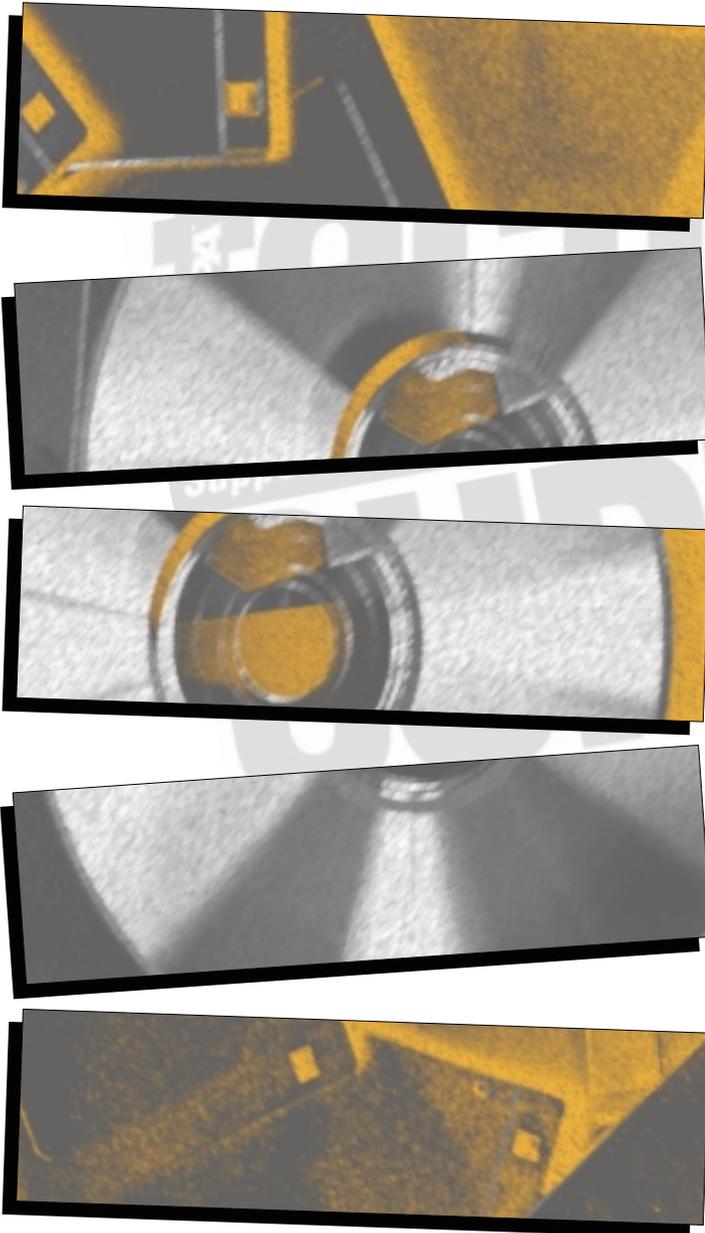


Implementing NetWare File Compression

BY JOHN E. JOHNSTON

This article examines a mechanism that can be used to control and monitor file compression on NetWare 4 and NetWare 5 file servers and to reduce the size of a file using standard compression algorithms, thus eliminating duplicated bits.



ONE of the most time-consuming tasks a network administrator must perform is file server disk management. Someone very wise once said, "If you install a disk drive on a NetWare server, it will be filled." Doesn't it seem that every available byte of disk space has a way of being gobbled up by the end users?

In the old days, network administrators would scan their NetWare disk drives once a week or so, looking for obsolete files that could be deleted or zipped to conserve space. Today, we can rely on file archival systems that automatically move infrequently used files to tape. We can also rely on the NetWare compression mechanism to compress files when they have not been accessed in a specified period of time.

NetWare disk file compression was introduced with NetWare 4 and continues to be used with NetWare 5. File compression reduces the size of a file using standard compression algorithms, which eliminate duplicated bits. Files that have many duplicated bits, such as BMP files, obtain a very high rate of compression. Other files that do not contain a great deal of duplicated bits will not achieve as high of a compression rate.

This article will examine the mechanisms that you can use to control and monitor file compression on NetWare 4 and NetWare 5 file servers.

FILE COMPRESSION OVERVIEW

The file compression feature works the same in NetWare version 4 and NetWare version 5. Before file compression can be used, it must be enabled on a volume level. When NetWare is installed, file compression is enabled by default on all volumes. Once a day, the server will initiate a file compression process. This process will scan all NetWare volumes attached to the server that have file compression enabled. When a file that matches the default file compression eligibility parameters is encountered, the file is compressed. This compression process continues to scan and compress files until it completes scanning all of the volumes or until the "Compression Stop Hour" is reached. You can control when the compression process starts and stops. You

FIGURE 1: MONITORING FILE COMPRESSION

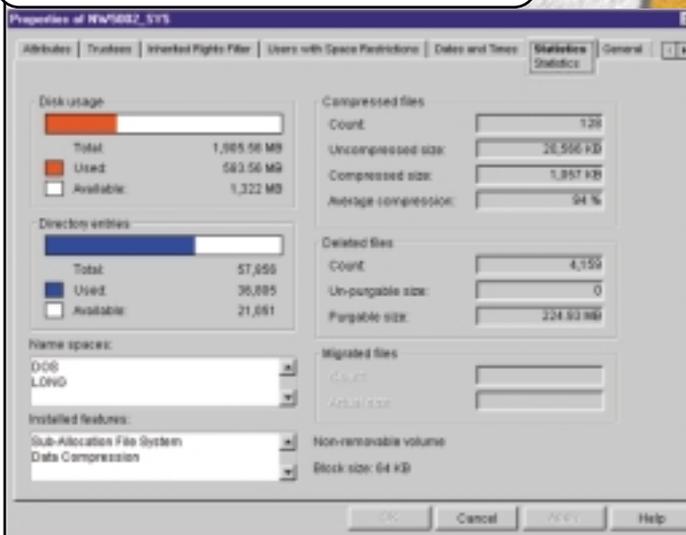
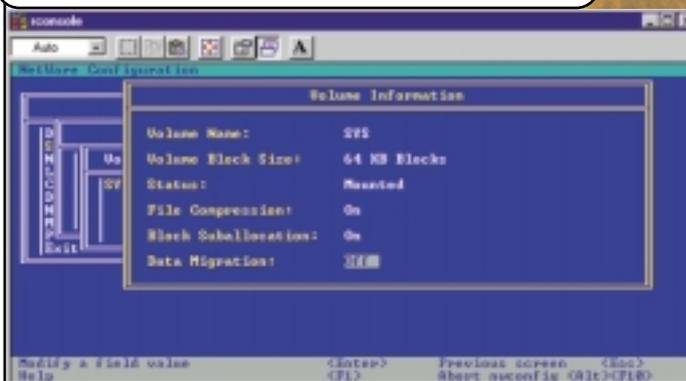


FIGURE 2: ENABLING FILE COMPRESSION ON NETWARE 5



can also modify the parameters used to determine a file's eligibility for compression.

When a user opens a compressed file, NetWare automatically de-compresses the file in memory. When the user closes the file, the file is written back to the file server in uncompressed format.

You can check the number of files compressed on your NetWare volumes using the NWADMIN and ConsoleOne utilities. The panel shown in Figure 1 shows the number of files compressed on a NetWare volume and the average compression rate achieved. Notice that 128 files are compressed on this volume with an average compression rate of 94 percent.

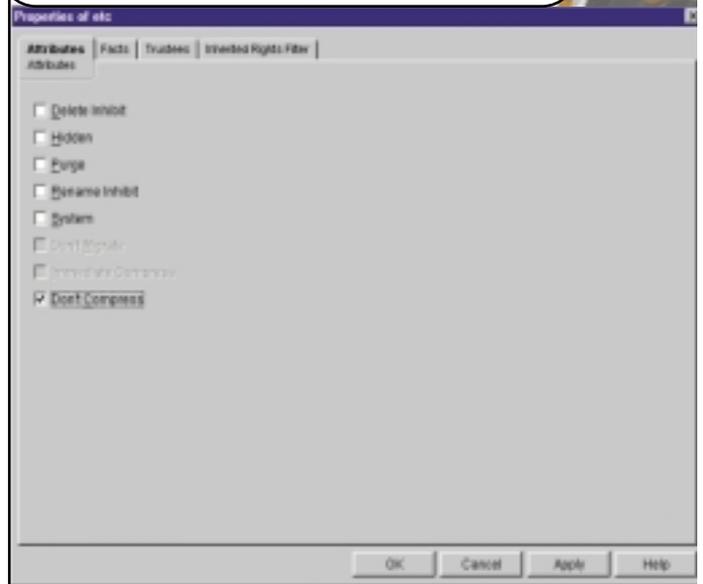
To view your compression statistics:

1. Start NWADMIN or ConsoleOne and then right click on the volume you wish to view.
2. In ConsoleOne select Properties > Statistics.
3. In NWADMIN select Details > Statistics.

ENABLING FILE COMPRESSION

File compression must be enabled on a volume level. Once the volume is enabled for file compression, you can control the file compression on a volume, directory or file level. File compression is automatically enabled on all volumes when NetWare is installed.

FIGURE 3: MARKING A DIRECTORY "DON'T COMPRESS" USING CONSOLEONE



Enabling File Compression on a NetWare 5 Server

To enable (or check the status of) file compression on a NetWare 5 server, perform the following from the system console:

1. Load NWCONFIG.
2. Click on Standard Disk Options > NetWare Volume Options. You will see a panel similar to the one shown in Figure 2. Notice that File Compression is enabled for this volume. If File Compression was not enabled for this volume, it could be enabled on this panel.

Enabling File Compression on a NetWare 4 Server

To enable (or check the status of) file compression on a NetWare 4 file server, perform the following from the system console:

1. Load install.
2. Click on Standard Volume Options.

You will see a panel similar to the one shown in Figure 2. You can enable file compress for the volume if it is not already enabled.

Note: Once file compression is enabled on a volume, it cannot be disabled unless the volume is deleted and then re-created!

GLOBAL FILE COMPRESSION SET PARAMETERS

There are several "SET" parameters that can change how file compression operates on a global level. These set parameters affect how file compression operates on all volumes attached to the file servers that have file compression enabled. These set parameters can be viewed and modified using the SERVMAN utility on NetWare 4 servers and the MONITOR utility on NetWare 5 servers. Let's take a look at the global file compression set parameters and their default settings:

Compression Daily Check Starting Hour

The default setting for this parameter is 0.
The valid range is 0-23.

This parameter tells NetWare when to start the daily file compression scanning process. The scanning process, as mentioned earlier, scans all NetWare volumes, which have file compression enabled, and compresses the files that meet its selection criteria.

The value in this parameter specifies the hour at which the file compression scanning process should start. The valid range for this parameter is 0-23, where 0 equals midnight and 23 equals 11:00 p.m.

Note: The file compression scanning process consumes a great deal of server resources and should be scheduled for off-peak times. The performance of your nightly backup jobs will be affected if file compression scanning is running at the same time as your backup.

Compression Daily Check Stop Hour

The default setting for this parameter is 6.
The valid range is 0-23.

This parameter tells the file compression scanning process when to stop scanning and compressing files.

Minimum Compression Percentage Gain

The default value for this parameter is 20.
The valid range is 0-50.

As stated earlier, some files compress better than others. Bitmap figures may compress to where they only consume 10 percent of their original space. Other files may only be compressed to where they consume 90 or 95 percent of their original space. This parameter allows you to specify the minimum amount of compression that a file must be able to attain in order to be compressed.

Enable File Compression

The default value for this parameter is On.
The valid values are On and Off.

This parameter is used to disable the file compression mechanism. If this parameter is set to Off, no more file compressions will take place until this parameter is set back to On. This parameter is useful on occasions where a special batch job will be running when the file compression scanning routine is run. You can disable the file compression mechanism so it will not interfere with the batch job.

Maximum Concurrent Compressions

The default value for this parameter is 2.
The valid range is 1-8.

This parameter allows you to control how many file compression scanning processes can run concurrently. Increasing this number can make the entire scanning process complete quicker if you have enough hardware resources to support the additional processes.

Note: Each NetWare volume can only run one file compression scanning processes at a time. This parameter will only help you if you have multiple NetWare volumes.

Convert Compressed to Uncompressed Option

The default value for this parameter is 1.
The valid range is 0-2.

This parameter tells NetWare what to do when the server uncompresses a file. When a compressed file is accessed, the server de-compresses the file in the server's memory. When the file is closed, the server must determine whether to leave the file compressed on the disk or to re-write the file in uncompressed format to the disk.

The valid options are:

- 0 - Always leave the compressed version on the disk.
- 1 - If the compressed file is ready only one time within the number of days specified in the "Days Untouched Before Compression" parameter, leave the file compressed on the disk.
- 2 - Leave the file uncompressed on the disk.

Someone very wise once said, "If you install a disk drive on a NetWare server, it will be filled." Doesn't it seem that every available byte of disk space has a way of being gobbled up by the end users?

Note: If the compressed file is modified, it will always be written back to disk uncompressed.

Days Untouched Before Compression

The default value for this parameter is 14.
The valid range is 0-100000.

This parameter specifies the number of days that must pass with a file not being accessed before the file is eligible for compression.

Decompress Percent Disk Space Free To Allow Commit

The default value for this parameter is 10.
The valid range is 0-75.

This parameter specifies the percentage of disk space that must be available on a volume for a compressed file to be written back to the disk uncompressed.

Note: If the compressed file is modified, it will always be written back to disk uncompressed.

Decompress Free Space Warning Interval

The default value for this parameter is 31 Min 18.5 Sec. The valid range is 0 Sec to 29 Days 15 Hours 50 Min 3.8 Sec

If there is not enough disk space to decompress a file (as specified in the "Decompress Percent Disk Space Free To Allow Commit" parameter), a warning message is issued. This parameter tells NetWare how often to issue this warning message.

Note: A value of 0 disables the issuing of the warning message.

Deleted Files Compression Option

The default value for this parameter is 1. The valid range is 0-2. This parameter tells NetWare how to handle the compression for deleted files.

- 0 - Don't compress deleted files
- 1 - Compress deleted files the next day after being deleted
- 2 - Compress deleted files immediately

FILE, DIRECTORY AND VOLUME COMPRESSION OPTIONS

All of the parameters and options we have looked at so far globally control the operation of NetWare file compression. You can achieve a more granular level of control by setting compression options on individual file's directories and volumes. We are provided two options to control compression on files, volumes and directories:

- Don't Compress
- Immediate Compress

Don't Compress

As its name suggests, the Don't Compress option tells NetWare not to compress any of the marked files, directories or volumes. If a volume or a directory are marked "Don't Compress," all files and subdirectories are affected.

Immediate Compress

The Immediate Compress option tells NetWare to compress the marked files and directories as soon as possible, without waiting for the global options, such as "Compression Daily Check Starting Hour" to pass.

Note: This option can severely impact your server. For example, if a volume is marked Immediate Compress, all files are compressed immediately. Files that are accessed frequently will be uncompressed, then recompressed each time they are accessed.

Setting Granular Compression Options

The granular file compression options can be set using the following three NetWare utilities:

- NWADMIN
- ConsoleOne
- FLAG

With NWADMIN and ConsoleOne, you simply navigate to the file, directory or volume, which you wish to modify and view the properties of the object. Then select either the "Don't Compress" or the "Immediate Compress" options.

Figure 3 shows a screen shot from the ConsoleOne utility, which is marking the SYS:ETC directory as "Don't Compress."

CONCLUSION

Managing file server disk space has always been a pain for network administrators. NetWare features such as file compression help make this chore a little less painful. Be very careful if you decide to modify the default SET settings or the granular file compression settings. Think through any changes you are contemplating, then check to make sure the desired results will be achieved. 



NaSPA member John E. Johnston is the owner of a web development company that specializes in e-commerce. John also performs contract work on Novell, NT and UNIX networks. He can be contacted at johnj@fast.net.