

File Systems

2 out of 3 rated this helpful

NTFS is a recoverable file system that guarantees the consistency of the volume by using standard transaction logging and recovery techniques. In the event of a disk corruption, NTFS runs a recovery procedure that accesses information stored in a transaction log file. The NTFS recovery procedure guarantees that the volume is restored to a consistent state. Transaction logging requires a very small amount of overhead.

NTFS ensures the integrity of all NTFS volumes by automatically performing disk recovery operations the first time a program accesses an NTFS volume after the computer is restarted following a failure.

NTFS also uses a technique called cluster remapping to minimize the effects of a bad sector on an NTFS volume.

Important

If either the master boot record (MBR) or boot sector is corrupted, you might not be able to access data on the volume. For more information about recovery from errors with the MBR or the boot sector, see [Disks Concepts and Troubleshooting](#) in this book.

Recovering Data with NTFS

NTFS views each operation that modifies a file on a volume as a transaction, and manages each one as an integral unit. After it is started, the transaction is either completed or, in the event of a disk problem, rolled back (the NTFS volume is returned to its state before the transaction was initiated).

To ensure that a transaction can be completed or rolled back, NTFS records the suboperations of a transaction in a transaction log file before they are written to the disk. When a complete transaction is recorded in the log file, NTFS performs the suboperations of the transaction on the volume cache. After NTFS updates the cache, it commits the transaction by recording in the log file that the transaction is complete.

After a transaction is committed, NTFS ensures that the entire transaction appears on the volume, even if the disk becomes corrupted. During recovery operations, NTFS redoes each committed transaction found in the log file. Then NTFS locates in the log file the transactions that were not committed at the time of the system failure and undoes each transaction suboperation recorded in the log file. Incomplete modifications to the volume are prohibited.

Important

NTFS uses transaction logging and recovery to guarantee that the volume structure is not corrupted. For this reason, all system files remain accessible after a system failure. However, user data can be lost because of a system failure or a bad sector.

[Top Of Page](#)

Caching and Data Recovery

The cache is the area of random access memory (RAM) that contains the most recently used data. When you

write data to disk, the lazy-write technique in Windows 2000 indicates that the data is written when it is still in the cache. There can also be cache memory on the disk controller, as with small computer system interface (SCSI) controllers, or on the disk unit, as with EIDE disks. The following information can help you decide whether to enable the disk or controller cache:

- Write caching improves disk performance, particularly if large amounts of data are being written to the disk.
- Control of the write-back cache is a firmware function provided by the disk manufacturer. See the documentation supplied with the disk or disk controller. You cannot configure the write-back cache from Windows 2000.
- Write caching does not impact the reliability of the file system's own metadata. NTFS instructs the disk device driver to ensure that metadata is written regardless of whether write caching is enabled. Non-metadata is written to the disk normally and can be cached.
- Read caching in the disk has no impact on file system reliability.

[Top Of Page](#)

Cluster Remapping

In the event of a bad-sector error, NTFS automatically implements a recovery technique called cluster remapping. When Windows 2000 detects a bad-sector, NTFS dynamically remaps the cluster containing the bad sector and allocates a new cluster for the data. If the error occurred during a read, NTFS returns a read error to the calling program, and the data is lost. If the error occurs during a write, NTFS writes the data to the new cluster, and no data is lost.

NTFS puts the address of the cluster containing the bad sector in the bad cluster file so the bad sector is not reused.

Important

Cluster remapping is *not* a backup alternative. Once errors are detected, the disk must be monitored closely and replaced if the detect list grows. This type of error is displayed in the System Log of Event Viewer.

FAT uses a form of cluster remapping, but only when the volume is initially formatted. If a bad sector occurs on a FAT volume after it is formatted, data stored within the associated cluster can be permanently lost. NTFS handles cluster remapping dynamically and continuously, ensuring the integrity of your data.

[Top Of Page](#)

Did you find this helpful? ☐ Yes ☐ No