

How Volume Shadow Copy Service Works

In this section

- [Methods for Creating Shadow Copies](#)
- [Volume Shadow Copy Service Architecture](#)
- [Components Used for Creating Shadow Copies](#)
- [How Shadow Copies Are Created](#)
- [Related Information](#)

The Volume Shadow Copy Service provides the backup infrastructure for the Microsoft Windows XP and Microsoft Windows Server 2003 operating systems, as well as a mechanism for creating consistent point-in-time copies of data known as shadow copies.

The Volume Shadow Copy Service can produce consistent shadow copies by coordinating with business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware. Several features in the Windows Server 2003 operating systems use the Volume Shadow Copy Service, including Shadow Copies for Shared Folders and Backup.

[↑Top of page](#)

Methods for Creating Shadow Copies

There are two methods for creating shadow copies: making either a complete copy (a full copy or clone) or copying only the changes to the volume (a differential copy or copy-on-write). Each method results in two data images — the original volume and the shadow copy volume. The functional difference between the two is that the original volume maintains full read/write capabilities, whereas the shadow copy volume is read-only. This read-only status ensures that the shadow copy volume remains a point-in-time copy until its status is changed by the administrator for a specific purpose.

Clone (Full Copy/Split Mirror):

A clone is a full copy of the original data on a volume. You can create a clone through either software or hardware mirroring. Clones remain synchronized until the mirror connection is broken for the shadow copy. From this point forward, the source data and the shadow copy volume are independent. The original volume continues to take application changes, while the shadow copy volume remains an exact read-only copy of the original data at the time of the break.

Hardware vendors offer different hardware-based implementations (sometimes called split mirrors, snapshot mirrors, or clones) for creating identical images of volumes that can be used for online backup, application development, and testing.

Copy-on-Write (Differential Copy)

The copy-on-write method creates shadow copies that are differential rather than full copies of the original data. Like the clone method of creating shadow copies, the copy-on-write method can produce shadow copies using either software or hardware solutions. This method makes a copy of the original data before it is overwritten with new changes, as shown in the following table. When a change to the original volume occurs, but before it is written to disk, the block about to be modified is read and then written to a “differences area”, which preserves a copy of the data block before it is overwritten with the change. Using the blocks in the differences area and unchanged blocks in the

original volume, a shadow copy can be logically constructed that represents the shadow copy at the point in time in which it was created.

The Copy-on-Write Method of Creating Shadow Copies

Time	Source Data	Contents	Shadow Copy	Contents
T0	Original data	1 2 3 4 5	No copy	—
T1	Original data overwritten	1 2 3 4 5	Differences and index stored on shadow copy	3

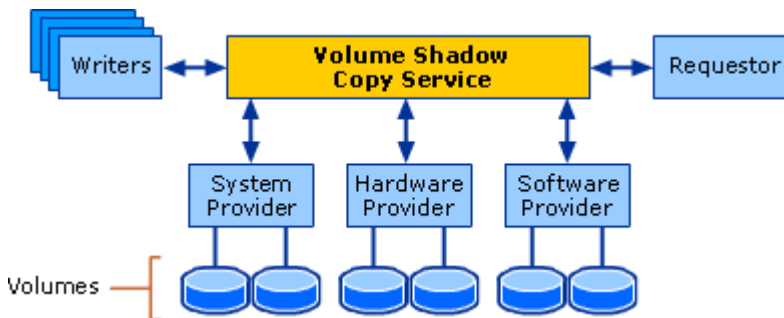
The advantage of the copy-on-write method is that it creates shadow copies very rapidly because it is only writing the changes to disk. The disadvantage is that in order to fully restore the data, the original data must still be available. Without the original data, the shadow copy is incomplete and cannot be used. Another disadvantage is that the performance of copy-on-write implementations can affect the performance of the original volume.

[Top of page](#)

Volume Shadow Copy Service Architecture

The following diagram and table describe how the Volume Shadow Copy Service coordinates with various components to create a shadow copy of a volume.

Volume Shadow Copy Service Architecture Diagram



Volume Shadow Copy Service Components

Component	Description
Volume Shadow Copy Service	A service that coordinates various components to create consistent shadow copies of one or more volumes.
Requestor	An application that requests that a volume shadow copy be taken. A backup application is an example.
Writer	A component of an application that stores persistent information on one or more volumes that participate in shadow copy synchronization. Typically, this is a database application like SQL Server or Exchange Server, or a system service like Active Directory.
Provider	A component that creates and maintains the shadow copies. Examples are the system provider included with the operating system and the hardware providers included with storage

Component	Description
	arrays.
Source volume	The volume that contains the data to be shadow copied.
Storage volume	The volume that holds the shadow copy storage files for the system copy-on-write software provider.

[↑Top of page](#)

Components Used for Creating Shadow Copies

This section outlines the requestors, writers, and providers that are necessary for creating consistent shadow copies. The Volume Shadow Copy Service provides coordination among these components.

Requestors: Initiating Shadow Copy Creation

The Volume Shadow Copy Service is invoked by the requestor, which is typically a backup application that creates shadow copy volumes to back up data while the source volume continues to operate in production. Requestors can also be management applications that manage shadow copy creation and usage, or fast recovery solutions which are specific products that reduce service level agreement (SLA) times for specific applications.

The requestor also communicates with the writers to gather information about what should be backed up and how it should be backed up.

Writers: Preventing Data Inconsistencies

Writers are software that is included in applications and services that help provide consistent shadow copies. Writers serve two main purposes:

- When applications and services are running, the writer responds to signals provided by the Volume Shadow Copy Service interface to allow applications to prepare and quiesce their data stores for shadow copy creation and to ensure that no writes occur on the volume while the shadow copy is being created. (During quiescence, applications make data on the disk consistent. For example, an application might flush its buffers to disk or write out in-memory data to disk.)
- The writer also provides information about the application name, icons, files to include or exclude, and a strategy to restore the files.

If an application or service is not running, and the writer cannot respond to the requesting backup application, it is assumed that all data on the volume is consistent, the databases are closed, and no additional effort is required to perform the backup.

A writer is associated with one or more components. A component is a group of files that must be backed up as a unit. For example, a database and set of log files. For a backup to be successful, all files associated with the component must be backed up. Writers also provide information about restoring the data on a component-by-components basis.

If an application has no writer, the shadow copy will still occur and all of the data, in whatever form it is in at the time of the copy, will be included in the shadow copy. This means that there might be inconsistent data that is now contained in the shadow copy. This data inconsistency is caused by incomplete writes, data buffered in the application that is not written, or open files that are in the middle of a write operation. Even though the file system flushes all buffers prior to creating a shadow copy, the data on the disk can only be guaranteed to be crash-consistent if the application has completed all transactions and has written all of the data to the disk. (Data on disk is "crash-consistent" if it is the same as it would be after a system failure or power outage.)

Under this design, the responsibility for data consistency has been shifted from the requestor application to the production application. The advantage of this approach is that application developers — those most knowledgeable about their applications — can ensure, through development of their own writers, the maximum effectiveness of the shadow copy creation process.

Applications that are not shadow copy-enabled

If the computer has no applications that are enabled for the Volume Shadow Copy Service, the data in a shadow copy is considered to be in a "crash consistent" state. All files that were open will still exist, but are not guaranteed to be free of incomplete I/O operations or data corruption.

While the crash-consistent state does not fully deal with all the issues associated with defining a stable backup set, it has several advantages over the backup set that conventional backup operations would have to use.

- For example, a shadow copy of a volume, even in crash-consistent state, still contains all files. A backup set created without a shadow copy would not contain all files open exclusively at the time of the backup. Files held open at the time of the backup operation are excluded from the backup.

The shadow copy of the volume is created at a single point in time and is synchronized across the whole volume set. In order to avoid inconsistencies, shadow copies are not taken file by file.

Providers: Creating Shadow Copies

In the context of shadow copy creation, a provider is a component that serves as an interface to the point-in-time imaging capabilities — either on the storage array (hardware-based) or in the operating system (software based). Providers manage running volumes and create shadow copies of them on demand. In response to a request from a requestor, the Volume Shadow Copy Service signals applications that a shadow copy is about to be created, then signals the provider to create and maintain that shadow copy until it is no longer needed. Hardware providers are implemented as a user-mode component which communicates with the hardware that will expose the shadow copy data. Windows Server 2003 includes a system software provider with shadow copy functionality. Alternatively, other hardware and software vendors can develop their own hardware or software providers to provide point-in-time imaging functionality. Windows Server 2003 supports multiple hardware and software providers that can be used in combination to solve many different IT operational scenarios.

The Volume Shadow Copy Service uses the following hierarchy to select the provider to use during shadow copy creation:

1. Hardware provider
2. Software provider
3. System software provider.

However, if a specific IT operational problem requires it, the requestor can override this hierarchy.

Hardware-based Providers

Hardware-based shadow copy providers act as an interface between the Volume Shadow Copy Service and the hardware level by working in conjunction with a hardware storage adapter or controller. The work of creating the shadow copy is performed by a host adapter, storage appliance, or RAID controller outside of the operating system.

Hardware providers always take the shadow copy of an entire LUN, but the Volume Shadow Copy Service will only expose the shadow copy of the volume or volumes that were requested.

While a hardware-based shadow copy provider makes use of the Volume Shadow Copy Service functionality that defines the point in time, allows data synchronization, manages the shadow copy, and provides a common interface with backup applications, the Volume Shadow Copy Service does not specify the underlying mechanism by which the hardware-based provider produces and maintains shadow copies.

Software-based Providers

Software-based shadow copy providers typically intercept and process I/O requests in a software layer between the file system and the volume manager software.

These providers are implemented as a user-mode DLL component and at least one kernel-mode device driver, typically a storage filter driver. The work of creating these shadow copies is done in software.

A software-based shadow copy provider must maintain a "point-in-time" view of a volume by having access to a data set that can be used to recreate volume status prior to the shadow copy. An example of this is the copy-on-write technique of the system provider. However, the Volume Shadow Copy Service places no restrictions on what technique software-based providers use to create and maintain shadow copies.

A software provider will be applicable to a wider range of storage platforms than a hardware-based provider and should be able to work with basic disks or logical volumes equally well.

This implementation sacrifices the performance that might be available by implementing shadow copies in hardware and does not make use of any vendor-specific features.

System Provider

One shadow copy provider, the system provider, is supplied as a default part of the Windows Server 2003 operating system. While a default provider is supplied as part of Windows, other vendors are free to supply their own implementations that are optimized for their own storage hardware and software applications.

To maintain the "point in time" view of a volume contained in a shadow copy, the system provider uses a copy-on-write technique. Copies of the blocks on volume that have been modified since the beginning of the shadow copy creation are stored in a shadow copy storage area.

The system provider can expose the production volume, which can be written to and read from normally. When the shadow copy is needed, it logically applies the differences to data on the production volume to expose the complete shadow copy.

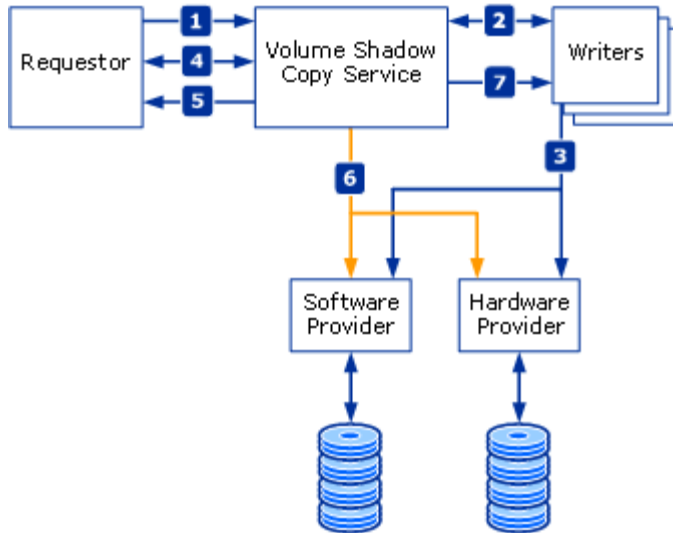
For the system provider, the shadow copy storage area must be on an NTFS volume. The volume to be shadow copied does not need to be an NTFS volume, but at least one volume mounted on the system must be an NTFS volume.

[↕Top of page](#)

How Shadow Copies Are Created

The various roles of the requestor, writer and provider are put into context in this section, which lists the steps that need to be taken to create a shadow copy. Overall coordination of the requestor, writer, and provider is controlled by the Volume Shadow Copy Service, as shown in the following diagram.

Shadow Copy Creation Process



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.