WHITE PAPER

Citrix® Provisioning Services™
Using Read-Only vDisk Storage

www.citrix.com

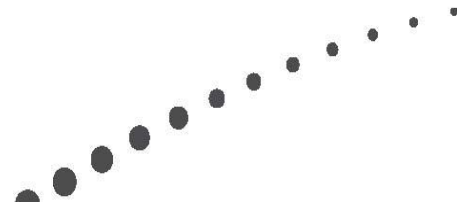

Table of Contents

Introduction

Citrix Provisioning Services™ uses streaming technology to dynamically deliver desktop images – operating systems and software stacks – on-demand to physical or virtual machines from a network service. In Standard Image Mode, Provisioning Services enables IT Administrators to create and manage a single, shared disk image that is streamed simultaneously to any number of machines. This dramatically simplifies and reduces the costs of desktop management.

- Operational costs are reduced. A single image is managed and patched rather than hundreds or thousands of individual desktops.
- Storage capacity requirements and costs are reduced by as much as 90 percent. Instead of allocating full image storage capacity for all of the desktops, IT need only allocate capacity for the single shared image (5 – 20 Gigabytes) and, optionally, write caches (less than 1 Gigabyte per desktop). Unlike deduplication, storage requirements start small and stay small throughout the deployment.

This document describes how to leverage the Read-Only vDisk Storage feature introduced in release 5.1 to further simplify Provisioning Services deployment and improve scalability and boot-time performance when combined with SAN storage such as the Dell EqualLogic PS Series SAN.

Dell EqualLogic PS Series SANs provide the flexibility to create LUNs and set them as read-only using the Group Manager console. In addition, the Citrix XenServer Adapter integrates EqualLogic PS Series SAN control interfaces into the Citrix XenCenter Management Client for a unified management experience.

This document assumes the user is familiar with the Provisioning Services, the Microsoft iSCSI Initiator software and the EqualLogic Group Manager GUI console. More detailed information on Group Manager is provided in the *Dell EqualLogic Group Manager Administration* guide, included as part of the PS Series Software Documentation set with each PS Series SAN array.

Read-Only vDisk Storage Feature Overview

Administrators can deploy provisioning services with a variety of storage solutions for vDisks and write-back caches. When deploying in a highly available configuration, with multiple Provisioning Services hosts, a shared storage solution, such as a SAN, NAS, or file share, must also be deployed. Many customers prefer to use SAN storage to maximize scalability and performance. Provisioning Services requires that SAN LUNs be formatted with a file-based system. However, the Windows NT File System (NTFS) does not support concurrent read/write access (doing so results in corruption of the volume). Instead, customers have traditionally deployed one of three solutions in conjunction with their SAN:

- NAS Gateway
- Cluster File System
- Windows Cluster Services

While all three of these solutions are viable, each adds a layer of management complexity, runtime overhead, and cost. With the 5.1 release, Citrix is adding a new feature to Provisioning Services: Read-Only vDisk Storage. With this feature, administrators can create read-only LUNs for vDisk storage that can be concurrently accessed by multiple Provisioning Services hosts without also deploying one of the three solutions listed above.

Use of the Read-Only vDisk Storage feature is limited to specific image mode configurations. Regardless of the configuration, only the vDisks may be placed in the shared NTFS volumes. Write-back caches must either be stored on the target device or in volumes that use one of the shared access solutions listed above. The following table lists the supported modes and deployment restrictions.

Image Mode	Cache Location	Allowed?	Restrictions
Private Image Mode		No	Private images disks must be stored on read/write volumes.
Standard Image Mode Difference Disk Mode	Server disk or Encrypted on server disk	Yes	Separate shared read/write write-back cache storage location required. vDisk properties cannot be modified while LUN is read-only; vDisk cannot be mapped on server.
Standard Image Mode	Device RAM	Yes	vDisk properties cannot be modified while LUN is read-only, vDisk cannot be mapped on server.
	Device HDD or Encrypted on device HDD	Yes	Fallback to cache on server disk does not function if device HD is not found or fails. vDisk properties cannot be modified while LUN is read-only, vDisk cannot be mapped on server.

As noted in the table above, the main limitations to placing vDisks on read-only storage are as follows:

- Private image mode boot from read-only storage is not allowed
- If cache on server disk is desired, a separate shared storage location that has read-write access is needed for the write cache files.
- Modifying the vDisk properties is not allowed when the vDisk storage location is read-only.
- Mounting the vDisk on the server is not allowed when the vDisk storage location is read-only.

Prerequisites and Example Deployment Architecture

The following items are prerequisites to using this new feature:

- The Provisioning Services hosts meet the minimum system requirements listed in the Provisioning Services installation and configuration documentation.
- The Microsoft iSCSI initiator software is installed on all Provisioning Services hosts having access to the SAN.
- The vDisk files being placed on the read-only shared LUN(s) have already been created and reside on a normal read-write storage location. Creating vDisk files in place on the LUN is more difficult than pre-building the VHD files in a normal read-write store and subsequently copying them to the shared LUN. Therefore this document describes the procedure assuming the vDisk files have been pre-made and reside in a normal read-write storage location. For documentation on creating vDisk files refer to the Citrix Provisioning Services 5.1 Administrator's Guide available at support.citrix.com.
- The SAN being used has the ability to set a LUN up for shared read-write access or shared read-only access without requiring a shared file system front end. Normally, using a LUN in shared read-write access mode without a shared file system front end results in a corrupt NTFS volume. Limiting the LUN access to read-only (configuring read-only volume access on EqualLogic PS Series via Group Manager) circumvents this problem.

The diagram below illustrates a typical deployment architecture meeting these requirements.

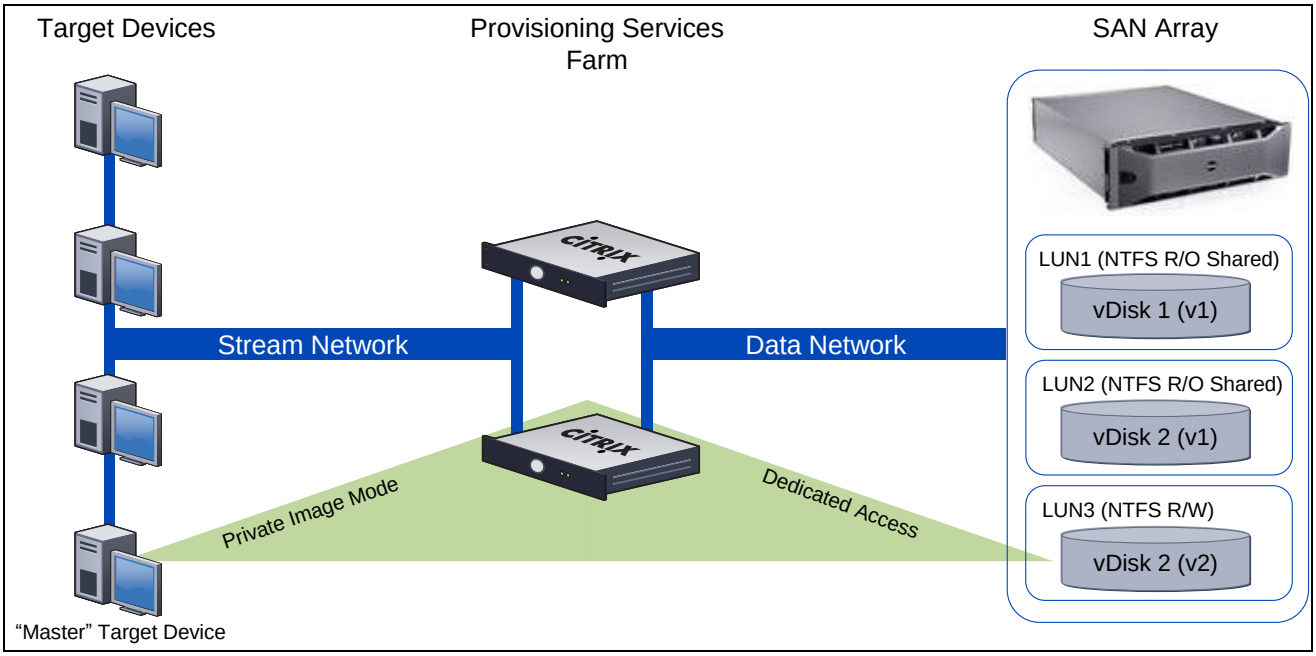

Figure 1 Typical Deployment Architecture

It is recommended, though not required, that separate read-only LUNs be created for each shared VHD file. This allows VHD files to be updated and taken on and off line independently of each other. Detailed recommendations can be found in the *Maintaining vDisk Images* section below.

Initial Configuration and Operation

This section provides guidance for configuring Provisioning Services and Dell EqualLogic PS Series arrays for use with the Read-Only vDisk Storage feature. Steps for configuring other SAN solutions are similar, using the management tools provided with them.

On the SAN

1. **Create a volume for each shared vDisk** using the EqualLogic Group Manager GUI. Make sure each volume is large enough to hold the associated VHD and associated PVP file.

2. **Set the access type for the volumes to read/write – shared**, again, using the EqualLogic Group Manager. The volume is made read-only through the NTFS attributes, not through the SAN access rights. While it is possible to configure the volume as read-only shared using the Group Manager GUI, it is not required for effective protection of the vDisks. Therefore this document only describes the process when the volume is set for read/write – shared access.

On the Provisioning Services hosts

For each shared access volume created in the previous steps, perform the following steps from the Provisioning Services hosts and console.

3. **Using the iSCSI Initiator, Log on to the volume from a single Provisioning Services host**.
Important Note: *Do not log on to the SAN Volume from more than one server simultaneously until the volume has been marked read-only using the instructions which follow.* Simultaneous host access through the iSCSI interface while the volume is in read-write mode causes corruption of the volume, resulting in loss of data and the need to reformat the volume. EqualLogic Group Manager enables restricted access to the LUN by specific IP addresses and iSCSI initiator names to help avoid simultaneous access by multiple servers.

4. **Format the volume with NTFS** through the Windows Disk Manager and assign a drive letter or mount point path. A mount point path is desirable if you have many LUN/Volumes exposed on a server because there will be no drive letter limitations. The drive letter/mount point should be identical on all hosts that access the volume; otherwise the host specific mappings must be specified using the Provisioning Services console. Once the volume is formatted and assigned a drive letter/mount point the volume should be accessible on the single Provisioning Services host as a read/write volume.

5. **Verify that all properties for the VHD/PVP file to reside on the volume are set as desired.** (such as HA enable, and so on)

6. **Copy the desired VHD and associated PVP file to the volume.** Lock files do not need to be copied. However, the PVP file must be copied along with the VHD file. The system cannot dynamically create a PVP file once the volume is set to read-only mode.

7. **Set the volume to read-only mode.** Close all Explorer windows that have access to the volume. Open a command prompt Window on the server that has access to the volume and run diskpart.exe. This starts an interactive session with diskpart.exe. Find the volume number for the volume by typing the following command:

```
list volume
```

Note the volume number of the target volume and select it by typing the following command:

```
select volume <volumeNumber>
```

where <volumeNumber> is the number of the volume identified in the previous command. Once the volume is selected, set the read-only attribute of the volume by typing the following command:

```
attributes volume set readonly
```

Verify that the readonly attribute was set correctly by typing the following command:

```
detail volume
```

Once this is completed, exit diskpart.

8. **Log off the volume on this server and then log on to the volume** again using the iSCSI initiator interface to force NTFS on the server to re-read the volume attributes so that it recognizes the volume as read-only. Ensure that the volume is marked as a persistent target, because this ensures the volume is accessible when the server reboots. It is now safe to mount the iSCSI volume on all Provisioning Services hosts.

9. **Mount the volume on all Provisioning Services hosts that need access**, marking the target for automatic connection using the iSCSI Initiator applet and Microsoft Disk Manager, as shown in Figure 2, below.

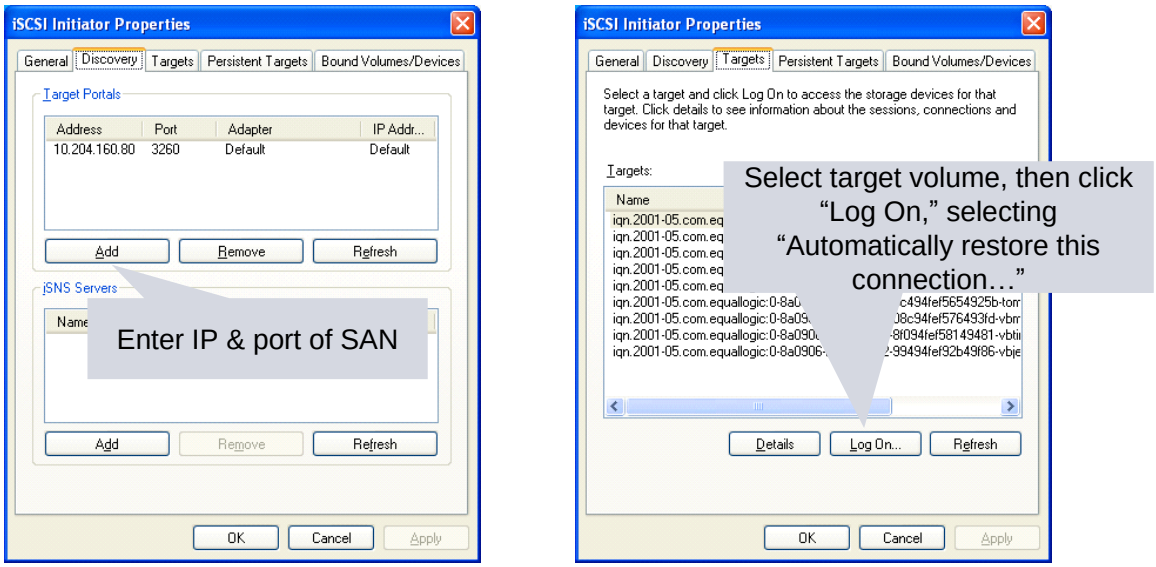

Figure 2 Mounting volume using iSCSI Initiator

10. **Set the same mount point/drive letter on all hosts** using the Disk Management snap-in of the Windows Computer Management tool to simplify configuration of vDisk stores in the Provisioning Services console.

11. **Mark the Provisioning Services Stream Service on all servers dependent on the iSCSI Service** to ensure that the volumes are available at the proper time if the server reboots while target devices are being streamed. This is done by editing the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\StreamService
```

    for the Streaming Service and adding a DependsOnService value pointing to the iscsiexe.exe service (MSiSCSI) as follows:

```
Value Name: DependOnService
Value Type: REG_MULTI_SZ
Value Data: MSiSCSI
```

12. **Create a vDisk store pointing to the drive letter/mount point for the volume** using the Provisioning Server console. Select which servers have access to the volume for this store.

    NOTE: If the vDisk being served from the shared read-only volume is set to "cache on server" or "difference disk" mode a "default write cache path" pointing to a separate shared read-write storage area must be specified. This shared storage area can be a Windows share or volume accessed through CIFS/SMB or a SAN volume formatted with a clustered/shared file system. If the vDisk is set for one of the "cache on target device" modes then no additional storage locations are required.

13. **Add the VHD to the Provisioning Services database** using the Store menu option "Add Existing Disk…" from within the Provisioning Services management console.

14. **Assign the VHD to target devices and boot them normally**. The VHD files on the read-only volumes always display in the Provisioning Server console as locked with the type "Read only media: Shared". You cannot remove this lock type. You cannot create a new vDisk on a store once it has been marked as read-only with diskPart.exe. You cannot edit the properties of the VHD once the store has been marked read-only.

Modifying vDisk Properties

vDisk properties cannot be modified while the SAN LUN location is marked read-only. To edit the vDisk properties or change/modify the vDisk files on the LUN use the following procedure:

1. **End all streaming sessions associated to the vDisk to be modified.** This can be done by shutting down target devices that have been assigned to the vDisk or assigning/booting them from a different vDisk.
2. **Log off all but one of the Provisioning Services hosts from the volume containing the vDisk** using the iSCSI initiator. Alternately, on some operating system types, the diskpart.exe utility can be used to mark the volume as offline on all Provisioning Services hosts.
3. **Mark the volume as read-write** using the diskpart.exe utility. This can be done using the same commands in *Initial Configuration and Operation* step 7, substituting the following command to clear the read-only attribute:

```
attributes volume clear readonly
```

4. **Log-off and re-login to the volume** using the iSCSI initiator as described in *Initial Configuration and Operation* step 8. Alternatively, mark it as offline and then online using diskpart.exe.
5. **Edit the VHD file attributes** through the Provisioning Server console normally and/or copy new files to the volume.
6. **Mark the volume as read-only** using the diskpart.exe utility as described in *Initial Configuration and Operation* step 7.
7. **Log off and re log on to the volume** using the iSCSI initiator as described in *Initial Configuration and Operation* step 8. Alternatively, mark it as offline and then online using diskpart.exe.
8. **Log on the remaining Provisioning Services hosts to the volume** using the iSCSI initiator.

Maintaining vDisk Images

Two approaches can be taken when performing vDisk images maintenance (operating system/application patches and installs, updating Provisioning Services drivers, and so on): update in place or deploy in new volume.

Deploy In New Volume

This approach follows the same procedure as described in *Initial Configuration and Operation*. The primary advantages to this approach are that it minimizes the down-time required for vDisk maintenance by keeping multiple versions of the vDisk on the array. However, this approach does require additional storage capacity on the array. The amount required depends on the size of the vDisks and the number of version copies that are maintained. To use this approach, perform the following steps:

1. **Follow steps 1 through 5 described in *Initial Configuration and Operation* above**.
2. **Copy the selected vDisk** from an existing read-only volume to the newly created read-write volume.
3. **Create a vDisk store pointing to the drive letter/mount point for the volume** using the Provisioning Services management console, specifying that only the host that has logged on to the volume will access it.
4. **Add the vDisk to the Provisioning Services database** using the Store menu option “Add Existing Disk…” from within the Provisioning Services management console.

5. **Set the vDisk to Private Image Mode** using the Provisioning Services management console.
6. **Assign the vDisk to a single target device.** This device applies the updates to the vDisk.
7. **Boot a target device from the vDisk.** Select a target device that is assigned to the vDisk and boot it normally using Provisioning Services.
8. **Update the vDisk as needed using the running target device.** Operating system and application updates can be applied as they would for a system booting off the local hard disk.
9. **Shut down the target device**.
10. **Follow steps 7 through 10 described in *Initial Configuration and Operation* above**.
11. **Grant access to the new vDisk store to the other hosts**, using the Provisioning Services management console.
12. **Assign the vDisk to target devices and boot them normally**. The VHD files on the read-only volumes always display in the Provisioning Server console as locked with the type "Read only media: Shared". You cannot remove this lock type. You cannot create a new vDisk on a store once it has been marked as read-only with diskPart.exe. You cannot edit the properties of the VHD once the store has been marked read-only.

Update In Place

This approach leverages the same procedure described in *Modifying vDisk Properties*. The primary advantages to this approach are that it does not require creation or configuration of any new volumes on the array and minimizes the number of image copies residing on the array, reducing overall storage requirements. However, this approach generally requires a longer outage time for target devices associated with the vDisk being updated, because there is no overlapping of versions available on the SAN. To use this approach, perform the following steps:

1. **Follow steps 1 through 4 described in *Modifying vDisk Properties* above**.
2. **Set the vDisk to Private Image Mode** using the Provisioning Services management console.
3. **Boot a target device from the vDisk.** Select a target device that is assigned to the vDisk and boot it normally using Provisioning Services.
4. **Update the vDisk as needed using the running target device.** Operating system and application updates can be applied as they would for a system booting off the local hard disk.
5. **Shut the target device down.**
6. **Follow steps 6 through 8 described in *Modifying vDisk Properties* above**.

Frequently Asked Questions

When should I choose this configuration?

Maximum benefit from the Read-Only vDisk Storage feature can be obtained in environments that use SAN for vDisk storage and are using client-side write-back cache (disk or RAM). In these environments, using this feature completely eliminates the need to deploy shared/clustered file system software, reducing deployment costs and complexity, and maximizing scalability and performance.

How does this feature improve performance?

Using a read-only volume for vDisk storage eliminates the overhead of the normal Provisioning Services disk locking mechanism. This impact is primarily visible during the target device boot. Tests have shown that boot time performance is improved 25 to 50 percent depending on the number of devices booting concurrently (greater improvement is achieved with higher device counts).

I need to use server-side cache with my SAN storage. Can I still use this feature?

Yes, however you must deploy a shared access solution such as a NAS Gateway, cluster file system, or Microsoft Cluster Services to enable shared access to the volume containing the write-back cache files. Future versions of Provisioning Services will eliminate this limitation.

Can I use this feature with storage solutions other than SAN?

Generally, this feature is not needed with other solutions (such as NAS or Windows File System shares) because access to these devices is through CIFS or SMB which support safe shared volume access. SANs use block-based IP or fiber channel protocols, which do not provide mechanisms to make shared access to volumes safe.

Conclusion

The Read-Only vDisk Storage Feature of Provisioning Services 5.1, when combined with SAN storage such as the Dell EqualLogic PS Series gives systems administrators the ability to manage the maximum number of target device systems with minimal investment in back-end supporting infrastructure. Future versions of Provisioning Services will expand on this feature, removing many of the current limitations and automating setup tasks. However, by following the steps described in this document, administrators can achieve most of the core benefits in scalability, performance, and simplified management today.

More Information

Information about Citrix Provisioning Services is available at
http://www.citrix.com/English/ps2/products/product.asp?contentID=1297541

Information about Dell EqualLogic PS Series SANS is available at:
http://www.dell.com/content/topics/global.aspx/power/en/dell_equallogic_ps_series_iscsi_sans?c=us&l=en