Storage Management - seven years and seven lessons with DMF

Robert C Bell, CSIRO Australia.
Deputy Manager, Bureau of Meteorology / CSIRO
High Performance Computing and Communications Centre
(HPCCC)
Robert.Bell@cmis.csiro.au

ABSTRACT: *CSIRO has been using Cray's Data Migration Facility (DMF) for storage management for over seven years. This talk will focus on seven lessons learned in the seven years. These are: Good storage systems are vital to productivity of users and administrators. Good storage systems cope with generational change transparently to the users. Good storage systems provide long-term safe keeping for users' data. Good storage systems need good hardware, software and support staff. Good storage systems have to be more than an archive. Good storage systems have to cope with many files as well as large files. Good storage systems provide good access paths.*

Introduction

CSIRO, the Commonwealth Scientific and Research Organisation is Australia's premier scientific research organisation, with about 7000 staff. CSIRO has been using Cray's Data Migration Facility (DMF) as a Hierarchical Storage Manager (HSM) for over seven years for its central high performance computing systems.

CSIRO has a long history in scientific computing – this year marks the 50th Anniversary of CSIRO's first computer, the locally designed and built CSIRAC, now the world's oldest surviving original electronic digital computer.

From the mid-1960s to 1990, Csironet (as it was finally called) provided central computing services to CSIRO. In the early 1970s, the CDC 3600 had a document region, which provided an HSM with manually mounted magnetic tapes. During the 1980s, Csironet developed a Terabit File Store, based on a Calcomp/Braegen Automated Tape Library, and hosted on a Fujitsu system. This provided automation, but functioned as an archive rather than an HSM system.

In 1990, CSIRO established the Joint Supercomputer Facility with a private company, based on a Cray Y-MP2/216. Little thought was given to storage needs when the facility was established, with only a small amount of disc (about 20 Gbyte) and 6250 bpi tape drives. It became clear within a year that better facilities were needed: manually mounted tapes under user control were not successful. A move was made to run Cray's Data Migration Facility HSM , as a means to provide overflow facilities for data when the discs became full. An HSM allows data to be automatically copied from disc to cheaper and higher capacity media and then removed from disc, and provides automatic retrieval of the data when needed. DMF was introduced into production on the home file systems on 14th November 1991, using StorageTek 4480 drives and manually mounted 3480 cartridge tapes. In 1992, a Cray Y-MP/4E was acquired and installed at the University of Melbourne to become the core of the CSIRO Supercomputing Facility. In June 1993, a StorageTek 4400 Automated Cartridge System (ACS) was brought into production, thus bringing full automation. In 1996 StorageTek Timberline drives replaced the 4480 drives.

In 1997, CSIRO and the Australian Bureau of Meteorology formed the High Performance Computing and Communications Centre (HPCCC). This featured an NEC

SX-4, but for transition and data management a Cray J90se and a StorageTek 9310 ACS with Timberline and Redwood drives were acquired by CSIRO and installed at the Bureau's site.

This talk will focus on seven lessons about storage management learned in the seven years running DMF.

Lesson 1. Good storage systems are vital to productivity of users and administrators.

It turned out that users of High Performance Computing (HPC) systems need more than just computing power. For many of our users, large scale storage is needed as well, to store the output of simulations. This is particularly needed for Atmospheric Sciences. However, just providing lots of capacity (which can be done quite cheaply with tape) is not sufficient.

Users just want to have fun, or at least want to store and access data. Users should not have to be overly concerned about the physical location of data, and should not have to mess with magnetic tapes with all their variants and multiple data formats. One of our users confided to me two years ago that he had wasted one month out of three messing about with tapes on his local system. I assert too that users should not be restricted in their science by the capacity of the storage system, or by the policies of their systems administrators. For maximum productivity, users should be able to access all their files through their usual file system.

I assert that storage systems without a management plan are a disaster. A management plan covers several items – capacity to meet needs, security, and resiliency. While users are fallible, we need to do backups. RAID systems are not enough – they guard only against some hardware errors, and not at all against human or software (or firmware errors). HSM systems are not a replacement for backups. But, HSM systems enable backups to be feasible for today's huge disc systems, whose capacities are growing much more rapidly than the speeds of backup devices. With HSM, most data is moved to tertiary storage, and backups have to deal with only small amounts of data and metadata.

Capacity to grow is important. I heard John Mashey of SGI say in a presentation at the end of 1997 that "discs are binary devices – either new or full". Inevitably, un-managed systems will fill. Quotas can help, but a great deal of time can be wasted in coercing or embarrassing users into removing files, or in coercing administrators into increasing quotas.

Administrators of storage systems should be freed from the tyranny of users filling up file systems – highly unproductive. They should be freed from chasing and pressuring big users (or vice versa). They should be able to provide backup (and retrieval) for large file systems. Administrators are users too – they need to store logs, reports and records, and so need good storage facilities too.

Good storage systems allow us to cope with growth. CSIRO's usage has been growing at about 100% per annum for the last two years, without any major crises.

DMF allowed CSIRO to improve the productivity of users and administrators.

Lesson 2. Good storage systems cope with generational change transparently to the users.

We are aware that storage media (punched cards, paper tape, 8" floppy discs, round tapes, etc) have a limited life. We know too that formats and software change – who can read CDC Cyber internal format data any more, or CP/M, or Microsoft Word 1. One scientific site I know expects users to care for data on their own commodity tapes, which they keep in their offices. And yet the same mistakes are being made, with people archiving data onto CD-ROMs, "because they last for 50 years".

I believe that good storage systems cope with the changes of media and storage systems transparently to the users. DMF has allowed CSIRO to cope with 7 'generations' of data storage since 1991. There have been three media formats (18-track, 36-track and SD-3), one major software change (to Advanced Tape Media Specific Process), and two changes of site (including machine upgrades). In the most recent move in 1997, two copies of data in separate pools enabled us to move sites, and move from a Y-MP to a J90se in a few hours, taking about 1.5 Tbyte of user data and the processing with us. Contrast this with rusting bytes in people's offices.

Part of the coping with generational change is coping with growth in storage capacity. On desktop systems these days, the usual way of coping with lack of storage space is to buy a new disc or new PC. It is more productive of users time to keep everything than to spend the time classifying and sifting through files.

To cope with increased demands, new storage technology needs to be introduced. This means copying from old media (e.g. 18-track) onto newer media (36-track, SD-3, 9840, etc). Good storage systems do this transparently to users, though with some pain. Clearly during transitions, there needs to be sufficient tape drives to not make a major impact on users. At these times, our analyst activated a task which checked for periods of low tape activity, and slotted copying activity into those periods.

Lesson 3. Good storage systems provide long-term safe keeping for users' data.

Long-term is important to not just archivists, but also to our scientists, though their time-scale is somewhat shorter. Some data needs to be kept for seven years for contract purposes. Other data is the output from climate simulations which may have taken 18 months computation; in theory, re-creatable, but in practice difficult. The data is needed for comparisons with recent experiments. I actually have data which I have carried over from the mid-1970s from earlier CSIRO systems. Data needs curation - like a museum collection – preservation and maintenance is needed. Unfortunately, individual users are rarely interested in this. To misquote Lord Tennyson, "flops may come and flops may go, but bytes go on for ever."

Of course, once an HSM is started, it is very difficult to withdraw the service, as users become attached to it! Thus it is important to choose a suitable HSM from the beginning.

Part of the long-term safe keeping is to protect users from themselves. Good storage systems provide backup, to guard against user and system error. With today's disc capacities outgrowing tape speeds, I believe it is essential to use HSM to trickle data to tape, so that backup only deals with small files and metadata. Our site does nightly backups of a 120 Gbyte file system, which is completed in about an hour. We also allow five weeks between a soft and hard delete of users' files, so that recovery is possible within this period. We also (as do several other sites) ensure that once a month, all used tapes in one pool and a complete set of backup tapes are taken off-site, to allow recovery in case of a major disaster.

Lesson 4. Good storage systems need good hardware, software and support staff.

HSM will not be successful without appropriate hardware. We waited for cartridge tape drives before we started DMF. We targeted automation at the start. We identified StorageTek tape equipment as the preferred solution for attachment to Cray systems at overseas sites through the Cray User Group. We knew we needed to provide high speed and fast access tape drives to make HSM workable.

We talked to other sites about their storage management systems, and concluded that for us and our budget, DMF was the best solution.

For the initial Y-MP set-up, Cray Research brought two analysts to Australia from the US. From 1992, CSIRO had a Facilities Management agreement with Cray Research to provide Help Desk and Systems Administration support. The Help Desk person dealt with the day-to-day problems, and the analyst had to deal with the entire systems administration for the site. This was beneficial to all parties – the analyst is now one of the key DMF support people outside the US. Here are some quotes from users received by the Help Desk and CSIRO Supercomputing Support Group. These show the importance of good support and the right systems.

"I continue to be very impressed with your amazingly quick and informative responses! I reckon you're the most helpful of all the people I have communicated with in this job!"

"Thanks to you and your team for all your assistance ... this is the best experience with a central computer facility that I have had in my 20-year history."

"A short note to express my appreciation of the excellent performance of the Cray Y-MP "cherax" computer and data migration system and the completely professional help and expertise from all staff. In the four years I have been using the system I have never lost any data or code or even a single model run from the system. The system has run with absolute reliability. Whenever I have needed assistance in solving coding or computer system related problems I have received prompt effective help. Further, when I have accidentally deleted runs they have without fail been promptly reinstated from system backups. And when I have suffered from annoying coding problems, I have received calm, tolerant help. Simply, it's the best quality computer system I have ever used."

Lesson 5. Good storage systems have to be more than an archive.

I define an archive system as a store which requires explicit actions by users to receive and recall data. Usually, no applications are run on archive systems except the storage management software itself.

Many people have said that we should run an archive system off the supercomputer, because it is cheaper, and that is all users need. Our experience is that archive systems don't work without making life difficult for users on the HPC system. Users are reluctant to copy their files away from where they work; managing files in yet another location adds complexity. (Our large users already deal with about eight different file systems - too many!) Users are even more reluctant to remove files from their primary disc area. Users have been known to leave unwanted files in the primary disc area until they need space themselves, and then remove the old files to be replaced by new files. Archive systems do not provide management of the disc space area where users do their work. Archive systems still have to have an interface to off-line media, and typically use an HSM. So my argument is, why not simplify things and run the HSM on the HPC system?

Archive systems can often provide a higher level of access methods than a file system can. They can often store more metadata, and provide more powerful searching and indexing methods. However, such systems tend to be discipline-specific, and in our mixed application environment, we do the best we can by providing HSM.

In an active system, users do experiments, and many of these are unsuccessful. However, the data is often saved for a while before being discarded. As well, we allow five weeks after a user deletes a files before we hard delete it from DMF. At one stage last year, we had a quarter of the data on our tapes in the soft-delete state, a very high churn rate.

The archive scheme is the old supercomputer centre model – separation of the compute engine from the storage system. Under this model, a user starting work has to transfer many files on to the compute system, and at the end of the job or session had better remember to save all changed or new files, or they will be lost. We do not expect our managers to have to do that on their desktop systems or laptops, so why should we do that to our users when a better solution is available?

Lesson 6. Good storage systems have to cope with many files as well as large files.

In recent years as the High Performance Storage System (HPSS) has come into operation, it has become clear that while it was designed to provide very high performance for very large files, it may not be a suitable system for large numbers of files. Creation rates of files are only around five per second, compared with hundreds for a conventional file system. At the Stuttgart CUG, Alan Powers presented data [1] for NAS which showed that for 90% of their files were less than 1 Mbyte in size, but these accounted for less than 3% of the data. In March 1999 on the Open side of LANL, 90% of files were less than 50 Mbyte in size, and 90% of the data was in files less than 15 Mbyte [2]. On the CSIRO system, 34% of the inodes are currently used for files smaller than 4096 bytes and directories. These account for less than 0.02% of the data!

Good storage systems need to cope with both large files and large numbers of files, preferably in the same file system. Good performance of the file system for both large and small files is needed: the Cray NC1 file system with primary and secondary allocation units in the one file system allows many files to be stored, and provides good performance for large files. Users want to store large and small files together – index files, source, small control datasets along with large data files.

Right from the start of our DMF experience, we were concerned about coping with large numbers of files, and finishing up with a file system full of small files and directories, with no working space for files. We used quotas to set a default limit of 5000 inodes per user (now increased to 10000). Like many other sites, we provided a utility (called tardir) which allowed a user to simply put all the files in a directory into a single tar archive, and removed the files. The archive is likely to get migrated, but the utility left behind small files containing a tar output listing and a long `ls` listing of the directory. These were likely to stay on-line, and meant the users could see what was in the archive file without retrieving it. For experienced users, the utility provides the convenience of dealing simply with a collection of files, and better performance – only one tape access is usually needed and a tar execution to retrieve perhaps hundreds of files, compared with many tape mounts for un-archived files. By means of quotas and providing utilities for users, we have reduced the overheads of the HSM by reducing the tape mounts and the size of the databases – DMF does not carry so often the relatively large overhead for small files.

At the beginning of May 1999, the 120 Gbyte CSIRO home file system on the J90se had 481 000 inodes in use, with 318 000 or about 66% being for migrated files. The average size of the files migrated or eligible for migration (> 4095 bytes) is 11.4 Mbyte.

Lesson 7. Good storage systems provide good access paths.

Users need to be able to retrieve files reasonably quickly. A one-minute attention span is probably typical. For small files, the average retrieval time is around 45 s. This has slowed down since our initial years with DMF, despite faster tape drives, but is still acceptable. We use StorageTek Timberline drives for files less than about 100 Mbyte, and Redwood drives for larger files, and old files larger than about 10 Mbyte, and for second copies of all files (to go off-site). (Two copies are kept of all files, except for dumps of the NEC SX-4 and some workstations, where only one copy is kept.) This strategy allows reasonable retrieval times while keeping our storage to one StorageTek 9310 Library Storage Module.

For HSM, we do not think that drives which take from a minute to over three minutes to access the first byte would provide acceptable performance for our interactive users. We need high performance drives with a low access time. Prior to acquiring our first StorageTek silo in 1993, I did a cost-justification based on saving the time of scientists and getting greater utilisation of our Cray Y-MP. The analysis showed that the investment would be recovered in about 18 months.

A site needs to provide sufficient tape drives to support the load. Last December, our four Timberline drives supported over 2700 tape mounts in a day, with the load coming almost entirely from one user. I have seen estimates of requirements from a site not running DMF, which suggested that the drives would be busy writing for 75% of the time. This allows far too little capacity for recalls and peak demands. Perhaps drives busy for 25% in each week would be reasonable. Spare drive capacity is needed for backups, and as mentioned in Lesson 2 above, for moving to new media. There is a need to merge tapes as they become sparse, and this requires sufficient tape drive capacity, particularly with high capacity tapes.

There is an increasing need to think about access patterns to data as capacities and demands for access grow at a greater rate than tape speeds. An HSM usually optimises the writing of data - it waits until it gets sufficient data, and then writes to tape in a single operation. This mixes users' files, files from different projects, files of different ages, sizes and types. The recall pattern would hardly ever match the writing pattern. Merging may further mix these files up. It may be useful to have the ability to group files according to the directory they are in.

Users can often help themselves by archiving up related files, or by issuing `dmput` commands for related files, thus increasing the likelihood that the files will end up on the same tape.

However retrieval patterns are harder to plan for. A climate model at our site creates about 50 files per model day. Some of these are tarred up to reduce the number of files, some are compressed, and some explicit `dmput` commands are done. Later, who knows in what order they may be retrieved! One user might want rainfall patterns over Australia for a period, another may want global mean temperature, another sea ice distributions. One of our ocean modellers said to me two years ago that the real bottleneck in these areas of science is not in the simulation, but in the analysis phase, and data access is a critical part.

Now users can help the accesses by issuing `dmget` commands in advance of when they need the files, and DMF can then optimise the tape mounts and access patterns on the tapes (though this is non-trivial for serpentine tapes). DMF also can be used to retrieve partial files, and this may suit some applications. Our new systems administrator, when faced with having to retrieve 3000 files and 35 Gbyte for a user to be written to a dedicated tape, devised a script which grouped the files by volume serial number, so that `dmget` commands could be issued sequentially for each tape, and the number of tape mounts was minimised.

One other solution is to consider tiling of data, as is used in image processing to circumvent the worst behaviour of virtual memory systems when data is accessed in various patterns. Indeed, there are many parallels between virtual memory and virtual disc (as provided by an HSM), and I think there should be some more lessons to be learnt from the vast studies done on virtual memory systems and page replacement algorithms.

One more solution is to store multiple copies of data (or at least of selections), organised for different retrieval patterns. So, for a climate model, one copy would be organised as perhaps a file for each field for each year, whereas another copy might be all fields organised into a file for each month. Even this organisation might reduce the tape mounts by factors of twelve for some access patterns.

Conclusion

I believe that good storage management is vital to the productivity of users and systems administrators, particularly for applications which are data-intensive or which produce large amounts of data. Hierarchical Storage Management is vital for providing capacity, easy access for users, and management of free space where users work. Archive systems do not provide the ease of use nor the space management which is required.

Careful management, education of users, and the provision of tools such as `dmget` and utilities which allow grouping of files can help us all to get the best out of storage systems. There is still a lot of work to be done on optimisation of data placement to provide efficient support for a variety of access patterns.

We have learnt a lot from seven years with DMF.

Acknowledgments

My thanks to Peter Edwards of SGI/Cray and Jeroen van den Muyzenberg of CSIRO for reviewing this paper. Special thanks to Peter Edwards and Eva Hatzi of SGI/Cray for their support for CSIRO's DMF operations.

References

[1] Powers , Alan K. 1998. The Rebirth of DMF on IRIX.
40th Cray User Group Conference Proceedings, 15th-19th June 1998, Stuttgart Germany

[2] Lynse, P., Lee, G., Jones, L., and Roschke, M. HPSS at Los Alamos: Experiences and Analysis.
16th IEEE Symposium on Mass Storage Systems / 7th NASA Goddard Conference on Mass Storage Systems and Technologies