

## VDI and snapshots: A winning combination

By Ray Lucchesi  
January 2009 InfoStor

The proliferation of user desktops is rapidly becoming an administrative quagmire for today's data centers. However, desktop virtualization products have recently emerged to reduce the negative economic impact of deploying a multitude of desktops.

VMware Virtual Desktop Infrastructure (VDI) is one such package, and provides desktop virtualization by using virtual machines to execute software while only running a remote desktop client on the user's desktop. Consequently, VDI provides easy desktop consolidation and reduces the overall cost of deploying enterprise desktop services.

But VDI implementation has at least one prominent, yet resolvable, concern. Specifically, VDI requires copies of boot image data on expensive, centralized storage for each virtual desktop. Snapshot services help resolve this concern by more efficiently replicating the boot image data, and thus significantly shrinking the incremental storage costs per virtual desktop.

On the other hand, not all VDI data is suitable for snapshots. For instance, virtual desktops also require storage for user workspace. Because user workspace data is unique to each desktop, it makes little sense to replicate via snapshots. Moreover, this data could easily be centralized and co-located on shared VMFS datastores.

### Snapshots for boot images

Normally, to create multiple boot images, boot image data must be copied. One boot image typically takes 10GB to 20GB of space and would take only a short time to copy. However, making 100 copies of 20GB boot images would take more than five hours at 100MBps (possibly double that if the source and target are on the same storage) and take 2TB of space.

A more effective approach is to use snapshot technology to create the multiple boot image copies. Snapshots provide for instantaneous, space-efficient replication of data on most storage subsystems.

There are three different approaches for snapshots that can be used to support VDI boot images:

- **Snapshots of raw device mode (RDM) physical mode LUNs**--VMware supports RDM, where a virtual machine boots directly from a physical LUN, bypassing VMware I/O virtualization. An administrator would typically create the boot image on a single LUN, and then this LUN could be snapshot as many times as needed.
- **Snapshot of a single boot image datastore**--The virtual desktop boots from a ".vmdk" file that is the sole disk image file on a datastore residing on a single LUN. An administrator would create a VMFS datastore, create a Gold boot image on the datastore, and then repeatedly snapshot the LUN holding the datastore to create the requisite number of boot image copies.
- **Snapshot of a multiple boot image datastore**--In this approach, multiple boot images reside on the datastore and can be replicated via snapshots simultaneously. As with the snapshot of a single boot image, the administrator creates a VMFS datastore on a single LUN and then creates the Gold boot image. At this point, the Gold boot image would be copied to create the multiple boot images on the datastore. The multiple boot image datastore would, as a final step, be snapshot as many times as needed. For example, to create a datastore of 10 boot images, the Gold boot image would be copied nine times, resulting in 10 copies of the boot image for each snapshot.

## Snapshot concerns

Snapshots of VDI boot images are constrained by LUN or file system granularity and thus are a major concern. VDI boot image data is typically just one of many ".vmdk" files on a shared VMFS datastore, configured over a number of LUNs. Snapshotting this datastore would replicate all the ".vmdk" files along with boot image data, and unnecessarily consume LUNs, wasting valuable resources. As such, VDI boot images should be isolated in a single LUN datastore for snapshot purposes to maximize VDI utility and minimize storage consumption.

A concern arising from this LUN proliferation caused by snapshots is the recognized axiom that more LUNs necessarily means more work; i.e., more storage configuration, more backup changes, and more space management monitoring. However, boot image snapshot data is basically non-growing data, and as such should not add to space management problems. Furthermore, boot image backup changes are only done once, and boot image data does not need frequent backup because changes are rare. Finally, the configuration changes needed to define additional LUNs can be mitigated because these are subsystem-created LUNs in response to snapshot commands. Given these factors, adding boot image snapshot LUNs should require substantially less incremental maintenance than other, non-static LUNs.

For larger data centers supporting numerous desktops, the potential for LUN proliferation could be critical. Using snapshots of multiple boot image datastores could mitigate this LUN proliferation concern. For instance, VMware Virtual Infrastructure 3 (VI3) only supports a maximum of 256 LUNs per ESX server. Not all of these LUNs can be VDI boot images, as some must be used for end-user workspaces, non-desktop virtual machine storage, and VI3 software. With four boot images per datastore, 125 snapshots would support 500 virtual desktops and still leave more than 130 ESX LUNs for other storage requirements.

Multiple boot images per datastore may not always be advantageous, however. In fact, when restoring a single desktop image, such an approach may be slower. In this case, instead of a quick two-step "point-and-shoot" mount and boot of a new snapshot, it becomes a slower three-step process of mounting the appropriate Gold image or backup volume, copying the correct boot image to the LUN in use (potentially a slow process), and restarting the client.

A second issue with the multiple boot image per datastore approach is LUN-level monitoring. Here, the granularity of storage monitoring only allows viewing activity at a LUN level, and as such cannot monitor the activity of a single desktop but instead views the aggregation of all the desktops assigned to the LUN.

Most vendors use "copy-on-write" technology to provide storage subsystem snapshots. This technology copies data only as it's modified and thus is ideal for rarely changed boot image data. However, vendors vary widely in their support of snapshot services. Specifically,

- *Not all storage subsystems support large numbers of snapshots per LUN.* For example, HDS USP-V limits the number of copy-on-write snapshots to one per LUN; IBM DS8000 limits the number of FlashCopySE snapshots to 12 per LUN, and EMC Symmetrix limits the number of Timefinder snapshots to 16 per LUN.
- *Not all snapshots are writable, and often there are limits to the number of writable snapshots from the same LUN.* For example, 3PAR allows 128 writable and 500 read-only snapshots per LUN. NetApp has both read-only snapshots and writable FlexClone volumes, but actual limits are not readily specified.
- *Not all storage snapshots reserve the same amount of disk space.* Some subsystems can reserve up to 40% or more of the original LUN for snapshot space, although a few vendors reserve no additional space for their snapshots
- *Not all storage products support space-efficient snapshots.* For example, when a source LUN is modified, a new space-consuming block would potentially need to be created for each LUN snapshot. Some systems, however, provide a non-duplicative feature, so only one copy of the update occurs regardless of the number of writable snapshots, resulting in less storage consumption.

---

***First and foremost, using snapshots for boot image data speeds up end-user desktop deployment.***

---

### Snapshot benefits

First and foremost, using snapshots for boot image data speeds up end-user desktop deployment. Copying TBs of data normally takes hours, but when using snapshots, it takes just minutes. Another snapshot advantage is desktop boot performance. Finally, using snapshots results in considerable storage space savings. For example, assuming a 20GB boot image, 2TB of storage would be necessary to support 100 desktops. Using 99 snapshots of the same boot image, and thus still supporting 100 desktops, may only take 20GB of storage -- a 99% reduction.

Given the compelling benefits and resolvable concerns of snapshots for VDI boot images, administrators using VDI boot image snapshots still have a difficult task. Once they have created a boot image ".vmdk" file, they need to locate the VMFS datastore it resides on, determine the LUN holding the datastore, locate the physical storage this LUN resides on, issue the requisite subsystem-specific requests to snapshot the datastore, and then export the new LUNs to ESX. Following the export, the administrator must signal VMware to re-scan for newly snapped LUNs, and to re-signature the snapshot volumes for VDI virtual machines use. Finally, the administrator must clone the virtual desktop configuration and attach the virtual machine to use the newly created boot image.

Recognizing the tediousness of many of these configuration tasks, some vendors provide scripts to automate much of the process. For example, 3PAR's Thin Copy Desktop for VMware VDI script, supplied as a customizable script, uses the VMware "perl API" to map the boot image from a VMFS datastore to a LUN, directs the subsystem to snapshot the source LUN multiple times, exports the new LUNs to ESX, and signals VMware to re-scan and re-signature the resultant LUN(s), which are made visible within the VCenter inventory. While not all configuration tasks have been scripted, much of the drudgery of traditional storage provisioning has been eliminated.

VMware VDI is a proven and successful product for deploying multiple desktops in an enterprise, and snapshots of VDI boot image data can significantly enhance VDI.

In fact, the case for boot image snapshots is compelling in a vast majority of situations. To reiterate the benefits:

- Lengthy copying processes can be performed in minutes rather than hours or days
- Boot times for multiple desktops can be significantly reduced
- Substantial storage space savings of boot image data can be realized--often 75% or more.

Examination of the implementation differences between subsystems is critical because vendors do not support snapshots in the same way. Decision-making factors should include automated desktop provisioning, space-efficient snapshots, efficient snapshot caching, and high availability.

© **InfoStor January 2009**

**About the author**

**Ray Lucchesi** is president of *Silverton Consulting*, a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.

<mailto:info@silvertonconsulting.com>

<http://www.silvertonconsulting.com>