



Tech Info Library

Retrospect: Streaming Tape Capacities Explained (6/96)

Article Created: 17 August 1993

Article Reviewed/Updated: 5 June 1996

TOPIC -----

This article answers the question, "Why do my tapes fill up sooner than expected?"

This article provides information about a non-Apple product. Apple Computer, Inc. is not responsible for its content. Please contact the vendor for additional information.

DISCUSSION -----

Dantz Tech Note No. 19

Streaming Tape Capacities Explained

Description

People often call us and ask "Why do my tapes fill up sooner than expected?" This technical note explains why.

Unlike 'block addressable' devices such as hard drives, cartridge drives, and some DC2000 tape drives, 'streaming' tape drives, such as DAT and 8mm drives, format the tape as they write to it, so the capacity varies, sometimes dramatically. Here are the five factors that influence streaming tape drive capacity:

1) Tape Condition

Streaming tape drives automatically verify data by reading each 'frame' of data immediately after writing it. If the data does not verify, the tape drive moves to the next part of the tape and attempts to write the frame again. As long as the drive successfully writes the data before trying too many times, it will not report a media failure and the backup will continue. This 'error rewrite' process, which can be caused by dirty heads on the tape drive, deteriorating media, temperature changes, or a host of other issues, reduces tape capacity invisibly to Retrospect or the user.

2) Media Failure

Retrospect writes to a streaming tape until the tape drive reports that it has reached the end of the tape or that it has encountered a media failure. Then Retrospect logs the media failure, if appropriate, and requests a new tape to continue the backup. So if a tape drive is not achieving expected capacity, the first step should be to look in the log to see if any media failures are reported. (See Dantz Technical Note #12 "Media Failures".)

3) Speed of the Source

The speed of the source volume is the single most important factor in determining streaming tape capacity. Each time the buffer in the drive runs out of data to copy, an "underrun" occurs and the drive goes through a process that uses up raw tape capacity. The more underruns, the greater the reduction in capacity.

A fast SCSI hard disk on the same SCSI bus as the backup drive will tend to result in fewer underruns and therefore allow the tape to hold more user data. Copying files over a network tends to be slower, resulting in more underruns and lower capacities, especially if that network is LocalTalk speed.

The size of files being copied has a dramatic influence on the source speed. The same Macintosh and hard disk may take four times as long to transfer 100 MB of small files as it does to transfer a few large ones also totaling 100 MB. CPU speed matters, too. A faster Macintosh may transfer files at twice the rate of a slower one, even though the hard disk and file sizes are the same.

4) Hardware Compression

In order to boost capacity and speed, many streaming tape drives contain a chip to compress the data coming into the drive before it is written to tape. The amount of extra capacity achieved, if any, is dependent upon how well the incoming data compresses, and this varies widely. Here, from best to worst, is a listing of approximate file compressibility: database, text, spreadsheet, graphics, applications, System files, previously compressed files. Some files will reduce to 25% of their original size while some will not diminish in size at all.

It is impossible to know how much compression will be achieved until the data is actually copied to the tape drive. For display purposes, Retrospect assumes that using hardware compression doubles capacity, which is somewhat optimistic. A 50% increase in capacity is a good and more conservative rule of thumb. However, we must emphasize the broad variation that makes any prediction suspect.

5) Tape Length

With a given tape format, a longer tape will achieve proportionally higher capacity. For example, most DDS DAT drives support both 60 and 90 meter tapes. Inadvertently using a 60 meter tape reduces capacity by approximately one third. DDS-2 DAT drives require specially formulated 120 meter tapes in order to write in their new format. Shorter tapes use the less dense DDS format causing an even more dramatic reduction in capacity.

Capacity Display

Retrospect displays an estimate of tape capacity to aid in managing the backup. Because DAT drives usually do not tell Retrospect what size of tape is inserted in them, and for all the reasons discussed above, Retrospect's assumptions of tape capacity may be inaccurate. Experience may indicate a more accurate number, at which time the capacity display may be changed by clicking the Set Capacity button in Configure:Catalogs. Note: no matter what tape capacity is displayed, Retrospect continues to write to the very end of the tape unless a media failure occurs.

Raw Capacity

We have performed extensive testing of DAT drives, and our experience is summarized in the following table. Note that these are typical ranges, and that it is perfectly normal, and expected, that any particular experience may be outside these ranges in either direction.

DDS DAT Drive Capacity (no compression)

Backup Source • Capacity (60 meter tape) • Capacity (90 meter tape)

600 MB Local SCSI Hard Disk • 900 to 1250 MB • 1350 to 1850 MB
AppleShare server over LocalTalk • 600 to 900 MB • 900 to 1350 MB
Retrospect Remote over LocalTalk • 600 to 1000 MB • 900 to 1500 MB
Retrospect Remote over EtherTalk • 800 to 1100 MB • 1200 to 1600 MB

DDS-2 DAT Drive Capacity (no compression)

Backup Source • Capacity (120 meter tape)

600 MB Local SCSI Hard Disk • 3000 to 3700 MB
AppleShare server over LocalTalk • 2800 to 3300 MB
Retrospect Remote over LocalTalk • 2800 to 3400 MB
Retrospect Remote over EtherTalk • 2900 to 3600 MB

These capacities will be increased by the amount of compression achieved, with the following caveat: if the source data compresses well, more of it is required in any given period of time. Therefore, the effective speed of the drive goes up and more underruns can occur, tending to reduce capacity.

Summary: Nominal vs. Actual

Tape drive manufacturers must provide a nominal capacity for the tape drives they sell. Yet there is no single accurate number for all users because so many factors influence drive capacity. Typically, drive manufacturers will choose the highest number that is reasonably achievable under some conditions: large, compressible files on a fast drive.

While it may not be possible to come up with a single capacity figure for everyone, we do know that all users with a correctly operating and configured

streaming tape drive will achieve some of the most economical storage available.

Copyright 1995-96, Dantz Development Corporation

Keywords: ksts,knts

=====

This information is from the Apple Technical Information Library.

19960606 07:04:28.00

Tech Info Library Article Number: 12985