

ingenieurbüro jörg grössler gmbh
- edv - beratung & system - unterstützung -



Dokumentation

1.84

rheinstraße 24
1000 berlin 41
030/852 7027

AN ALTERNATE TAPE BACKUP SYSTEM
REPLACING STORE AND SYSDUMP

=====

AUTHOR: DIPL.-ING. JÖRG GRÖSSLER

DATE: 11/15/83

An Alternate Tape Backup System Replacing STORE and SYSDUMP

1. Introduction

Since 1975, I am using the HP3000 Computer System. Everybody knows that the only available system to backup disc files on magnetic tapes is STORE and SYSDUMP which both have the same internal structure. Very soon after I had started using the system the idea for an alternate backup system was born which should eliminate most of the disadvantages of STORE and SYSDUMP.

Now, in 1984, these plans have come to a point where we can start implementing a new tape backup program. This paper will present the basic philosophy of the new tape backup system and outline the major differences to the existing STORE and SYSDUMP.

2. How STORE and SYSDUMP work

First of all it has to be said that STORE and SYSDUMP produce tapes with compatible format. The difference between SYSDUMP and STORE tapes is that SYSDUMP uses the first two files of the tape which are empty on ordinary STORE tapes. In the first file, SYSDUMP puts in permanent table information (for example the RIN table) and the system directory. In the second file system programs, the SL and non standard driver are stored.




Fig. 1: general STORE/SYSDUMP format

Figure 1 shows the general STORE/SYSDUMP format for the first volume of a STORE/SYSDUMP tape. The so called STORE/RESTORE label contains an identification which enables RESTORE to identify a STORE tape. The different formats (HP, for example, switched from 1 K to 4 K blocks in the late '70s) are also mentioned here.

The contents of the tape directory is very easy to explain. It just contains the fully qualified file names of each file on the tape (file name, group name, account name). No further information is stored.

For each file on the disc there is one file on the tape. It contains the information sectorwise bound together in 4 K blocks. Each file starts with the MPE file label (one sector of information) possibly followed by user labels and the actual contents of the file. In case the file contains less than 4 K words of information the block stored on the STORE tape is also reduced accordingly.

Figure 1 shows the format of STORE/SYSDUMP tapes for the first volume. Each additional volume contains as the first file an additional STORE/RESTORE label and the volume directory in one file. The volume directory is similar to the tape directory except it only contains those file names which are stored on this tape and the rest of the tapes since STORE cannot evaluate how many files will fit on a specific tape reel.

The format of STORE tapes for labeled tapes is very similar for the first reel except that additional files are produced by the tape label subsystem. The format of each additional volume is different due to the fact that the tape label subsystem handles multiple reels itself. The volume directory for subsequent volumes, therefore, is not stored as the first file of each additional volume. This fact is not very important to understand how the alternate backup system works.

3. The Tape Format of the New Tape Backup System

The new tape backup system uses a very enhanced tape format. Much more information is stored in the tape directory and we will see later on how this effects the handling of this new system.

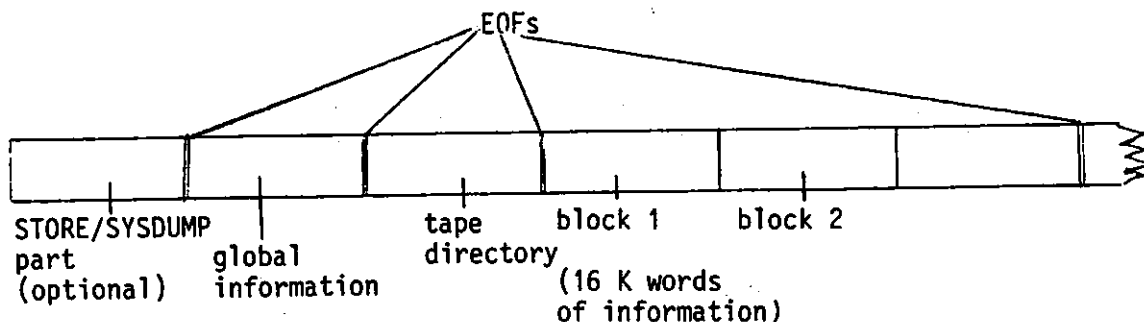


Fig. 2: format of new backup tape

First of all you will see that there can be any STORE/SYSDUMP prior to the actual new backup tape. This is done for two reasons:

- On an ordinary new backup tape the program to restore the tape will always be stored in ordinary STORE format at the beginning of the tape. This ensures compatibility with systems where the new backup system is not installed.
- This format also makes it possible to produce a complete backup on one single set of tapes. In this case the STORE/SYSDUMP part contains an ordinary SYSDUMP including important files from PUB.SYS and the operating system.

There is not much to say about the part which contains the global information. Its structure will be similar to the STORE/RESTORE label.

The tape directory will contain much more information about the files stored on the tape. For each file there will be a 64 words entry which contains most of the information out of the file label.

The actual contents of the files are stored in continuous blocks. Each of these blocks contains 16 K words of information. There is one EOF after every 10th block.




Fig. 3: identification of files

Figure 3 shows how files are identified in the new backup tape. Positioning a file for a RESTORE can be done very easily. Knowing the block number, the system will perform 'BLOCKNUMBER/10' number of SKIPEOF. Then it will skip 'BLOCKNUMBER MOD 10' blocks to reach the actual start point.

4. Results

The introduction of a new format for a tape backup system makes it possible to provide better features than the existing STORE system.

- There is less tape consumption because the blocks of information are larger and small files are packed into the same block and not split into even smaller blocks with one EOF at the end of each file. This measure saves between 10% and 20% of the tape needed for a backup.

- Compressing the data (stripping off zeros and blanks) gives an other 40% to 50% of less tape quantity.
- A considerable improvement of the restore speed can be expected because of the new layout of the tape directory. Since all the information of the file label is stored in the tape directory, a parallel program can go ahead and start building files before the actual location of the file is reached. In the existing system, RESTORE is only able to build files when the file label which is part of the file itself, is read from the tape. This means an average delay of approx. 1 sec. for each file on the tape, regardless of the speed of the tape drive.
- The new tape directory layout also makes it possible to append and replace files to/on an existing tape. At the moment it is necessary to restore the entire tape and store it again when you want to append one single file which is a very time consuming procedure.
- The new backup tape will have its own label identification which means that the standard tape label subsystem will not be used. This makes sense since the tape label subsystem produces a lot of overhead and inhibits the system to use the faster ATTACHIO instead of the file system, and since the tape label subsystem can never be used with SYSDUMP tapes.
- The additional information in the tape directory will also make it possible to decide whether or not a RESTORE can be performed without problems prior to the point where the actual file is on the tape. This means that the user will have the information that there are files on the tape with a creator which is not stored in the system directory at the beginning. So far, this information came out at the very end of the RESTORE and in most cases the entire procedure had to be re-done. With the new system, the user can decide at the beginning whether or not he wants to proceed or may put the creator into the system directory.

All these new features make it worth to rebuild the entire layout of a backup tape. The disadvantage of not being compatible with the existing STORE/RESTORE system is made less severe because on each tape there will be at least one file in the existing STORE format, and this is the program to restore the tape.

I am very optimistic that this new backup system, once it has been introduced, will solve many of the existing problems to backup discs on tapes, and that this is something the users have already waited for a long time.