



Tintri
Intelligent Infrastructure

Citrix XenDesktop™ v7.1

CITRIX®
READY

Machine Creation Services (MCS) Deep Dive

Exploring Citrix XenDesktop™ v7.1 Machine Creation Services
on VMware vSphere™ 5.1 and Tintri VMstore™

Revision 0.1, April 2014

Contents

Executive Summary	1
Consolidated list of practices	1
Machine Creation Services Overview	3
The Good, the Bad and the Ugly	4
Pros	4
Cons.	4
Gotchas	5
Big Bang Theory – The creation process	6
Pandora’s Box – The Identity Disk	12
Consideration Factors	14
Space Utilization & Capacity Considerations	14
Paging and swap files	14
Persistence	15
Audit & Compliance	15
Additional Factors	16
Management & Housekeeping	17
Adding VMs to an MCS-based Machine Catalog	17
The Image Update Process	17
The Roll Back process	19
Maintaining Image Versions	20
Environmental Changes to the VMware Infrastructure	21
Data Protection & Disaster Recovery (DR) Overview	22
Backup & DR of MCS-based Machine Catalogs	22
Conclusion	23
Additional Resources – Reference URLs	23
Tintri Links	23

Citrix Links	23
Other	23
<u>Appendix A – Environment Details</u>	<u>24</u>

Executive Summary

This paper explores Citrix Machine Creation Services (MCS) in Citrix XenDesktop 7.1 deployed on Tintri VM-aware storage and is intended to complement the paper: [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#).

The target audience includes IT admins, architects, and other technical roles that already have some experience with XenDesktop, vSphere and storage. If you do not have any experience with XenDesktop yet, we strongly suggest you follow our [XenDesktop 7.1 Quick Start Guide](#) to get a functional XenDesktop environment in place, and then continue reading this paper.

There are three methods available to provision XenDesktop virtual desktops:

- [Citrix Provisioning Services \(PVS\)](#).
- **Machine Creation Services (MCS) – focus of this paper.**
- [Tintri Native space and performance efficient VAAI Clones](#).

All three methods are all fully supported and work very well on Tintri VMstore systems. Deciding which method(s) to use for your organization depends on many factors, and the goal of the series of papers is to give you more insight into each technology to help you make a well-informed decision for your enterprise.

We recommend you try each of the methods covered in the XenDesktop Deep Dive series of papers. Without further ado, let's dive into MCS!

Consolidated list of practices

The table below includes the recommended practices in this document. Click the text in the *Recommendation* column or see the section later in the document that corresponds to each recommendation for additional information.

Tags	Recommendation
Pandora's Box - The Identity Disk	<i>Do: Review all of the information listed and make a note of it somewhere. If you plan on changing your environment in anyway, remember to keep your XenDesktop VMs in mind and test well to determine whether the proposed Infrastructure changes you want to make will break your VMs or not.</i>
Space Utilization & Capacity Considerations	<i>Do: When cleaning up old base disk folders and files, be diligent and investigate thoroughly to make sure the files are not needed before deleting.</i>
Paging and swap files	<i>Do: From time to time, test your images with different amounts of memory assigned to each VM to determine how much RAM is the right amount and add RAM to your hosts in necessary.</i>
Persistence	<i>Do: If persistence is important in your environment, consider using Tintri Native clones. Citrix Personal vDisks can help you, but are out of scope of this paper.</i>
Audit & Compliance	<i>Do: Check the logs to ensure that actions you expect to be recorded are in fact recorded, and are sufficient evidence to meet the audit and compliance requirements of your environment.</i>
Audit & Compliance	<i>Do: If the Event logs or other monitoring software within the guest require retention, consider a centralized logging solution to record and maintain yours logs outside of your desktops VMs.</i>

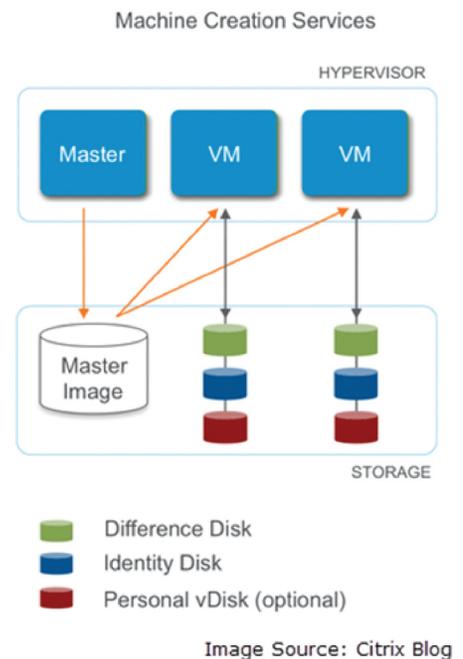
Tags	Recommendation
Adding VMs to an MCS-based Machine Catalog	<i>Do: Once you've created new VMs, don't forget to also add them to your delivery group(s) to make them available for users to connect to.</i>
The Image Update Process	<i>Do: Staggering your updates is desirable as it gives an opportunity to catch any bugs or issues specific to the new master without having 100% of your user population stumbling on it at once.</i>
Environmental Changes to the VMware Infrastructure	<i>Do: Test and be diligent, and add these gotchas to a checklist of things to check when undertaking a migration, major upgrade, or other changes that are required to stay current and supportable.</i>
Backup & DR of MCS-based Machine Catalogs	<i>Do: Consider using 2 or more active sites in a global Citrix topology, each configured with their own unique machine catalogs.</i>

Machine Creation Services Overview

Citrix MCS provides the simplest means of creating a machine catalog. MCS tools and services are baked into a default XenDesktop install and is ready for immediate use after install and the basic configuration steps of adding a VMware vCenter server and host resources to your XD deployment. A step-by-step walkthrough on how to create/configure this can be found in the [XenDesktop 7.1 Quick Start Guide](#).

A Machine Catalog created with MCS has the following characteristics:

- A master image is prepared and created based on a standard VM that has all desired software and customizations that an admin wants to make available in users' virtual desktops.
- The finalized master image is a .vmdk virtual disk file that resides in a folder in the root of the datastore it is deployed it (i.e. Tintri VMstore).
- An admin defines the size of the VMs by choosing how many vCPUs and how much vMemory of each desktop VM.
- The admin chooses how many VMs to be created and added to the machine catalog and defines a naming convention, such as "Win81-MCS-###" (where ### will automatically increment as 001, 002, etc.).
- VMs are all created as "**Linked Clones**". Each VM created has two or three disks assigned to it (see figure for reference):
 - Master Image, A single Base Disk (.vmdk file), mounted as Non-persistent on SCSI 0:0 to each VM. This allows read-only access to the base .vmdk file, and all data change is redirected to a .REDO file in the VM's home folder. It is the relationship between each VM and this read-only disk that is mounted in each VM that earned the descriptive name "**Linked Clone**".
 - Personality disk, a small 16 MB VMDK mounted on SCSI 0:1. This disk contains some basic information such as the machine name, SID, domain computer account password, and any other unique info that needs to be injected into the OS on boot-up.
 - Personal vDisk (option) – This disk is used in cases where you want persistence in your VM (out of scope of this paper).
- Personalization of each cloned VM in the machine catalog is done automatically and there are no special requirements.
- During creation, the Tintri VAAI plugin that has been installed on every host is automatically leveraged which speeds up the process and keeps the copies space-efficient.



The Good, the Bad and the Ugly

This section is devoted to highlighting Pros, Cons and Gotchas specific to MCS. In some environments weakness listed may be a showstopper, whereas in other environments, it may not cause any inconvenience at all. One size does not fit all, so one needs to weigh the importance of each strength/weakness based on the environment in which the solution is going to be deployed.

Pros

Provisioning and maintenance

- Easy to create the initial machine catalog.
- Initial Machine creation process leverages VAAI, which dramatically reduces the time and storage capacity required to create a master base image to be linked to clones.
 - **NOTE:** Only applies if there are no VMware snapshots on the Master Image selected. More information is provided in an item in the Cons section below).
- Simple to add new machines to an existing catalog
- Image update process is simple
 - Image updates can leverage a rolling reboot update when set to “update on next reboot”. More detail to follow under the Image Updates section within Management & House-keeping.

Cons

Provisioning and maintenance

- Some advanced VM settings, such as CPU/MMU Virtualization settings aren't defined and applied to VMs created.
- Only a single master base disk is copied, prepared, and added to the MCS machines created. Additional disks configured on the master VM (such as a dedicated disk for a pagefile) do not appear in your cloned VMs.
- Machine Catalogs are tied to the single host or cluster that is defined in the Hosting Configuration.
- Machine Catalogs are tied to the same datastore initially configured in Hosting Configuration.
- A VM template (.vmtx) can't be selected for the Master Image, only a Virtual Machine (.vmx). Although a template can easily be converted into a VM at the time of machine creation, this is listed as a Con because Templates do not show in Hosts & Clusters view, and are clearly identified in VMs and Templates view by having a different logo. During routine sweeps to clean up old VMs, it is possible you would see a VM powered off that hasn't been used for some time, but should not be removed or your image roll back process will be impacted. To avoid having VMs in your inventory and being subjected to accidental deletion, we recommend placing VMs to be used as Master Images in a folder that identifies them as such.

Storage and IO consumption

- Hidden storage consumption. Base disks for each version collect in the datastore and readily visible unless browsing the datastore.
- Image updates are full copies of base vmdk (wastes capacity).
- There is some IO overhead associated with non-persistent disks (REDO.vmdk).

Interaction with snapshots

- MCS creation process creates and leaves VMware snapshots on the VM selected as the master image. This can negatively affect you in the following ways:
 - If VMware snapshots are present on the VM selected as the Master image:
 - VAAI will not engage and a full copy of the vmdk will be made.
 - Creation process will take longer and consume more storage.
 - Subsequent updates to the same initial master image will not engage due to the VMware snapshot created during the initial MCS machine catalog creation process.
 - If VMware snapshots are deleted or the original Master VM is removed, “Image rollback” will fail. In this case, an error is given and you will be prompted to select a new Master VM. While this in itself may not seem problematic, there’s a catch: you can only select an alternate Master VM that has a VMware snapshot on it, otherwise you will continue to get errors that prevent a rollback.
 - **WORKAROUND:** Use the standard update procedure to present the previous image version as a new one to promote your VMs to. To create a new Master VM based on the previous image version, copy the previous base image (.vmdk) to a new VM and select it as a new Master VM during the update process.

Gotchas

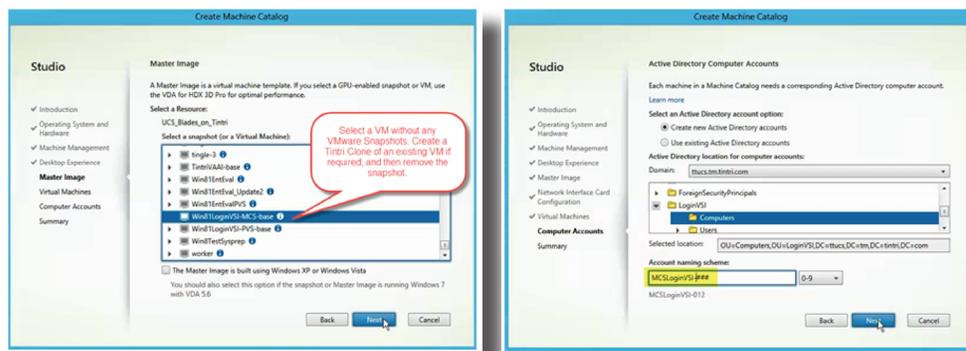
MCS does NOT handle VMware environmental changes well! Changes to cluster names, datastore names and network names can all “break” MCS configuration. Let’s look at one example of this: renaming a datastore. In order to be able to rename a VMware datastore successfully (and still have a fully functional XenDesktop deployment), the hosting resources need to be deleted and added back.. –However, dependencies from delivery groups prevent deletion of the resource. In order to remove the dependencies, the MCS machine catalogs needed to be removed first, but can’t be removed if they are in use (i.e. shutdown all your virtual desktops!).

And to make things worse, if you’ve already renamed your datastore, you won’t be able to delete the machine catalog unless the old datastore name is found in order to complete the process. Once all dependencies are off the original storage resource, you can remove it. Once removed, you can rename your datastore and then re-add a storage resource back under Configuration – Hosting and then setup up Machine Catalogs again on the new resource.

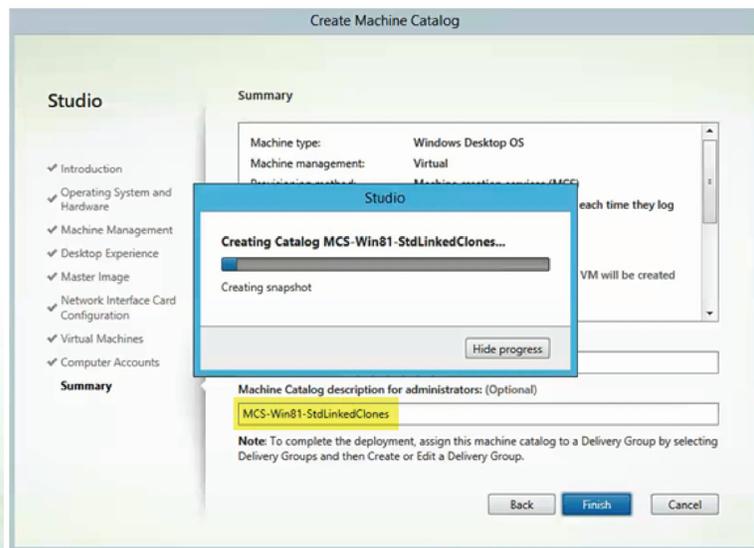
Big Bang Theory – The creation process

The MCS creation process is very simple from an end-user perspective. The process is wizard driven and all we need to do is fill in a few fields (refer to [Quick Start Guide](#) for play-by-play screenshots), then sit back and let the magic happen. But what is happening, and how do Tintri features benefit the process? Let's dig in and have a look!

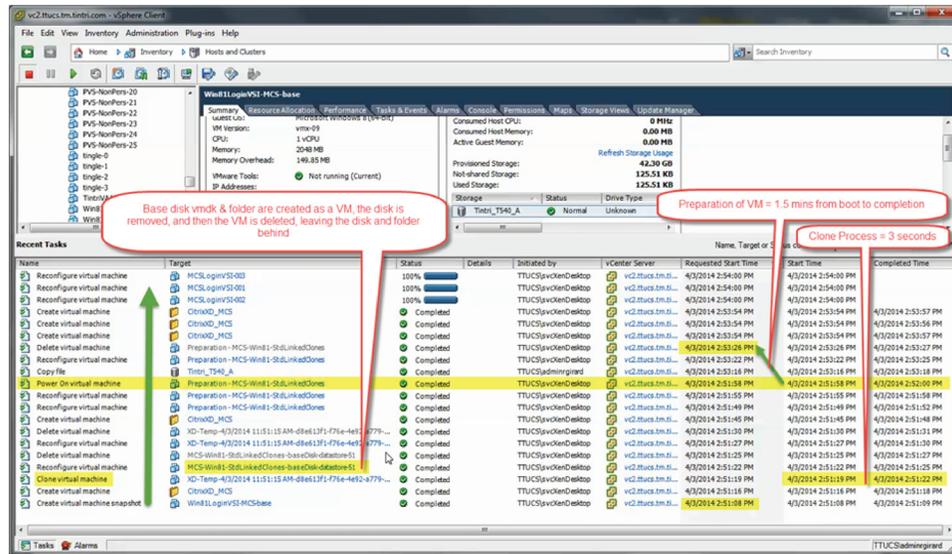
During the MCS Machine Catalog Creation wizard, we select a master image VM, as shown in the screenshots below. This needs to be a VM (VMware templates will not show in the list), and make sure that the VM does NOT have any VMware snapshots or else VAAI may not activate. If VAAI does not activate, you will not reap the benefits (increased speed, reduced host IO & no additional space consumed on VMstore) of offloading the copy/clone process to the VMstore.



After selecting a VM to be the Master image and providing a naming scheme, you're prompted for a Machine Catalog Name (shown below). Choose wisely... this name is what will be used to name the base disk and folder for your base disk that will be mounted non-persistently to every VM in your new catalog (as you'll see below). After clicking finish, the creation process begins:



In the Citrix Studio console, there's very little to see, other than a progress meter. During the creation process, we've captured activity taking place elsewhere to see what is happening under the covers, primarily using the vCenter tasks for monitor the steps. Let's follow along:

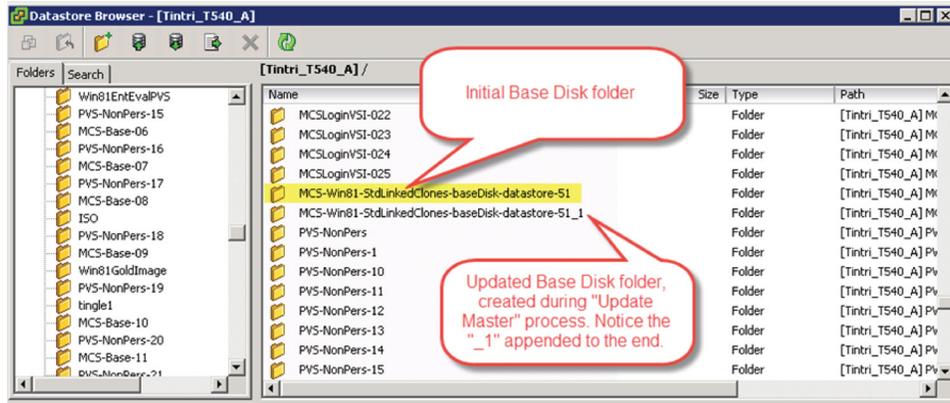


From a timeline perspective, read the vCenter tasks from the bottom up, taking note of start and completion times for all of the steps.

1. MCS takes a VMware snapshot of the Master VM.
2. MCS creates a Temp VM (XD-Temp-4...), which is configured with the vmdk from the original master image that was just snapshotted.
3. It then clones the Temp VM (XD-Temp-4...) to another VM, using some the name supplied in the MCS creation wizard: "MCS-Win81-StdLinkedClones", and then appends "-baseDisk-datastore-51" to it. The VM is created with this name so that the folder and base_disk.vmdk inside have the naming convention specified by Citrix MCS.
 - If you check the timestamps on the clone, the procedure was completed in only in **3 seconds!** Not only was it fast, but an extra ~10 GB (approx. footprint of base image) was NOT consumed from storage.
 - Assuming no VMware snapshot existed on the master image VM (Win81LoginVSI-MCS-base), VAAI is leveraged for cloning. BUT, if this image was already used for MCS, it likely has multiple VMware snaps on it from previous iterations. If you do NOT see the same result, and you've confirmed there is no VMware snapshot on your master VM, ensure all the hosts in your cluster have the Tintri VAAI Plugin installed, as mentioned in the pre-requisites.
4. After cloning, the new VM (**MCS-Win81-StdLinkedClones-baseDisk-datastore-51**) is reconfigured to have the disk removed.

5. MCS issues a command to delete the VM is deleted.

- By first removing the disk from the VM, it prevents the base disk (.vmdk file) and folder from being deleted when the VM is deleted. The result is all of the VM files are purged, but the base disk sticks around, as well as the descriptive name selected for it during the MCS creation wizard.

6. A new VM is then created, called "Preparation – + "name specified in MCS": "**Preparation – MCS-Win81-StdLinkedClones**".

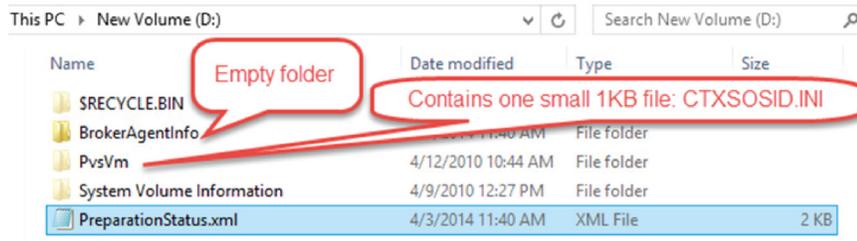
- The Prep VM is created with the amount of vCPU and RAM we specified in MCS creation wizard:



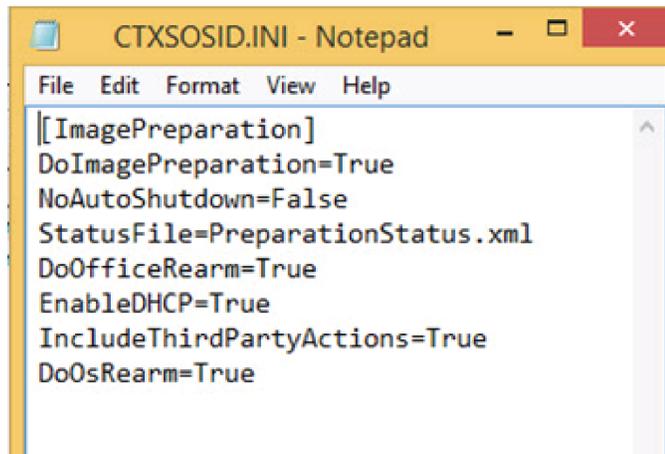
- The new prep VM is reconfigured to add our basedisk created in earlier steps (first reconfigure seen in the task log). Unlike all the linked clones that will be created, the vmdk is NOT mounted as Independent: non-persistent. The contents of the base disk still need to be altered to make it suitable for multiple MCS linked clones.
- The prep VM is reconfigured again (2nd reconfigure seen in task log) and a 2nd vmdk is added. This time, the .vmdk file added is a small disk file called "prepare-identity.vmdk". The file is provisioned as 16 MB, but is only 16 KB on disk and holds a simple file structure with a couple files to be used for the preparation process.

9. The Prep VM is then powered on... I “think” this is to validate that the clone process was successful and a single clone comes up without any issues.

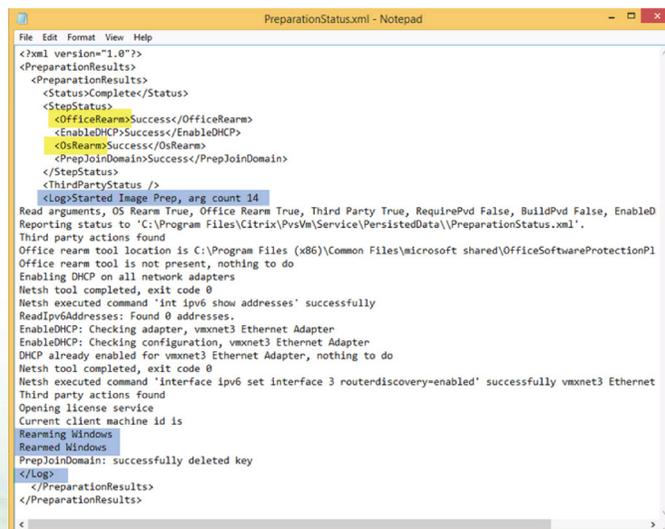
- During this process, we manually copied the prepare-identity.vmdk file that was found in the Preparation VM home folder so we could inspect the contents before it gets deleted. This is what was found:



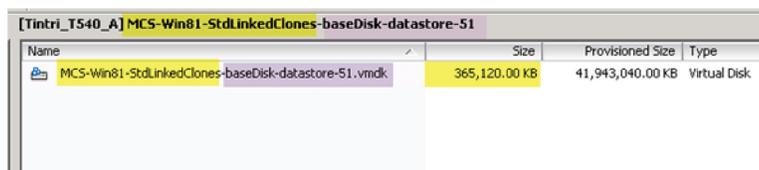
- Contents of PvsVm\CTXSOSID.INI:



- Contents of .\PreparationStatus.xml

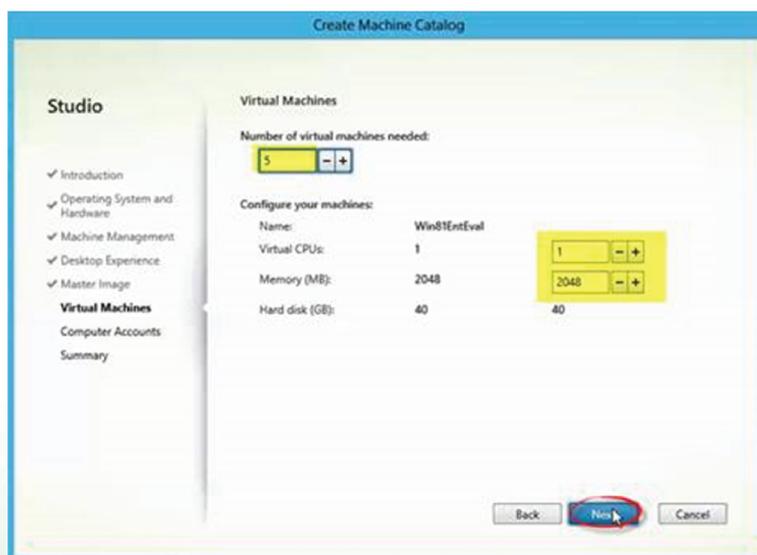


10. After the Preparation VM is powered on it is prepared with the information found in the identity disks.
 - The OS and Office installations are Re-armed so that each MCS linked clone created later is ready to run with their own custom identity files, similar to how sysprep works, but without requiring time to sysprep on every boot.
11. After in-guest processing (Windows preparation process) of the powered-on Prep VM, the Prep VM the VM is shutdown from within the guest (which is why a power-off command does not show in the log).
12. Once the Preparation VM is powered off, it is reconfigured to remove the base disk (to prevent deletion).
13. After reconfiguration completes, the Preparation VM is deleted from disk.
14. At this point, the base disk is left behind and has now been prepped and is ready to be a read-only disk that all MCS linked clones will mount as their C: drive.
 - Notice the size of the vmdk file in the store? This is because VAAI was leveraged and this disk is the product of a Tintri Clone. The 365MB shown is attributed to data change in the vmdk during the preparation process while mounted exclusively to the Preparation VM in the previous step:



Name	Size	Provisioned Size	Type
MCS-Win81-StdLinkedClones-baseDisk-datastore-51.vmdk	365,120.00 KB	41,943,040.00 KB	Virtual Disk

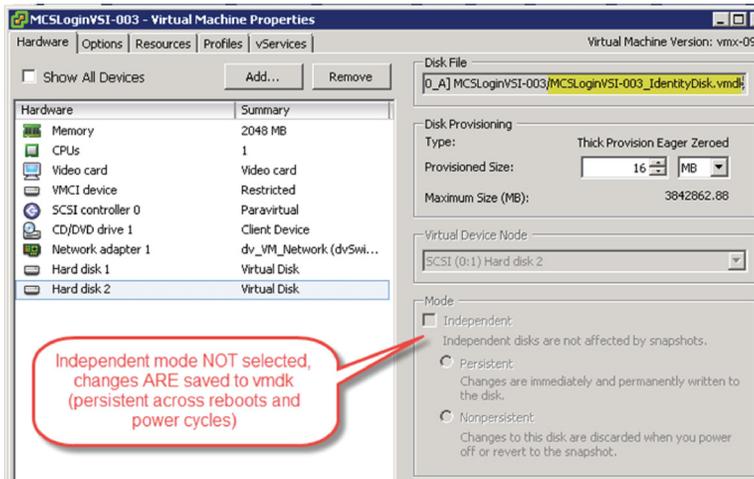
15. Next, VMs are created, 3 at a time, using the naming convention specified during the MCS creation wizard.
16. After creating the VMs, they are reconfigured 3 times each (as seen in the Task log):
 - **First Reconfiguration** - VMs are sized according to the specs provided in the wizard:



- **Second Reconfiguration:** The base disk ([Tintri_VMStore]\MCS-Win81-StdLinkedClones-baseDisk-datastore-51\MCS-Win81-StdLinkedClones-baseDisk-datastore-51.vmdk) is

mounted in the VM as SCSI0:0 in “Independent –Non-Persistent” mode (i.e. read-only). This is to prevent changing/locking the base disk and all write activity is redirect to a REDO file in the VM’s home folder, similar to the writes are redirected to a VMware snapshot file. The difference between the REDO file and a standard snapshot is that the REDO file is deleted when the VM is powered off or reset. And change that took place in the VM is lost, and all space consumed by the REDO file is released back to the datastore.

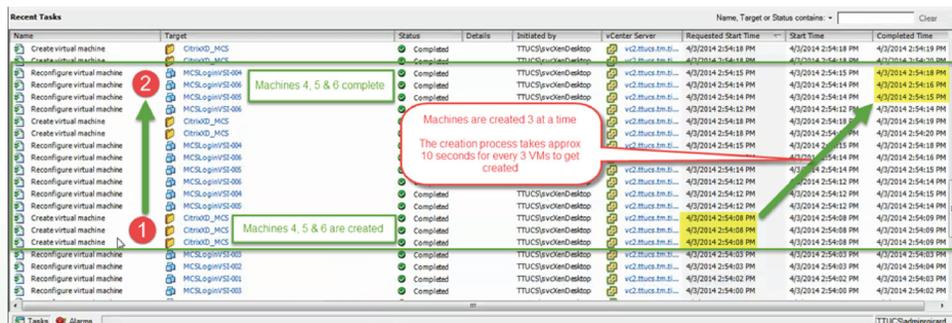
- **Third Reconfiguration:** A small (16 MB) identity disk is added, similar to the one added to the Preparation VM. The disk is placed in the home folder of the VM and is named “%VMNAME%-IdentityDisk.vmdk”:



- The contents of this identity disk are at the heart of what allows machines created with MCS to be non-persistent linked clones that can have their master disk updated, yet still persist and keep an identity on the network. This disk will be explored further, after we wrap up analyzing the tasks involved during the creation process.

17. The process in the previous step (Create – Reconfigure – Reconfigure – Reconfigure) is repeated, 3 new VMs concurrently, until the total number of VMs specified in the MCS creation wizard have been created.

18. Taking a look at the task log as it scrolls (see image below), a set of 3 VMs are created and fully configured in approximately 10 seconds (~3 seconds per VM):



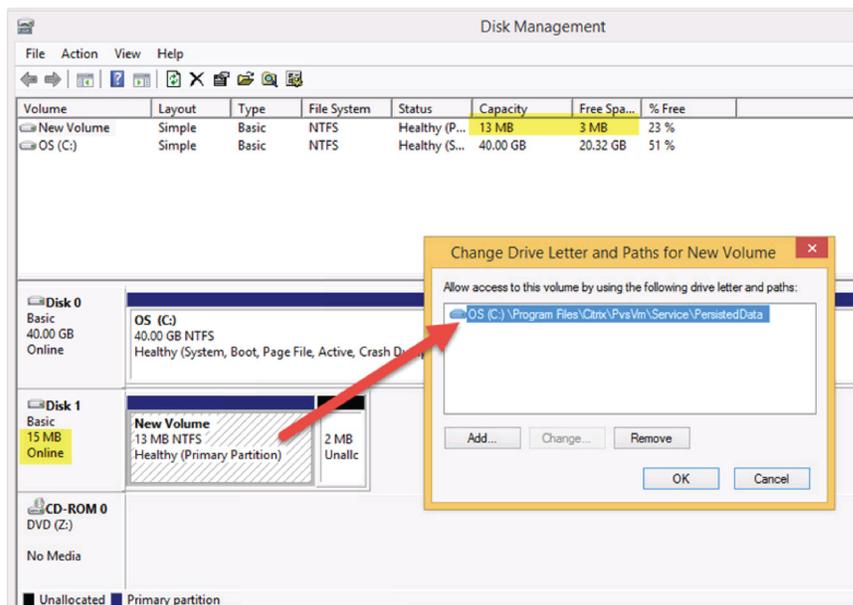
At this point, you now have a complete set of VMs in your MCS-based Machine Catalog. To use this, add them to a delivery group and they will power-on as user-demand for them dictates.

Pandora's Box – The Identity Disk

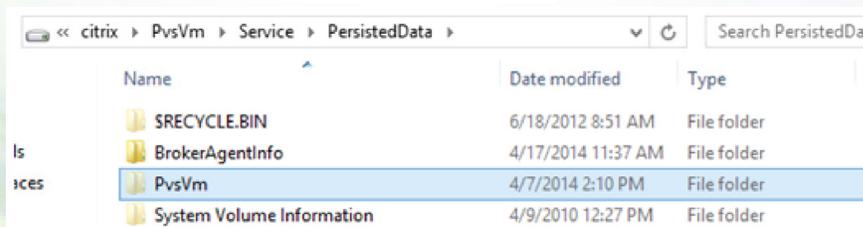
You're curious... what is in that small identity disk that gives each VM its own unique and persistent identity on boot-up? We've dug in and had a look for you. We won't cover the contents in complete detail, but you should get a good understanding of the key info held in these disks and can explore Citrix's KB article to learn more on your own.

WARNING: In order to view the files and folders below, Windows permissions needed to be altered. Altering windows permissions could lead to unintended consequences such as weakened security or functionality failures. When investigating the disks below, the .vmdk files were mounted non-persistently to another VM to investigate the files and prevent any unintentional changes to the files and folders within the VMDK. **Proceed with caution in your own environment.**

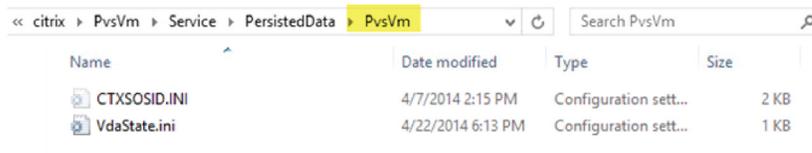
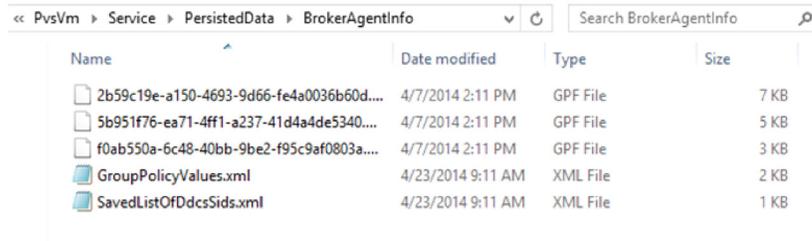
First, we have a look in one of our powered-on MCS Virtual Machines to see where the disk is mounted. In the screenshot below, you can see the disk is mapped to **C:\Program Files\Citrix\PvsVm\Service\PersistedData**:



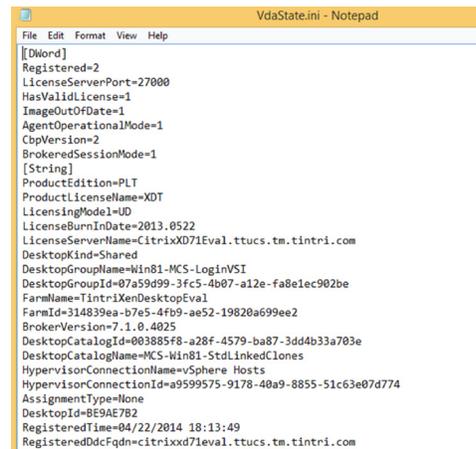
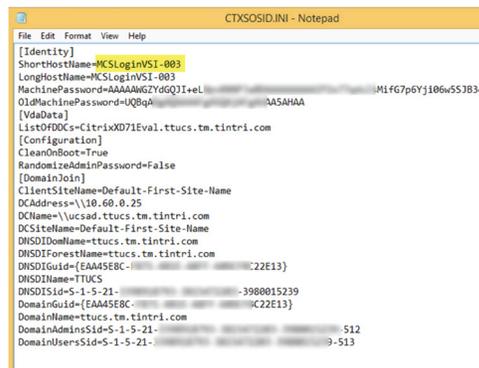
Looking instead that folder, we can see a similar layout of folders and files found in the "pre-prepare-identity.vmdk" disk that was mounted within the temporary Preparation VM. 2 Folders can be seen, **BrokerAgentInfo** and **PvsVm**:



Within these folders are files unique to the identity of this VM:



The BrokerAgentInfo folder contains information specific to GPOs applied to the VM. The PvsVm folder contains the answer files that are injected into the registry and other windows files that give this VM its unique identity and map it the computer account object in the Active Directory. Here are the contents of the two files:



Any further details about identity files are out of scope for this paper. As mentioned in the Gotchas section, MCS does not respond well to environmental changes. With all of this hard-coded information in each VM, you can understand why. How can you use this to your advantage?

-  **Do:** Review all of the information listed and make a note of it somewhere. If you plan on changing your environment in anyway, remember to keep your XenDesktop VMs in mind and test well to determine whether the proposed Infrastructure changes you want to make will break your VMs or not.

TIP: DR Planning: The above information is useful in DR planning and tests. If you're planning for a disaster and can't figure out why replicated machine catalogs don't work in the DR site, refer back to these files.

In most cases, altering these files is more pain and effort than simply re-building an MCS-based machine catalog. Worried about your AD accounts? Not to worry, the MCS creation process allows you to map your VMs to pre-existing computer accounts.

Consideration Factors

This section provides additional factors specific to this method that are useful in evaluating how best to run your XenDesktop infrastructure.

Space Utilization & Capacity Considerations

One thing to consider with MCS is how overall disk space and capacity is utilized. While MCS is quite lean, lightweight and efficient, especially when combined with VAAI features inherent in a Tintri VMstore, it's easy for space to get consumed quietly. The base disk folders are not visible in any VMware or Citrix console, and are only visible when browsing the datastore. If a base disk or subsequent base disk update (which is really just another copy of a base disk) is in use, space recorded as "Live VM" data in the VMstore space fuel gauge in the main dashboard:



However, with frequent updates, there is no automatic cleanup mechanism to delete old versions of base disks. Unused folders and vmdk files will start to collect in your VMstore and should be manually cleaned up from time to time. Once these files are no longer used, their space will be counted as "Other" space in the datastore. Check your "Other" space periodically and try to account for where it's coming from. We often create a folder to contain ISO images for installs, which also falls into this "other" category. If you keep your legitimate "other" data in a single location, such as a folder called "ISO_Files", you'll find it easier to figure out where other space is being wasted, such as old base disk folders and files.

- ✓ **Do:** When cleaning up old base disk folders and files, be diligent and investigate thoroughly to make sure the files are not needed before deleting.

VMware vSphere will prevent you from deleting a vmdk that is mounted to any running VM, which is a great safe guard, but doesn't help in cases where a disk is assigned to VMs that happens to be powered off.

Paging and swap files

If there are no memory constraints in your hosts, you will not need to worry about data being consumed by swap files. On a Tintri VMstore, the swap files are thinly provisioned, even if they appear to be the full size as the amount of RAM allocated to your VM when looking in the VMware datastore browser. More details on this can be found in [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#) document.

Space that collects in REDO files is reclaimed when a VM is power cycled. After creating a new machine catalog, prior to booting up the VMs, take note of the overall capacity free in your VM store and compare that to the value seen when your desktops are being used during peak

working hours. Depending on the size of user profiles, how they are configured, and other data specific to your image, the up & down capacity flux seen in a REDO per VM can vary greatly, anywhere from <100 MB to several GBs each.

Another component that contributes to space consumed by REDO file is the Windows page file, found in each VM. If VMs don't have enough memory assigned, memory will be paged, which can negatively impact performance.

 **Do:** From time to time, test your images with different amounts of memory assigned to each VM to determine how much RAM is the right amount and add RAM to your hosts in necessary.

Too much RAM will result in minimal page file usage, but will increase the likelihood that VMs will swap to the VMware swap file. Unfortunately, every master image is configured differently so there is no magic amount of RAM to suggest for your VMs. Try starting with 2 GB and adjust up or down as needed.

Persistence

The provisioning techniques in this paper have covered mainly stateless, non-persistent desktops configured in delivery groups as random pools. If your desktops require persistence across reboots and power cycles, determine if it's the whole desktop that needs to be persistent, or just the user's profile (covered in the next section).

You can achieve some persistence, as demonstrated in the section above that discusses the small 16MB identity disk assigned to each MCS catalog linked clone. In cases where you need to keep a persistent ID for each VM, such as a GUID for a global anti-virus deployment (ex. McAfee ePO), an OID for WSUS identification, or a unique database for Citrix Edgesite, refer to each vendor's specific docs on how this can be accomplished with MCS. Most large vendors with products common to traditional desktops and VDI deployments have procedures to follow that are designed to make them work in MCS deployments.

 **Do:** If persistence is important in your environment, consider using Tintri Native clones. Citrix Personal vDisks can help you, but are out of scope of this paper.

Audit & Compliance

Environments with strict audit and change control may need to provide evidence for certain administrative activities. MCS has audit logging built in, which is enabled by default.

 **Do:** Check the logs to ensure that actions you expect to be recorded are in fact recorded, and are sufficient evidence to meet the audit and compliance requirements of your environment.

Does activity within your desktops need to be maintained for security or compliance purposes? If so, you can't depend on the event logs of stateless systems as all traces of guest activity is lost on reboot or power cycle.

 **Do:** If the Event logs or other monitoring software within the guest require retention, consider a centralized logging solution to record and maintain yours logs outside of your desktops VMs.

Additional Factors

In addition to these, there are many other factors to consider that are common to all of the methods of provisioning that are detailed in the [XenDesktop 7.1 Provisioning and Machine Creation Deep Dive](#) paper and we recommend you review that paper as well.

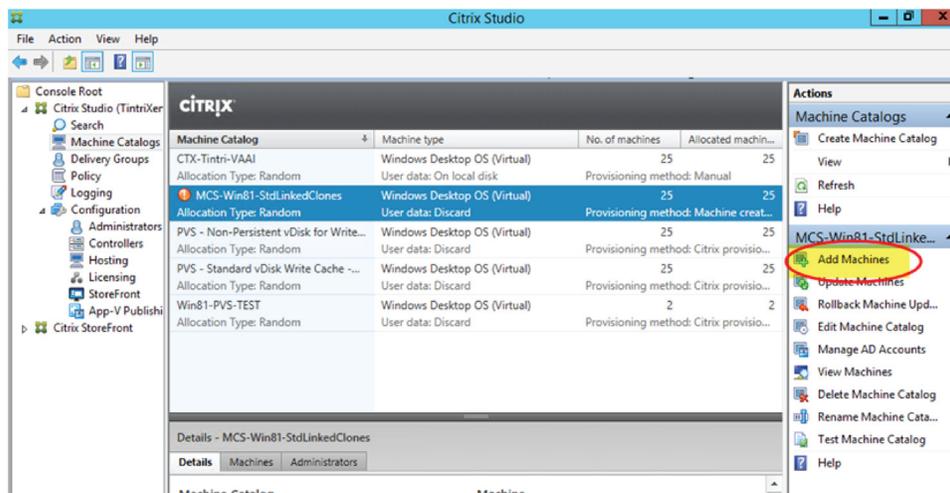
Management & Housekeeping

“The Only Thing That Is Constant Is Change” – Heraclitus

You don't need to be in IT very long to appreciate this quote! IT environments constantly change; staying on top of the changes is critical. This section is dedicated to raising awareness to some of the overall support & management costs.

Adding VMs to an MCS-based Machine Catalog

MCS makes it easy to add additional VMs to an existing machine catalog. From within Citrix Studio, simply highlight your MCS-based machine catalog, and choose Add Machines:



You'll be prompted for how many VMs you want to add, and whether to create new AD accounts or re-use existing ones. On the AD screen, if you are creating new ones, leave the name as-is and your new machines will follow your original naming convention, incrementing from where the last batch of VMs that were created had left off.



Do: Once you've created new VMs, don't forget to also add them to your delivery group(s) to make them available for users to connect to.

The Image Update Process

Updating an MCS master image can be straight-forward, but may result in some of the Tintri enhancements, such as VAAI, not being used.

WARNING: Any new Master VM selected with a VMware snapshot on it will prevent VAAI from activating and full copies of data are made (wasteful and slow).

To make changes to your master image, first make a Tintri clone of the current master image. This will allow you to keep a copy to track various versions, and also allow you to remove any VMware snapshots on your new clone (if any exist) without losing a way to rollback any acci-

dental changes when modifying the master. Boot up your Tintri clone and make the necessary configuration changes.

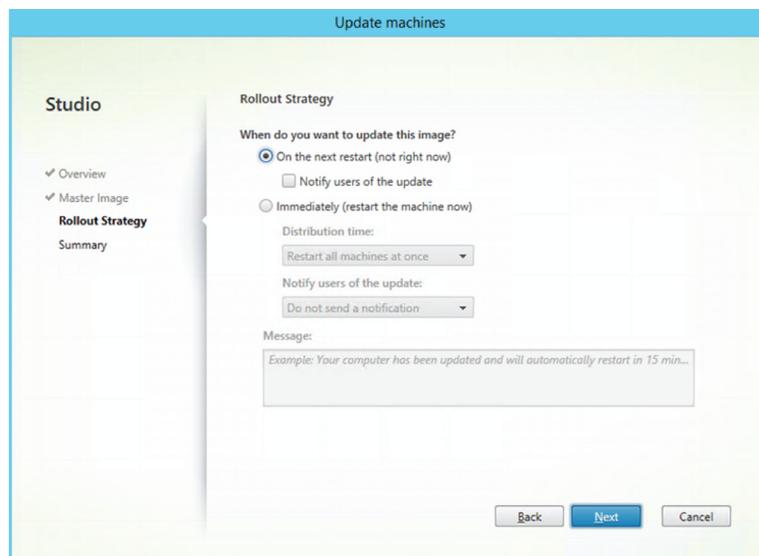
TIP: Prior to completing your configuration changes to a new Master image, always download the latest anti-malware definitions and run a full scan before shutting down the new master. More information can be found in [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#) document, under the section that discusses the impact of aging images.

With a new Master image prepped and ready to deploy to your MCS machine catalogs, first apply it to a controlled test group of VMs. After successful testing, proceed to updating your production catalog.

TIP: If your change control procedures don't already include test steps, we suggest you have an additional "Test" Machine catalog with a few machines assigned to various test users within your organization that can be used to test images before going live to all users in production.

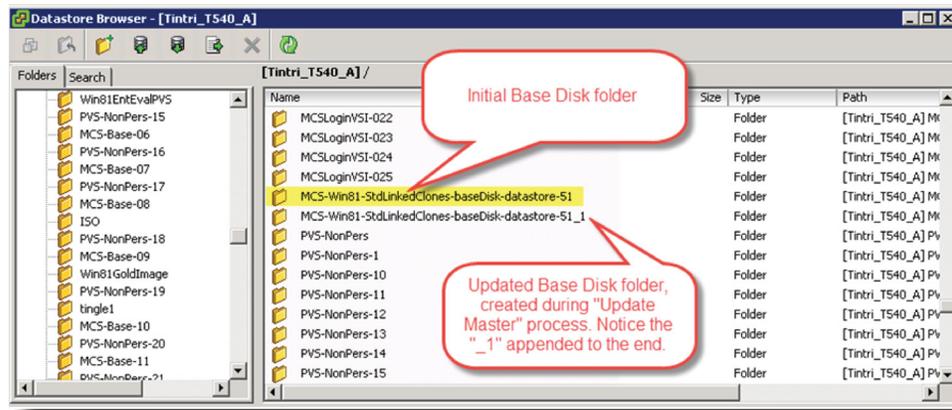
To use apply a new Master image to your machine catalog, simply choose "Update Machines" on your machine catalog within Citrix Studio. The process is similar to the initial creation process and will prompt you to navigate within your hosts/clusters to select a new Master VM.

Once you've selected a new VM to become your new master image, you are presented with several options for how to schedule the update. We recommend you choose "On next reboot", and coordinate reboots using some automated scripting.



A sample "rolling reboot" process is described in the Management & Housekeeping section of [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#). Essentially, it's a suggested process that puts a percentage of your VMs into maintenance mode daily (to prevent users from connecting) and then reboots them once any existing sessions have disconnected. Repeat this process daily until all VMs in your catalog have been updated. When new updates are required, simply set the VMs to update on next reboot and they will automatically get updated when their day of the week/month comes up for reboot.

After selecting a new Master image, a new base image folder containing a new vmdk file will be added to your VMstore and can be seen in the Datastore Browser:



Upon the next reboot, your VMs will be reconfigured to use mount the new base disk non-persistently, instead of booting off the existing image

- ✔ **Do:** Staggering your updates is desirable as it gives an opportunity to catch any bugs or issues specific to the new master without having 100% of your user population stumbling on it at once.

If there are issues, the affected machines can be put in maintenance mode to prevent users from connecting to those machines, and then roll back to a previous image, as described below.

The Roll Back process

No matter how well you test the latest changes to a master image, sooner or later, something may slip through the cracks and negatively impact users. In these cases, rolling back to a previous last known good build is a necessity. More thorough and higher quality testing will certainly reduce the need to rollback, but is unlikely to eliminate it altogether, so rolling back changes needs to be considered when deciding upon which deployment method to choose.

What happens when a change goes wrong and needs to be rolled back? How many users are impacted? Will there be any downtime as a result? The answers to these questions go back to how you introduce your updates in the first place and your overall change control procedures.

There is an option on your MCS-based machine catalogs to “Rollback machine update”. This process works fine, as long as the previous VM still exists in your datastore and still has a VMware snapshot on it. This process has been listed in the Gotchas section because VMware snapshots can interfere with VAAI and the standard image update process.

In the event rollback fails due to changes made to the original VM used, use the regular update process instead, but choose a Tintri clone of one of your VMs configured with the previous image version. Don't have one? Create a new VM, use the datastore browser to find the VMDK from the base disk folder of the version you want to rollback to, and then copy the base disk .vmdk into the home folder of the VM just created. VAAI should automatically get invoked, so the copy process should be near-instant and no additional storage should be consumed. Edit your new VM and mount the vmdk you just copied so that your VM is configured as any standard VM without any special disk configuration (i.e. do NOT add your disk as non-persistent). Follow the update process and select this new VM as your new master image.

Maintaining Image Versions

MCS does not allow for multiple versions of an image the way that PVS does. Administrator options are fairly limited when it comes to versions: update to a new one (immediate, scheduled or on next reboot), or roll back to a previous one.

You can make your own versions by keeping master copies of every change using appropriate labels. In the event you want to roll back two versions, simply update to a new version, and reuse older version instead of provisioning a brand new image. Keep in mind the impact of booting up an old image and the catch-up process that is required, which is explained in more detail in [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#) paper.

Environmental Changes to the VMware Infrastructure

Do you have dependencies hard-linked between your XenDesktop infrastructure and the virtual environment it is running on? Most likely, the answer is yes! This is something to keep in mind to avoid problems in your XenDesktop environment that are caused by seemingly innocent changes to the environment.

MCS doesn't handle changes well at all: cluster names, datastore names, network name changes can all "break" the MCS environment. In short, in order to be able to rename a VMware datastore successfully (and still have a fully functional XD deployment) we need to delete the hosting resources and re-add them. Easier said than done - there are dependencies from Delivery groups that prevent deletion of the resource. In order to remove the dependencies, the MCS machine catalogs needed to be removed first, but can't be removed if they are in use (ie. shut-down all your virtual desktops!).

And to make things worse, if you've already renamed your datastore, you won't be able to delete the machine catalog unless the old datastore name is found in order to complete the process. Once all dependencies are off the original storage resource, you can remove it. Once removed, you can rename your datastore and then re-add a storage resource back under Configuration - Hosting and then setup up Machine Catalogs again on the new resource.

Similar gotchas apply to changing vCenter name and/or SSL certificates associated with it.



Do: Test and be diligent, and add these gotchas to a checklist of things to check when undertaking a migration, major upgrade, or other changes that are required to stay current and supportable.

Data Protection & Disaster Recovery (DR) Overview

Backup and DR for XenDesktop environments need to consider more than simply ensuring good copy of data is available somewhere, such as a backup tape or offsite location. Desktops aren't often backed up for their unique data, but are instead the tools (OS + Applications + Configuration) that allow users to access their data (databases, email, application-specific files, etc.). A full and comprehensive DR strategy is out of scope for this paper, but this section provides some direction on what should be considered for backups.

XenDesktop infrastructure consists of 3 major components:

- XenDesktop Infrastructure Servers – XenDesktop Controllers, StoreFront, SQL servers, etc.
- **XenDesktop Machine Catalogs – The virtual workstation VMs.**
- Dependencies – User profiles, redirected folders, data, and the rest of the corporate infrastructure (i.e. the goods that users access their virtual desktops to access).

In the paper, we'll look only at the Machine Catalog. In [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#) paper, we'll cover some backup & DR aspects of XenDesktop Infrastructure Servers, but the dependencies are largely out of scope other than mentioning them so they are not overlooked.

Backup & DR of MCS-based Machine Catalogs

MCS provides an easy to use, quick, and space-efficient method of provisioning machine catalogs. Unfortunately, it does not allow for a simplified DR process as PVS or Tintri Native Clones do. One of the objectives we set out to achieve with our DR copies is to maintain a single set (i.e. production) of XenDesktop infrastructure VMs & machine catalogs and use ReplicateVM to keep copies on a remote site that can be brought online in the event of a disaster in the production site.



Do: Consider using 2 or more active sites in a global Citrix topology, each configured with their own unique machine catalogs.

However, this means you have a bigger footprint of active VMs to maintain (patching, AV updated, operational, security, licensing, etc.). Please refer to Citrix KB articles, blogs, and other online references for more information on how to devise a DR plan for your Citrix environment.

Conclusion

We hope that you have found this deep dive on MCS useful and enlightening. While there is no “right” or “wrong” method to use for your XenDesktop deployment, there are pros and cons associated with each and we recommend you explore the different options available in order to find what works best for your organization.

The overall ease of use of MCS is the most compelling reason to choose it as the deployment method. On the downside, the ease of use hides a lot of information hard-coded into your MCS deployment, which may make troubleshooting more difficult in the event that something does not work as planned. MCS is prone to errors and failures caused by environmental changes, so be sure to test carefully before making seemingly harmless changes within your VMware infrastructure.

Additional Resources – Reference URLs

Tintri Links

- [XenDesktop 7.1 Quick Start Guide](#) – Covers all the prerequisite tasks required to get a XenDesktop environment up and running on Tintri storage.
- [XenDesktop 7.1 Provisioning & Machine Creation Deep Dive](#) – Referenced throughout this paper as “the main” provisioning deep dive paper and contains information that applies to all XenDesktop provisioning methods, and compares the method discussed in this paper to alternative methods.
- [XenDesktop 7.1 Provisioning Services \(PVS\) Deep Dive](#) – A similar paper to this one that focuses on PVS as an alternative method of provisioning virtual desktops.
- [XenDesktop 7.1 Tintri Native Clones Deep Dive](#) – A similar paper to this one that focuses on using Tintri Native Clones as an alternative method of provisioning virtual desktops.
- [XenDesktop Best Practices white paper](#).
- [NFS Best Practices](#).
- [Additional resources](#).

Citrix Links

- [Citrix XenDesktop eDocs](#).
- Citrix Blog series that focus on image management:
 - [Part 1: PVS - The myths and tales of image management](#).
 - [Part 2: MCS - The myths and tales of image management](#).
 - [Part 3: MCS vs View Composer](#).

Other

- [VMware FAQ for VAAI](#).
- [Login VSI – Load Testing tool for VDI environments](#).
- [FSLogix](#).

Appendix A – Environment Details

This guide was written to be applicable to any server hardware and all supported versions of VMware vSphere. In case your own results vary from what you've seen in this paper, here is a high-level overview of the lab equipment used:

Hardware

Storage: Tintri VMstore T540 running Tintri OS v2.1.2.1

Servers: Cisco UCS B22-M3 blades

Software

VMware vSphere 5.1 U2 – Build 1612806

Tintri vSphere Web Plugin v1.0.0.1 (includes the Best Practices Analyzer tool)

Server OS: Windows Server 2012 R2 – Datacenter Edition

Desktop OS: Windows 8.1

Citrix XenDesktop 7.1

Citrix PVS Provisioning Services 7.1