



## Vista's Volume Shadow Copies

### What is a Shadow Copy?

As we've said before, Windows Vista, Windows Server 2008, Windows Server 2003 and Windows XP all have the ability to make Shadow Copies of NTFS Volumes. System Restore and Backup solutions provided by Microsoft and other software vendors make use of these Shadow Copies because it is possible to make a reasonably consistent image of an NTFS volume.

Volume Shadow Copies use the well-known copy-on-write technology. As Microsoft has implemented it, this technology operates at the block level of the storage stack. A baseline copy of the volume is generally not taken, only a history of the blocks which are changed is collected and becomes the shadow copy. The changed blocks can be stored anywhere, but the "diff-area" containing them will only be meaningful if the volume's current state (including the data) is available, since the unchanged blocks are needed to complete an accurate picture of the volume at some prior time.

### Volume Shadow Copy Consistency

A consistent copy of a volume means: a copy in which all the data is as the application requires. Note that, only the application can know what this really means. Hence, taking a consistent copy of a storage volume using the copy-on-write approach requires synchronizing the application's view of the data with the shadow copy. Microsoft has done this by implementing a Volumes Shadow Service (VSS) that coordinates the shadow copy synchronization. Whenever a shadow copy is declared, VSS asks all the running applications to flush in-process data. The fact of cached data in the many applications (processes) that must be synchronized noticeably complicates this task. (See the Appendix, which recaps <http://technet.microsoft.com/en-us/library/cc785914.aspx>). In the form that Microsoft delivers Vista, shadow copies are created whenever a new application or system component is installed. These shadow copies are "Restore Points." They are no different from other shadow copies, except that the method of taking the shadow copy is fixed as follows. Restore Points always invoke all the VSS writers in order to achieve the highest level of consistency. Shadow copies without this guarantee that all VSS writers completed successfully may have inconsistent files or folders in them. Restore points have other properties, which are insured during the restoration of the shadow, such as not restoring user data files when restoring Vista to a previous configuration. Nevertheless, the shadow does contain this data and it can be restored by some other means e.g. the Previous Versions tab in the properties dialog or by third party applications like **TimeTraveler™**.

### Speed versus Consistency

The completion of all the applications' VSS writers with cached or otherwise inconsistent data may take some time. By invoking the "No writers" form of shadow copy creation the shadow copy can be created much more quickly. If we do not expect to use this copy as a restore point or system backup, but rather as a point in time from which we can retrieve an applications data files (e.g. the data in our Microsoft Office files), then

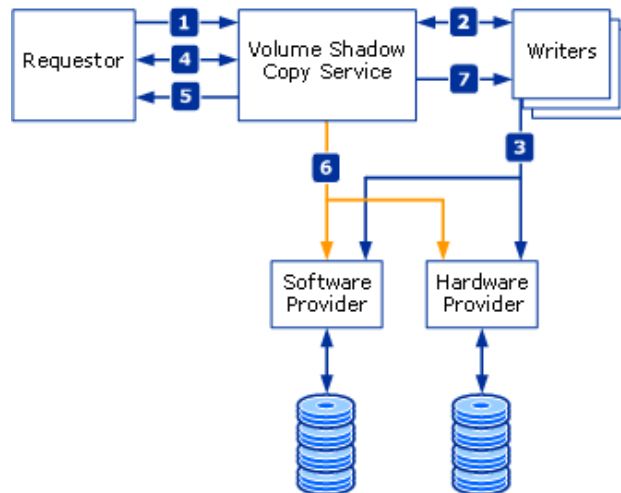


waiting for the writers is probably unnecessary. Microsoft Office applications like Word 2007 try very hard to preserve consistent versions of their files. If they make a mistake, then a warning and an opportunity to restore the file(s) is presented when the Office application next opens the file. This behavior was present before VSS became available, but many users still rely upon it to insure against corruption of their documents. VSS extends this protection against corruption to mistakes we make in editing our files. Now we can “undo” those mistakes by traveling back to a previous time.

### TimeTraveler™

Time Traveler™ takes advantage of VSS to quickly create shadow copies without requiring all application writers to flush their cache (“No Writers”). These relatively fast shadow copies are possible either manually or automatically. Either way, they are the perfect protection against “user errors.” Fast and efficient, they are easy to use and manage. Vista still takes restore points whenever the system or an application is changed or installed, and these shadow copies are also available to us if we need them.

### Appendix: Volume Shadow Copy Service Architecture



1. The requestor asks the Volume Shadow Copy Service to enumerate the writers, gather the writer metadata, and prepare for shadow copy creation.
2. The writer creates an XML description of the backup components to the Volume Shadow Copy Service, and defines the restore method. The Volume Shadow Copy Service notifies the application-specific writer to prepare its data for making a shadow copy.



3. The writer prepares the data in whatever way is appropriate, such as completing all open transactions, rolling transaction logs, and flushing caches. When the data is prepared for shadow copy creation, the writer notifies the Volume Shadow Copy Service.
4. The Volume Shadow Copy Service initiates the “commit” shadow copy phase.
5. The Volume Shadow Copy Service tells the writers to quiesce their data and temporarily freeze requestor (application) I/O write requests (I/O read requests are still possible) for the several seconds required to create the shadow copy of the volume or volumes. The application freeze is not allowed to take longer than 60 seconds. The Volume Shadow Copy Service flushes the file system buffer and then freezes the file system, which ensures that file system metadata is written and that the data is written in a consistent order.
6. The Volume Shadow Copy Service tells the provider to create the shadow copy (a maximum of 10 seconds).
7. The Volume Shadow Copy Service thaws the file system. After the shadow copy is created, the Volume Shadow Copy Service releases the writers from their temporary inactive phase and all queued write I/Os are completed.
8. The Volume Shadow Copy Service queries the writers to confirm that write I/Os were successfully held during shadow copy creation.
9. If the writes were not successfully held (meaning that the shadow copy data is potentially inconsistent), the shadow copy is deleted and the requestor is notified.
10. The requestor can retry the process (go back to step 1) or notify the administrator to retry at a later time.
11. If the copy is successful, the Volume Shadow Copy Service gives the location information for the shadow copy back to the requestor.