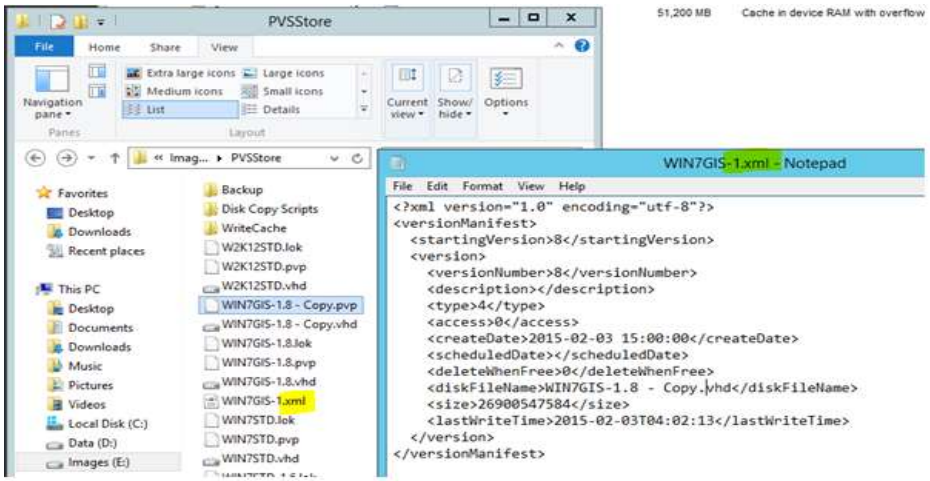

You should have a set of VHD/AVHD/PVP/XML files with the same file name. And an XML which refers to the new VHD files

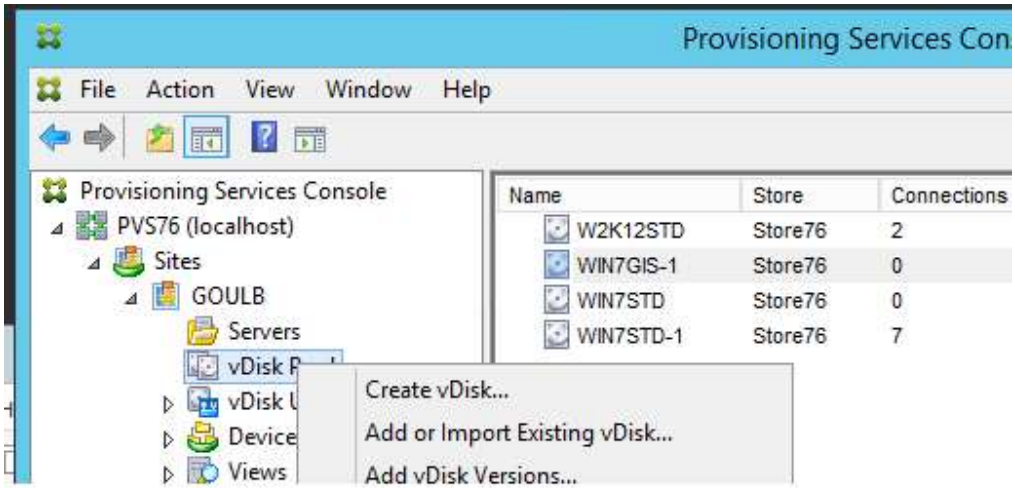

Ref:
http://docs.citrix.com/en-us/provisioning/7-13/managing-vdisks/vdisk-create.html

https://docs.citrix.com/en-us/provisioning/7-1/pvs-vdisks-lifecycle-wrapper/pvs-vdisks-versioning-import-export.html

How the Target Device as a template works in a collection?

A target device can be set as the template for new target devices that are added to a collection. A new target device inherits the properties from the template target device, which allows you to quickly add new devices to a collection

Note: Target devices that use personal vDisks are created and added to a collection when the XenDesktop Setup Wizard is run. If a target device template exists, it is ignored when the target device that uses a personal vDisk is added to the collection.

To set a target device as the template device for a collection, in the Console, right-click on the target device, then select Set device as template.

Note:

- Disable the target device that serves as the template to permit all target devices using this template to be added to the database, but not permit the target device to boot.
- Target devices receive a message requesting that they first contact the administrator before being allowed to boot.
- A 'T' appears in light blue on the device serving as the template. New target devices automatically have a name generated and all other properties will be taken from the default template target device. No user interaction is required.

I want to boot vdisk only from specific PVS Server? Is it possible? How to change?

Yes, it is possible to boot vdisk from specific PVS server

Right Click vdisk -> Load Balance -> Use Specific PVS server option

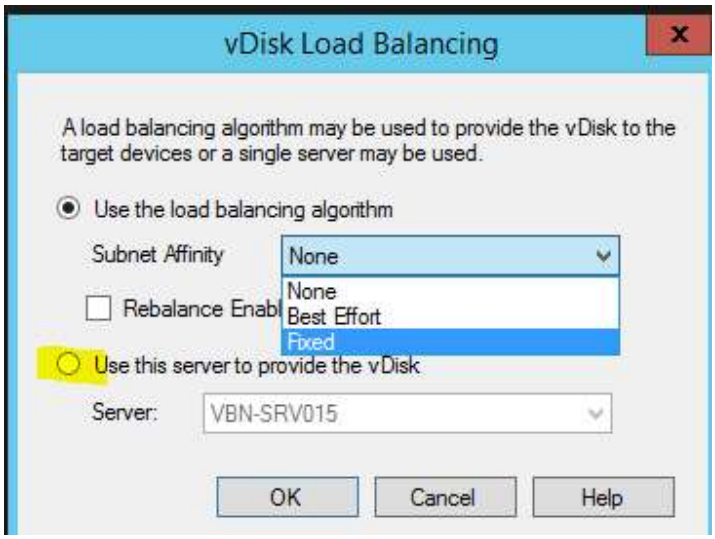

What are the Core Services installed with PVS Server? Explain SOAP and STREAM Services roles and Impact if service is not available or restart of service.

The Soap Service is for Console and management operations only and will not effect streaming at all in case of service restart.

The stream service will helps in streaming images and booting vdisk's. Restart of PVS will not effect any existing connections, however failure of this impacts vdisk booting issues.

- When there is only one pvs server, and the devices are already connected, even if we restart stream service and soap services, there will not be any problem. We will see a tftp error only when devices are booting up.
- If there are two or more, pvs servers, new connections might be pointing to other servers, and the existing connections wont have any issue, when restarting stream and soap services.
- No matter you have one pvs server or multiple pvs servers, restarting stream and soap services doesnt effect anything. When there is only one pvs server, for new connections, we ll see a tftp error when restarting those services. And when there are multiple pvs servers, new connections will be pointed to next pvs server.
- when you create device in pvs, and after device is created, right click on that device, active directory, create machine account. If you do this, when ever a machine is booted up, its hostname management will be done by pvs. This will be very useful while updating your pvs images.

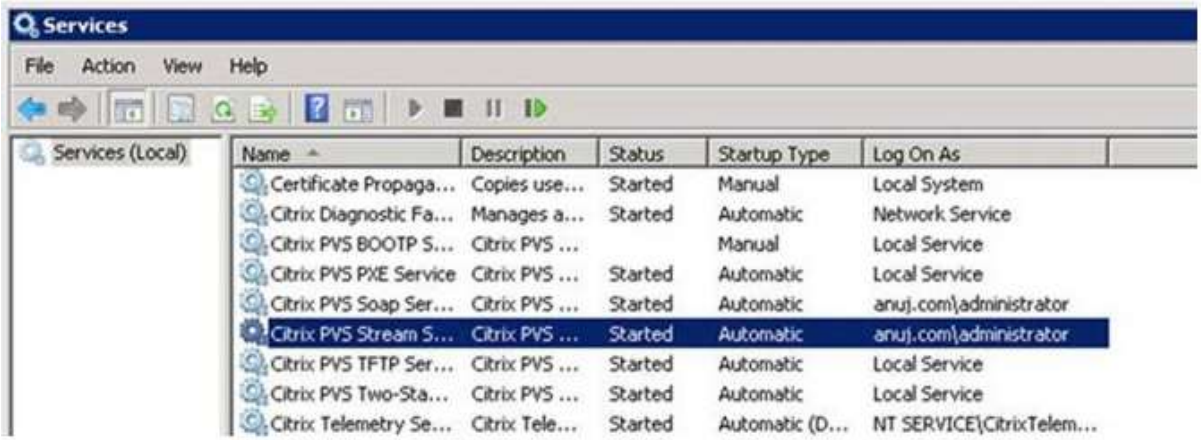

Ref:
https://docs.citrix.com/en-us/provisioning/7-13/managing-servers/server-services-start-stop.html

Explain the Pvs vDisk Logon (Boot) Process

Citrix Provisioning Server (PVS) streams the contents of the vDisk to the Client VMs (target device) on demand, in real time. The Client VMs boot directly across the network and behaves as if it is running from its local drive. This article talks about various steps involved in network boot and how the vDisk is streamed.

Steps:
Let's say, a Client VM is powered ON and it is configured to boot from Network

Since it is configured to boot from network, it gets IP address from DHCP Server.

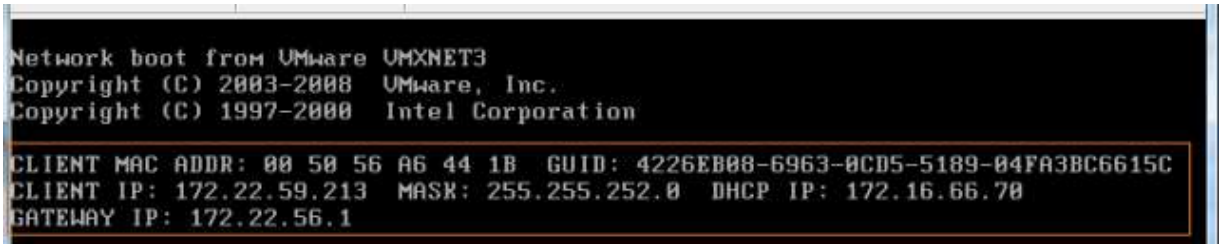

After getting IP Address from DHCP, Client sends broadcast request for PXE servers in the network. In general both PXE and TFTP services are provided by the PVS server itself.

Let's say there are 2 PVS servers in the network having PXE service enabled. Both the PVS servers will respond to the client's broadcast request.

Client contacts the PXE server from which it receives the response first. PXE session is established

PXE service on PVS server sends the information of TFTP server IP address and name of the Bootstrap Program file to download from TFTP. ***The only use of PXE is to provide TFTP server details to clients***.

Bootstrap program is a small program which initializes the streaming session between Client and PVS server.

Client establishes the TFTP connection to PVS Server and downloads the Bootstrap Program. ***The only purpose of TFTP service is to provide Bootstrap program to clients.***

Bootstrap program file contains the information of which PVS servers to contact for vDisk download. In Citrix terminology, these servers are called Login servers (shown in above screenshot).

Bootstrap program contacts the PVS servers for vDisk. PVS checks its database for the target device mapped to client's MAC address and the vDisk assigned to it.

Client establishes the streaming session to PVS Server. Client boots up with the OS image in the vDisk file

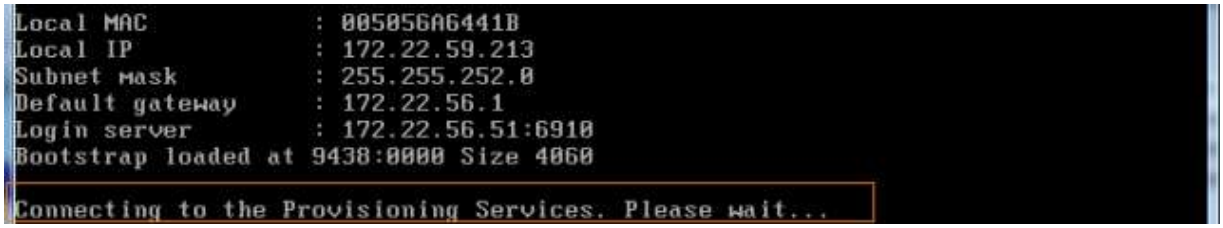

Just to summarize, PXE is used for getting the TFTP Server IP and bootstrap file name details by the clients and TFTP is used for downloading bootstrap program file

Is there a way to eliminate PXE from the communication process?
Yes, by specifying TFTP Server IP Address and bootstrap file name in DHCP Scope options on DHCP Server. DHCP Scope options 66 and 67 can be used. DHCP server provides TFTP details to the clients along with dynamic IP address. Client VMs will directly contact TFTP server and download the bootstrap program.

Is there a way to eliminate both PXE and TFTP from communication process?

Yes, by using Boot Device Manager (BDM) both PXE and TFTP can be eliminated from communication process. Using Boot Device Manager utility on PVS Server, bootstrap file can be written into a ISO image, USB flash drive, or local hard drive. Client VMs need to be configured to boot from ISO image file. As ISO image file itself has bootstrap program, clients doesn't use PXE or TFTP.

What are the different (3) ways to load bootstrap program (PVS Boot process options) and name few pros & cons of these methods? (Like PXE,BDM,BIOS embedded Bootstrap)

Different options to boot a target device (vdisk) with Citrix Provisioning Services are TFP,PXE,BDM(Boot Device Management), BIOS-embeded bootstrap) and target device should know from where the streamed VHD to request.

To know difference quickly ,

One of the most common boot scenario is using A: PXE or B: DHCP both in combination with TFTP.

Option A:
With PXE ,you don't have to configure the DHCP options to provide the TFTP server to the targets, you can create a kind of redundant configuration by setting up multiple PXE and TFTP servers, but please note that this is not a real HA configuration because there is no logic involved which controls the way a TFTP server is

provided to the targets. For example a broken or unavailable TFTP server can be provided to your targets, you can compare this a little with the way how DNS round robin works.

Option B:
With the DHCP options ,you can only provide one TFTP server to your targets, to enable HA here you can configure a load balancer (NetScaler for example) in front of the TFTP servers, this is more HA then option A because you can configure the load balancer to check the health of the TFTP servers and bypass TFTP servers that are currently in down state. But you will need to add multiple nodes in your load balancing configuration so you don't have a single point of failure.

With both option A and B TFTP is used to deliver the bootstrap to your targets, but what if we can't use PXE or TFTP because of network restrains or just because we want to eliminate the whole PXE and TFTP dependency… Yes, we also have the option to create a bootable disk or ISO with the bootstrap embedded, it contains a list of the Provisioning Servers to provide HA for your VDISK.

To create the ISO we use the Provisioning Services Boot Device Manager which is part of the Provisioning Server installation.

Using a bootable ISO is a great way to overcome network related issues that can be the case when using TFTP and\or PXE. DHCP will always play an important role in your Provisioning Services environment, use split scopes or clustering there to guaranty the uptime of your Provisioning environment. Provisioning Services comes with a lot of "moving parts" compared to MCS, but with this boot option you can at least eliminate a few of them. It's nice to see PVS is still alive in Excalibur so there is room for MCS to grow, the power is choice!

To Know in depth of Boot Options

TFTP

I think that is the most common in Citrix PVS environment and very easy to set up. We only need a DHCP Server and PVS Server.

The DHCP server must be configured with options 66 (boot server) and 67 (file name) and the Citrix PVS TFTP Service must me running on the PVS Server and the target startup must be setup to network device.

The option 66 define the boot server. (IP of PVS Server)

The option 67 define the filename that will receive by TFTP (ARDBP32.bin)

TFTP Installation and configuration

During the installation of Citrix PVS Server will ask to use the TFTP Service

If you want to change some TFTP configuration, you can re-run the PVS service configuration wizard. Other useful tool is TFTP CPL on C:\Program Files\Citrix\Provisioning Services\tftpcpl.cpl. with this you can change the path of ardbp32.bin, enable logs and chose the network bind of TFTP Service.

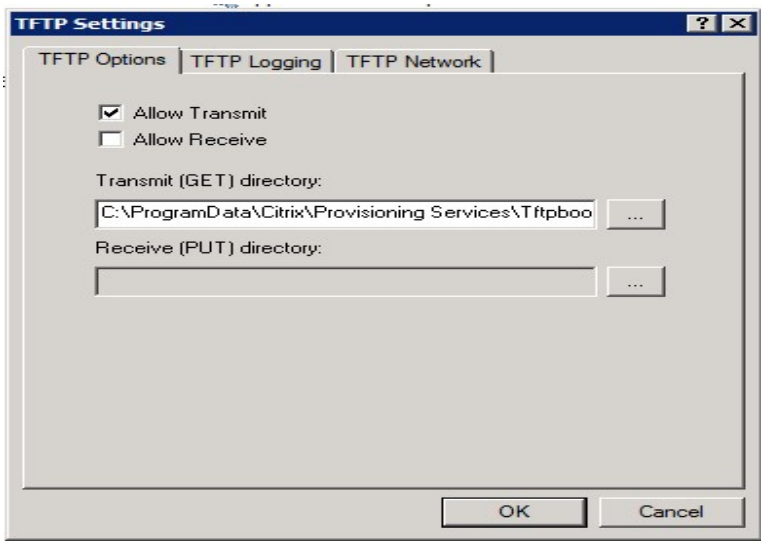

For those environments that request a load balance can be easier accomplish with NetScaler. If you don´t have a NetScaler you can set up multiple IP from PVS TFTP Server on DHCP option 66 separated by a semicolon (if your DHCP support). Another think that you can keep in mind that you client must be able to interpret the option 66 with multiple entries.

The flow is like this. (After the TFTP Data Packet this flow was resumed)

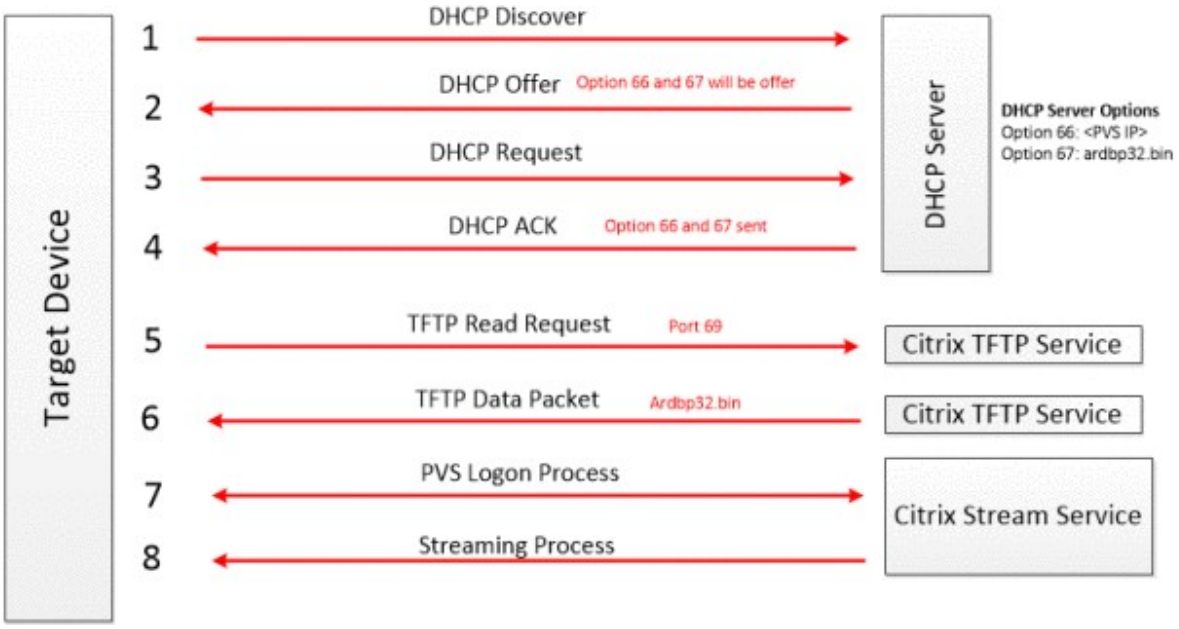

Citrix PVS PXE Service

With this type of boot, you won´t need a DHCP with the option 66 and 67. This option can be initial configures when you run the Citrix PVS Configuration Wizard (or just enable the service on control panel > services)

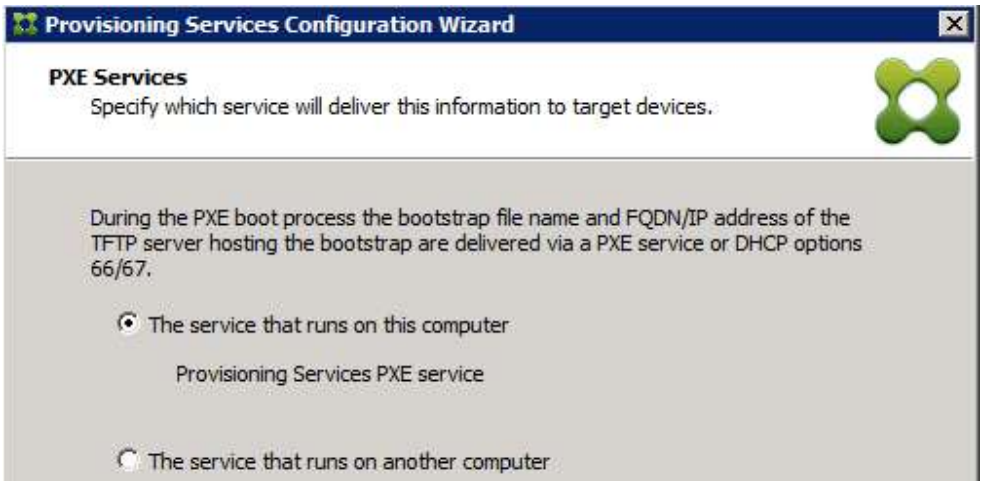

With Citrix PXE built-in service the client during the startup, send broadcast requesting the option 60 (PXEClient) and all PVS Server that handle this service will answer this request and send the file called bootptab. The file bootptab make reference of ardbp32.bin file that contain the image file for the target device.

You can use the PXE tool to configure some option like Logs, Network Bind and Proxy DHCP.

The location is C:\Program Files\Citrix\Provisioning Services\BNPXE.cpl

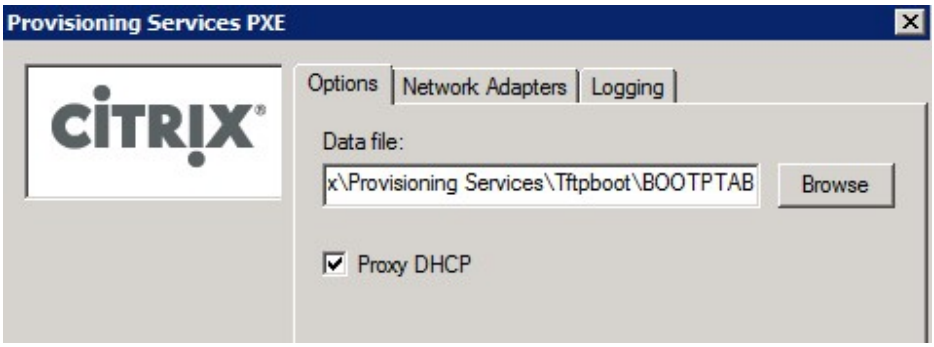

The boot process flow is like this:

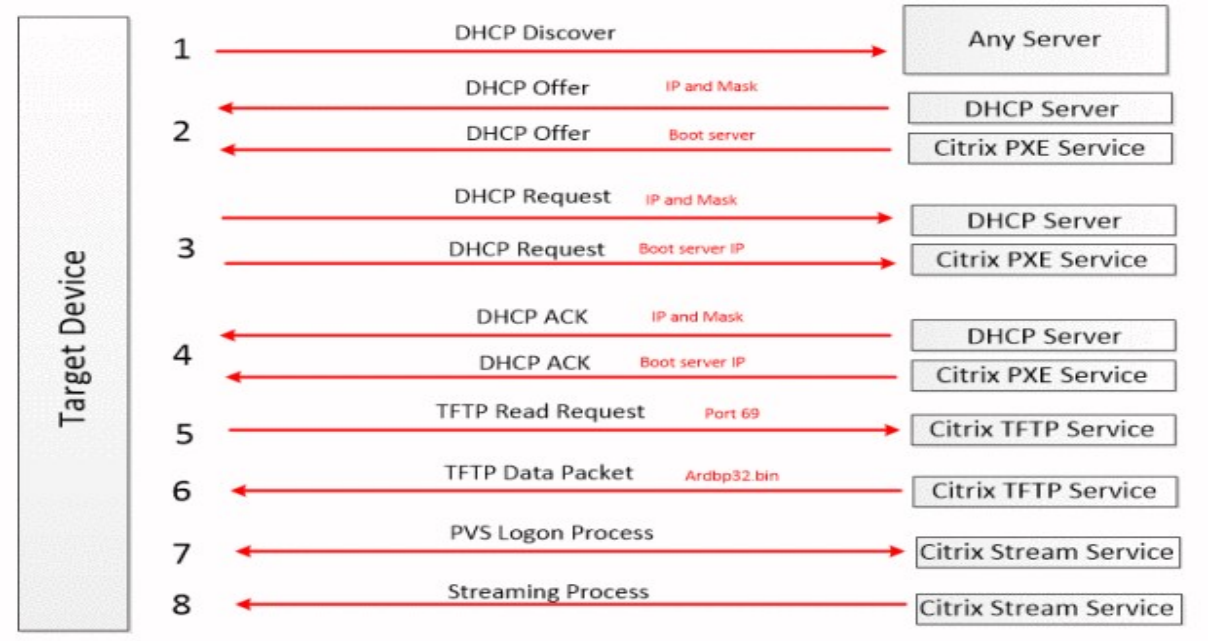

Load balancing for this type of boot process is really simple because when a target send broadcast searching for some dhcp, automatically all the PVS Server with PXE service enabled will respond the request with the boot file. So in case of environment with 2 PVS Server with PXE enabled and one of this servers goes down, the remain PVS Server will answer the dhcp request.

Citrix ISO Image Recorder

In this type don´t need a DHCP option 66 and 67. During the process to burn an ISO, you need to insert the IP or DNS Name of PVS Server, so the target already know the IP to get the boot file. This can easily combined with DNS to make load balance with round robin fashion.

It´s use the boot file TSBBDM.bin and the Citrix PVS Two-Stage Boot Services.

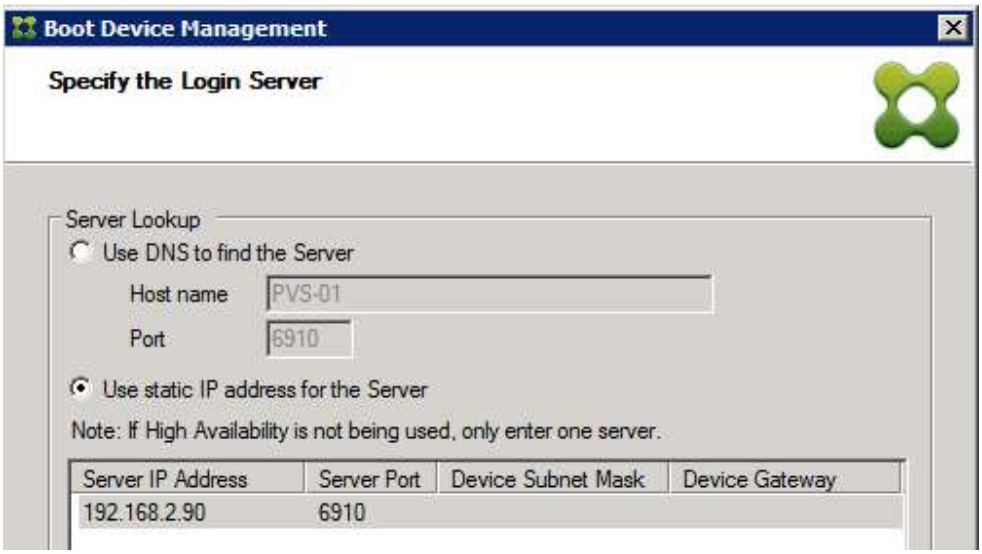

The Target Device must configured in BIOS startup through CD/DVD, so in virtual environments could be a bad solution because sometimes is we need to vmotion virtual machines, and if the virtual machine has and local ISO attached, this vmotion will fail.

the boot process flow is like this

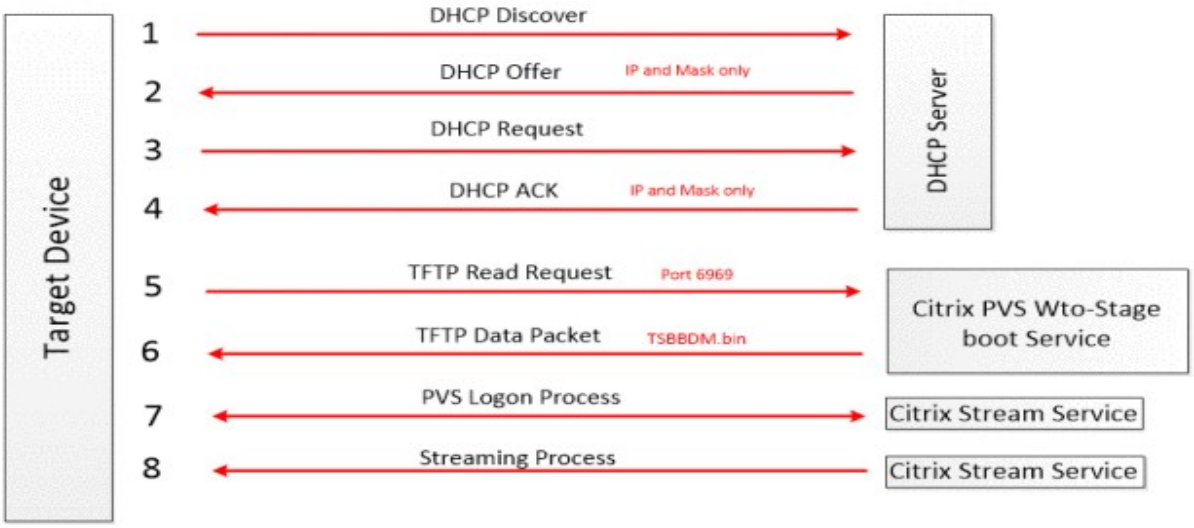

Virtual HD

The Virtual HD act exact how ISO works. The process to burn the information on a virtual disk drive is diferent but equal easy.

To save the information on vhd you will need to add this vhd to a Provioning Server and then execute the Citrix Boot Device Management.

The process to create is:

1- Add a virtual disk on you PVS Server (the vhd can be very small. 100mb is recommended by Citrix)

2- Format and mark this disk with ACTIVE PARTITION

Volume	Layout	Type	File System	Status	Capacity
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page File, Crash Dump, Primary Partition)	39.90 GB
(D:)	Simple	Basic	RAW	Healthy (Active, Primary Partition)	10 MB

3- Open the Citrix Boot Device Management and enter the information about your PVS enviroment

4- Select your 10mb virtual disk to burn and click and Burn

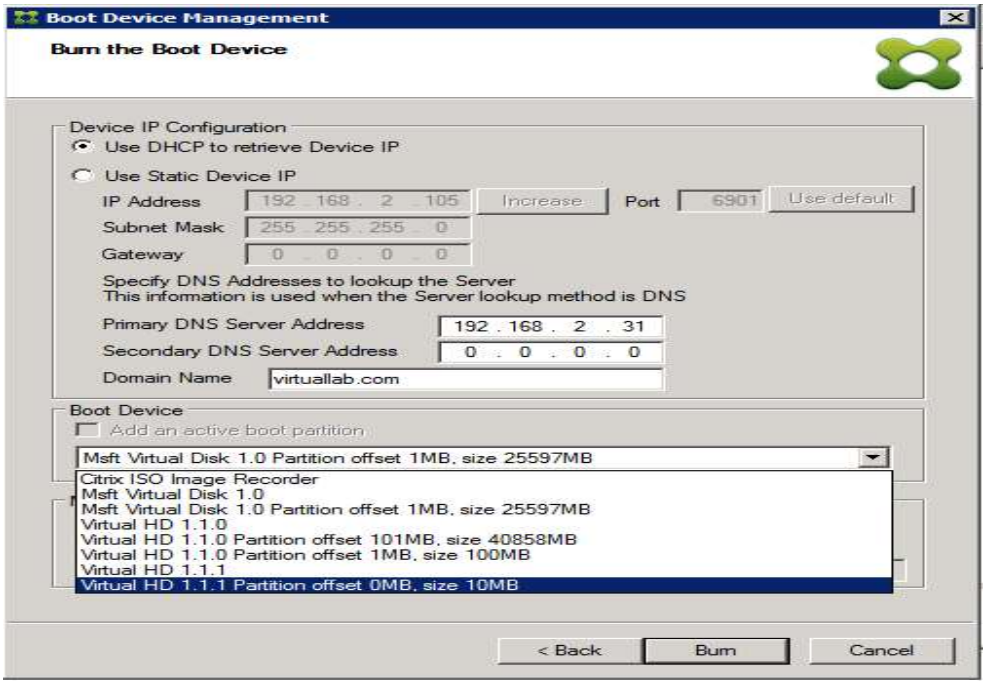

Remember. If ask you to enter “Yes” to complete the process, you need enter exactly the the word Yes with the Y in lowercase.

5- After that you need to detach the vdisk of PVS Server and attach to your target device and set up a startup order to hard disk

Load Balancing with this option is the same with Citrix ISO Image

Citrix PVS BootP Service

This service in very useful if you cannot change your DHCP. This service is responsible to delivery IP to Target device acoording wit BootP Tab.

C:\Program Files\Citrix\Provisioning Services\BNBOOTP.cpl

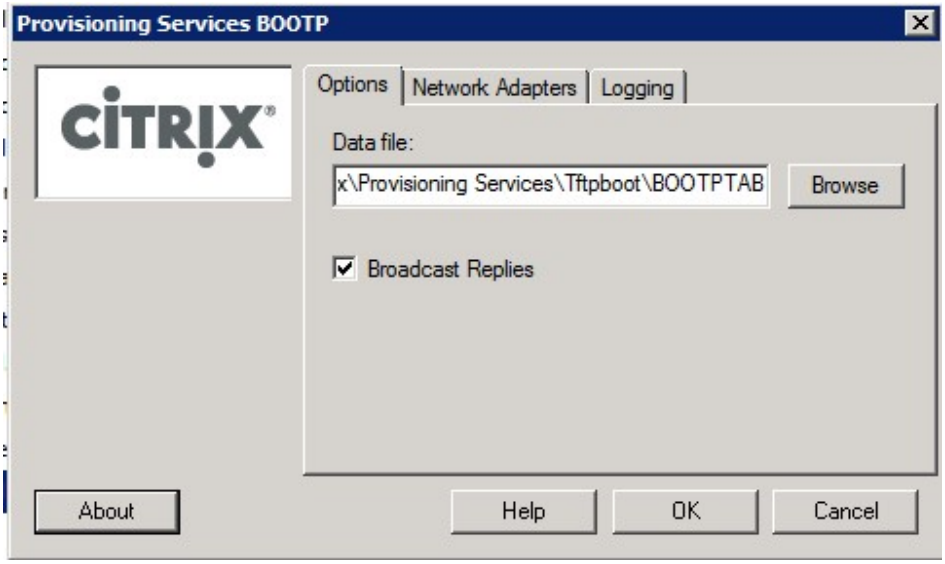

Configuring the BIOS-embedded bootstrap

This feature is OEM specific and provides end users with systems preconfigured with Provisioning Services, allowing customers to deploy an Provisioning Services-enabled environment with minimal effort. This feature becomes an alternative to the standard PXE boot method.

As part of this solution, the OEM embeds the bootstrap within the target device's BIOS at the factory. The OEM also pre-configures the device with product license keys.

If the target device boots using the BIOS-Embedded Bootstrap, the configuration settings are obtained from the device's BIOS. These BIOS settings may indicate using DHCP with DNS to lookup the IP and server information (dynamic), or it may list up to four server IP addresses in the BIOS (static).

The first time a target device boots, it reads the product license key and configuration information from the BIOS, locates the Stream Service, and then sends a device registration message to the server. This message contains the information, in addition to the information inherited from the device collection template, necessary to add the device to the Provisioning Services database.

Ref:

http://docs.citrix.com/en-us/provisioning/7-9/managing-target-device/pvs-bootstrap-bios-embedded.html

https://docs.citrix.com/en-us/provisioning/7-6/pvs-install-wrapper/pvs-bootstrap-from-console.html

https://venthusiastic.wordpress.com/2014/01/11/citrix-provisioning-service-boot-options/
https://bramwolfs.com/tag/pvs-bootstrap/
http://docs.citrix.com/en-us/provisioning/7-9/managing-target-device/pvs-bootstrap-bios-embedded.html

How to update a Master image in Dedicated (Static) catalog in MCS? Explain procedure updating catalog for Static and Pooled

For obvious reasons, XenDesktop does not provide the ability to update Static (Dedicated) machine catalogs, only Random (Pooled). On the other hand, this is ok, because once deployed by MCS, Static desktops will thereafter be maintained by SCCM, so we simply don't touch the already deployed machine catalog. But where does that leave new desktops?

Let's assume, we would have to deploy a new machine catalog pointing to the new image/snapshot whenever we re-baseline each quarter. This is where it starts to get interesting. This approach would mean 4 new machine catalogs per year, per customer. And we need to maintain these catalogs as long as there are people connecting to them. The list of catalogs will become huge over time or we may have a situation where SCCM is not in place

If the frequency of your image updates are more then it is recommend you to use the pooled/random catalog types only, one the backend either use MCS or PVS). And use romaing profiles to store some user preferences for different application.

In this way, your users will always get the latest application from the VDA and user data comes from some other central place, not trying to any particular VDA machines.

In addition, it is also good to consider non-persistent + personal vDisk. The decision was made however to manage static virtual desktops (once deployed) pretty much like we would a standard traditional desktop

In simple, Pooled desktops allow us to easily update our master image and ensure that all machines receive that update.

Persistent(Static) machines are made from a master image but then must be managed going forward with other ERD (electronic Distribution tools) like SCCM & Altiris.

With Dedicated Catalog of machines the snapshot chosen at catalog creation time is forever more used to create new machines in that catalog. This can prove a challenge if the master image over time becomes out of date, or possibly could contain something that is wrong/irrelevant going forward

Updating a Persistent Catalogs tie to its master machine cannot be performed in the Citrix Studio console but can be performed using Powershell commands

Procedure to Update Static Catalog

If you pay close attention when you initially deploy the image, you will notice that MCS will do a full VMDK copy of your snapshot chain into a folder of every datastore that is defined in your hosted XenDesktop environment. This makes desktop creations extremely quick when scaling out additional VMs because it 1) negates the need to potentially copy VMDKs across datastores during desktop creation and 2) negates the need to consolidate snapshots during creation. The folder will typically be the **machine catalog name + basedisk + random datastore** identifier assigned by XenDesktop. This applies to all MCS images; static and pooled.

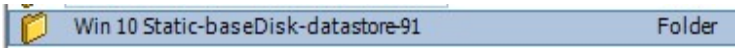

We obviously want to keep the master of dedicated machines up-to-date to avoid unnecessary SCCM pushes, Windows updates, missed software, etc. when we deploy new desktops. Unfortunately, Citrix does not give a GUI option for this, like we get on our pooled desktops in Studio:

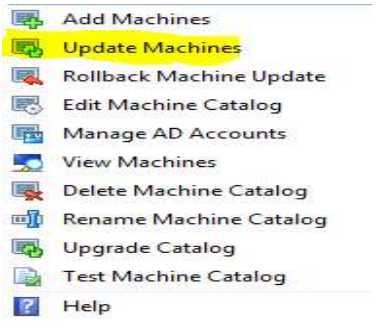

So, what is usually the method of action when no GUI option is available? That is right – PowerShell!

There are two main things to consider here: the “Provisioning Scheme” and the new “Master Image.” The provisioning scheme name almost always matches the machine catalog name. It keeps track of the master image location and some other metadata. The master image is just the snapshot of your master machine that MCS does that full VMDK copy to each datastore that we talked about earlier.

Let’s get right to it. First, open PowerShell on your DDC, and get the provisioning scheme name and the current snapshot that is being used for the master:

add-pssnapin *citrix*

Get-ProvScheme

This will return two very important things for each MCS machine catalog: 1) the ProvisioningSchemeName and 2) the MasterImageVM. You will notice that this contains the name of the

snapshot that mirrors the name you provided in vSphere, followed by .snapshot. This makes it easy to locate!

Let's assume our current snapshot is named "v1" and our master is named "XDMaster1." Therefore, the MasterImageVM should look like:

XDHyp:\HostingUnits\<Cluster Name>\XDMaster1.vm\v1.snapshot

Note: If your VM is in a resource pool, this path will also contain that as a "directory."

We will create a snapshot named "v2" on the master after making some changes, updates, etc. (it is good to create snapshot before changes too for easy rollback but give proper name for easy identification) and shutdown the master. Let us verify that XenDesktop now sees this snapshot in our hypervisor environment:

get-childitem -path "XDHyp:\HostingUnits\<Cluster Name>\XDMaster1.vm\v1.snapshot"

You will see that v2.snapshot is now a child item of your v1.snapshot! Good deal! So how do we point MCS to this snapshot? Simple:

First, let's make it easy on ourselves and create a couple of variables. The two important ones that I touched on earlier: ProvisioningSchemeName and MasterImageVM:

$ProvScheme = "Windows 10 Static"

"Windows 10 Static" will be the ProvisioningSchemeName from earlier, or usually the name of your Machine Catalog.

$NewMasterImage = "XDHyp:\HostingUnits\<Cluster Name>\XDMaster1.vm\v1.snapshot\v2.snapshot"

That will be the full path to your new snapshot. *Remember to use get-childitem to ensure that the DDC sees your new snapshot.*

Now, we will use the Publish-ProvMasterVMImage cmdlet to wrap it all up!

Publish-ProvMasterVMImage -ProvisioningSchemeName $ProvScheme -MasterImageVM $NewMasterImage

After running this command, pay attention to your vSphere tasks. You will see a temporary VM get copied, VMDKs get copied to the various datastores, and you should finally get a response from PowerShell that states 100% completion and where the new master image location points.
If you see the dreadful red text, pay attention and make sure you got your paths correct. It is easy to mistype the XDHyp path, forget quotes, etc.

Procedure to update Pooled Catalog

Citrix recommends that you save copies or snapshots of master images before you update the machines in the catalog. The database keeps an historical record of the master images used with each Machine Catalog. You can roll back (revert) machines in a catalog to use the previous version of the master image if users encounter problems with updates you deployed to their desktops, thereby minimizing user downtime. Do not delete, move, or rename master images; otherwise, you will not be able to revert a catalog to use them.

For catalogs that use Provisioning Services, you must publish a new vDisk to apply changes to the catalog. For details, see the Provisioning Services documentation.

After a machine is updated, it restarts automatically.

Update or create a new master image

Before you update the Machine Catalog, either update an existing master image or create a new one on your host hypervisor.

1. On your hypervisor or cloud service provider, take a snapshot of the current VM and give the snapshot a meaningful name. This snapshot can be used to revert (roll back) machines in the catalog, if needed.
2. If necessary, power on the master image, and log on.
3. Install updates or make any required changes to the master image.
4. ***If the master image uses a personal vDisk, update the inventory.***
5. Power off the VM.
6. Take a snapshot of the VM, and give the snapshot a meaningful name that will be recognized when the catalog is updated in Studio. Although Studio can create a snapshot, Citrix recommends that you create a snapshot using the hypervisor management console, and then select that snapshot in Studio. This enables you to provide a meaningful name and description rather than an automatically generated name. For GPU master images, you can change the master image only through the XenServer XenCenter console.

Update the catalog

To prepare and roll out the update to all machines in a catalog:

1. Select **Machine Catalogs** in the Studio navigation pane.
2. Select a catalog and then select **Update Machines** in the Actions pane.
3. On the **Master Image** page, select the host and the image you want to roll out.

4. On the **Rollout Strategy** page, choose when the machines in the Machine Catalog will be updated with the new master image: on the next shutdown or immediately. See below for details.
5. Verify the information on the **Summary** page and then click **Finish**. Each machine restarts automatically after it is updated.

Tip: If you are updating a catalog using the PowerShell SDK directly, rather than Studio, you can specify a hypervisor template (VMTemplates), as an alternative to an image or a snapshot of an image.

Ref:

http://nicksitblog.com/2017/02/update-a-staticdedicated-mcs-image-vsphere/

http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-7/install-configure/machine-catalogs-manage.html#par_anchortitle_f39b

Is It possible to Configure SQL HA for existing farm either Mirroring or SQL Always On> If yes? What needs to be run at PVS?

Enabling Mirroring Within an Existing Farm

To enable mirroring within an existing farm:

1. Confirm that the primary and failover database servers are up and running.
2. Using MS SQL server tools, mirror the Provisioning Services database to a database on the failover database server.
3. Run the Configuration Wizard on each server.
4. Identify the farm by choosing either the Farm is already configured or the Join exisiting farmoption on the Farm Configuration page.
5. On the Database Server page, select the primary and failover database servers and instance names, then enable the database mirror failover feature .
6. Complete the remaining wizard pages.

Ref:

https://docs.citrix.com/en-us/provisioning/7-13/managing-high-availability/ha-db-mirror.html

Citrix Licensing

What are the different types of license models available in XenApp and Xen Desktop and mention difference.

Types of Licenses are varies with version, below type of license available from 11.13

Ref:

https://support.citrix.com/article/CTX123762

https://www.citrix.co.in/buy/licensing/product.html

http://www.basvankaam.com/2016/04/18/citrix-licensing-and-microsoft-demystified-i-need-your-help/

http://blog.lakesidesoftware.com/entries/citrix-licensing-deciding-between-concurrent-and-userdevice-licenses/

https://docs.citrix.com/en-us/licensing/11-13-1/lic-license-types.html