
Adager's Reblocking of IMAGE datasets

Ken Paul

Adager Corporation

Sun Valley, Idaho 83353-3000 • USA

<http://www.adager.com>

This month I will talk about the concept of reblocking IMAGE datasets. The process to reblock an IMAGE database using Adager is the same whether you are on a Classic or a RISC HP3000. What is different, however, is your objective in reblocking your IMAGE database on these two very different platforms.

CISC or RISC?

An IMAGE dataset's file is always going to have a blocking factor of 1 when viewed from a LISTF or MPE standpoint. This is the MPE blocking factor of the file and not the internal IMAGE blocking factor of the dataset. Within each MPE record of a dataset, also known as an IMAGE block, there can be anywhere from 1 to 255 IMAGE entries. If a dataset has a blocking factor of 10, every MPE record of the dataset's file contains 10 IMAGE entries. When a program does a DBGET it is IMAGE's responsibility to de-block this MPE record and return the appropriate IMAGE entry to the program.

Blocking factor

The difference between a Classic HP3000 and a RISC HP3000 is how the logical IMAGE I/O is done. On the Classic HP3000, the logical IMAGE I/O is done at the IMAGE block level, otherwise known as the MPE record level. On the RISC HP3000, IMAGE I/O is done in units (pages) of 4096-bytes (or 2048 16-bit words). We have to define the reblocking objectives for each kind of system to address the different I/O methods.

How IMAGE does Input/Output (I/O)

**Classic (CISC)
HP3000s**

Classic HP3000 users should attempt to reblock all datasets to have the same MPE record size. The first thing to do is to check a database by doing a LISTF <DBNAME>@,2 and look at the SIZE column. The root file will always have a record size of 128W (which means 128 16-bit words) but all of the dataset's files will probably have several different sizes (ideally, they should have the same size such as 512W).

IMAGE on a Classic HP3000 allocates buffers in memory based on the size of the largest IMAGE block, which is known as the database buffer length. As an extreme example, let's say we have a database that has one dataset with an MPE record size of 2048 words while all the other datasets have an MPE record size of 512 words or less. When this database is opened, IMAGE will allocate buffers in memory that are 2048 words long. When I/O is done on the set which has a block length of 2048 words these buffers will be 100% utilized. When I/O is done on any of the other datasets the buffers will only be 25% utilized and 75% of the buffer space will be wasted. At this point you should do one of the following:

1. You should reblock all the datasets up to 2048 words.
2. You should reblock the one dataset (whose MPE record size is 2048 words) down to 512 words (if possible).
3. You should reblock all the datasets to be some "common" size which should be a multiple of 128 words.

Which of these options you take will depend on your particular system configuration, on your database structure and on the degree of utilization of the datasets with the smaller MPE record size. Some people say that you should reblock everything to 2048 words and others say you should stick with the IMAGE default of 512 words (if the largest media entry will fit in 511 words, allowing 1 bitmap word for a total block length of 512 words). I don't have a magic number but you should definitely reblock all of your datasets to a consistent size under MPE/V for best performance and best utilization of your system memory. You may want to try experimenting with different sizes to see what works best on your system with your databases. Adager allows you to carry out these experiments at very high speeds.

RISC HP3000s

On RISC machines, IMAGE no longer allocates buffers in memory based on the size of the largest dataset block. With mapped file access, buffers are allocated to contain only the pointers to the data. Because of this, efficient disc space utilization of each dataset is the objective when it comes to reblocking IMAGE datasets on a RISC HP3000.

What do I mean by “efficient disc space utilization?” Let me explain. If you do a LISTF <DBNAME>@,2 as described above and look at the SIZE column you may see that some dataset's files have different MPE record sizes. Even if these sizes are different, they all should have a record size which is a multiple of 128. This is the “Athena” format. Every IMAGE dataset's file may have the same MPE record size (e.g. 512W) but, internally, every dataset may have a different IMAGE entry length. Because of this, the goal is to fit as many IMAGE entries into a block without having too much space left over or “wasted” to get us to the next multiple of 128 words.

Let's look at an example. Let's say that we have a dataset which has a media entry length of 29 words (this is the data plus all of the IMAGE pointers). If this dataset had a blocking factor of 17 it would have a buffer length of 495 and an MPE record size of 512 words because 512 is the next multiple of 128 after 495. Now let me explain where I got some of these numbers.

Each IMAGE block contains a bit map followed by each of the entries. A bit map is one 16-bit word for every 16 entries and, since we have a blocking factor of 17, our bit map needs to be 2 words long; so the size of our block is:

$$2 + (17 * 29) = 495 \text{ words}$$

With each MPE record only using 495 of its 512 words, we are wasting 17 words for every 17 entries in the dataset. Now if we were to reblock this set to 640-word MPE records, the blocking factor would change to 22 and the IMAGE block would also equal 640 words:

$$2 + (22 * 29) = 640 \text{ words}$$

Instead of wasting 17 words of disc space for every 17 entries we are now wasting no disc space at all. This may not seem to be a big deal but if our dataset contains over a million entries we are talking about saving over a million 16-bit words of disc space (2 megabytes). This is also an example to show that certain MPE record sizes work better with certain IMAGE media entry lengths and for larger datasets it helps to find the most efficient size so that the least amount of disc space is wasted.

One thing to look for in all of your databases is whether more IMAGE entries can fit in the MPE record that is already allocated. I have seen several examples of databases where a blocking factor was 10 and the buffer length was 521 words and the MPE record size was 640 words. This MPE record could contain 2 more IMAGE entries and have a buffer length of 625 words. I believe that this situation came about when HP converted databases from IMAGE/3000 to TurboIMAGE with DBCONVERT. The program kept the same blocking factor for each set and because master datasets had some words added to

their media entry length the total buffer length went just over 512 words so IMAGE built a record of 640 words but wasted a lot of space.

How Adager helps you in this process

I don't expect you to get out your calculators and start figuring out what the best blocking factor is for each of your datasets. You can use Adager's REBLOCK function to test all the different multiples of 128 against your database. The output from REBLOCK shows you your dataset's current blocking factor and block length and the proposed dataset blocking factor and block length. More importantly, the REBLOCK output shows you the percentage of the block that is being utilized at both the current and the proposed levels so these numbers can easily show if a REBLOCK of a given dataset would help disc space utilization. If you wish to minimize the disc space utilization, Adager offers you the option of specifying a target blocksize and then reblocking all of the eligible datasets so that the resulting blocking factor for each dataset minimizes the dataset's file disc space (while keeping its blocksize less than or equal to the target blocksize specified). This option is available on Adager versions newer than 940101.

Summary

The goals of reblocking are different depending on which HP3000 platform you are running. If you are running on a Classic HP3000 under MPE/V you should reblock all of your datasets to the same size MPE record length such as 512W, 1024W or 2048W. If you are running on a RISC HP3000 under MPE XL or MPE/iX you should reblock datasets individually to achieve the best disc space utilization. You should also concentrate on the larger datasets as they will give back more disc space. Reblocking your database on a RISC HP3000 will not have an impact on your database performance because you are not affecting how many entries are being brought into memory with each I/O as is the case on Classic HP3000s.