# Decommissioning a Large Data Archive: Lessons Learned from Cleaning out the Attic

August 2019

Richard L. Moore[1], David Minor[2], Subhashini Sivagnanam[3], Mahidhar Tatineni[3], William S. Young[4]

## Abstract

This paper describes key elements of the decommissioning of a large tape-based data archive that the San Diego Supercomputer Center (SDSC) operated for its users from the center's inception in 1985 until ~2010. This 25-year period covered many generations of supercomputers and correspondingly many generations of tape and storage technologies, with Moore's-law growth in supercomputing power and associated storage capacity/bandwidth. Over the archive's last decade, data volume grew exponentially with a doubling period of ~16 months to a maximum size of ~10 PB. In ~2010, the National Science Foundation terminated funding for SDSC's tape archive and SDSC proceeded with decommissioning the archive over a ~2-year period. This paper briefly describes the principles and process by which we decommissioned this large archive, key issues that arose during this process, and implications for institutions that operate data archival systems and suggestions for operating archival systems in the FAIR data environment.

## Background

The San Diego Supercomputer Center (SDSC) has operated many generations of supercomputers from its inception in 1985 to the present, primarily funded by the National Science Foundation (NSF), with supercomputing time allocated by a national peer-reviewed proposal process at no cost to users. In parallel SDSC operated a large tape-based data archive for its users from the center's inception in 1985 until ~2012. This 25-year period covered many generations of supercomputers, with Moore's-law growth in supercomputing power, data generation and data storage technology. As an example, over the archive's last decade, data volume grew

[1] University of California San Diego, San Diego Supercomputer Center (retired); r3moore@ucsd.edu.
[2] University of California San Diego Library.
[3] University of California San Diego, San Diego Supercomputer Center.
[4] Genentech Inc., A Member of the Roche Group (formerly with UCSD/SDSC).

exponentially with a doubling period of ~16 months to a maximum size of ~10 PB. In ~2010, the NSF terminated funding for SDSC's tape archive and SDSC decided to decommission the tape archive and either delete or find a new home for all data in the archive. This was a challenging and time-consuming process, with many lessons learned for resource providers that operate large-scale data archives.

Tape-based archival storage was offered to all SDSC users who received time on the supercomputers. There were some direct and indirect guidelines to constrain use of the archives – e.g. inactive users were expected to move/delete their archival data after a reasonable period, 'excessive' archival storage users were identified and requested to curtail storage, and at times dormant files were identified and flagged to users. (Sometimes these approaches backfired – e.g. some users developed scripts to periodically 'touch' their files simply to avoid being flagged as dormant, consuming tape drive resources with no productive purpose.) But at the bottom line, tape-based storage was considered relatively inexpensive (~$500/TB/year in 2006, see Moore *et al.* 2007) compared to the computational cost of generating the data and, in the spirit of providing good service to the users, hard constraints were rarely imposed. As an operator, the exponentially-decreasing costs of storage made it feasible to retain older files at modest cost, and *requiring* users to spend their precious time dispositioning old files was deemed an unnecessary imposition on our user community. The result was that data accumulated in the archive over that 25-year history, with ~10 PBs of data, millions of files, and thousands of users. A significant fraction of these files were 'write-once-read-never' (WORN).

The subtitle of this paper is '*cleaning out the attic*.' As with an attic, our users put their data in the archive for decades knowing it was there if they ever needed it. Meanwhile they could safely forget it. A box in the attic may contain highly valuable heirlooms or useless junk – and only the box's owner may know the difference. And even though we had usernames associated with all the files, it was not trivial to find those users over decades, get them to look through their files and make disposition decisions. For anyone who has ever cleaned out an attic, it's frustrating and tedious and the last thing anyone wants to spend their time doing. And for scientific researchers, reviewing old data files for retention/deletion has never gotten anyone tenure or a Nobel Prize.

This paper briefly describes the principles and process by which we decommissioned this large archive, key issues that arose during this process, and implications for institutions that operate data archival systems.

The authors acknowledge that this paper is written retrospectively, about 7 years after the decommissioning was complete and with many specifics long lost; priorities at the time unfortunately precluded a contemporaneous documentation of the decommissioning process. However it remains valuable to briefly document the key elements and conclusions.

## Principles and Process

In the early 2000's, SDSC maintained two hierarchical storage management systems in parallel across its tape archive – IBM's High-Performance Storage System (HPSS, now a collaborative project described at http://www.hpss-collaboration.org/) and Sun's Storage Archive Manager (SAM-QFS, now by Oracle, see https://docs.oracle.com/cd/E22586_01/html/E22570/glebg.html). In order to reduce costs, in July 2009 we decided to consolidate to one system (SAM-QFS) and issued a notice to all HPSS users that they would need to move their data across to the SAM-QFS system or out of the archive. Most of the data in the archive was managed under HPSS, particularly the older files, and this notice started the process of working with users on dispositioning their data and getting confirmation of successful transfer/disposition. In mid-2010, NSF curtailed all funding for the archive and SDSC decided to decommission the entire system, both HPSS and SAM-QFS. At that point SAM-QFS users and the remaining HPSS users were notified to disposition their data out of the system (see below). The initial expectation was that the migration from HPSS to SAM-QFS would take about a year, but in fact we did not retire HPSS until the end of 2011, about 2.5 years from the time we first notified users. SAM-QFS and the archive system were retired in June 2012, about two years after we notified users that we would decommission the system.

We approached the decommissioning process with the simple guiding principle to not delete any data without the data owner's permission. SDSC had acted as stewards of the users' data and felt that positive confirmation for deletion was the right thing to do – even though it made this process much harder and time-consuming. Alternative approaches – e.g. to send an email to all users and say they had until a specific date to move the data they wanted to retain – would have been simpler and much less time-consuming.

A key issue was what users should do with data they wanted to retain. SDSC provided tools for transferring data (e.g. pre-Globus gridftp), high-speed network connections, and user support personnel to work individually with any users that wanted to move their data. Some users were able to readily transfer data to their local work environment (often inexpensive USB drives with dubious reliability and replication procedures). A few users took advantage of commercial cloud-based storage systems, but cost was often a barrier with for-fee services. Large-scale users faced significant challenges but many of them moved the bulk of their files to archival storage systems at other NSF or DOE supercomputing facilities. After the tape archive was decommissioned, SDSC has offered annual, renewable storage allocations on the multi-petabyte scale disk-based *Data Oasis* storage system (www.sdsc.edu/support/user_guides/oasis.html). However, although this system currently has ~2.5PB capacity, it is much more limited than the tape archive, especially when one considers 8 years of advances in computational scale.

When the decisions were made to decommission first HPSS in mid-2009 and later SAM-QFS and the archive in mid-2010, we posted logon notices and sent emails to thousands of compute/archive users notifying them of the reasons for the decisions, the alternatives for where they could store data to be retained, the tools and support available for them to move/delete their data, and a generous timeline of about a year for dispositioning their data. We asked that people contact us and provide confirmation that the data they wanted to retain had been moved and that we could delete any of their remaining files. In practice, we had slightly different procedures and timelines for different categories of users – e.g. traditional users that had received computer time allocations, users that were part of our more recent 'Data Central' project for storing data other than computational simulations, and SDSC-internal users (e.g. backups and system data). But the process of notification and positive confirmation of disposition applied to all categories of users.

A significant fraction (~10-20%) of that initial round of emails bounced with undeliverable addresses, and this initiated the first round of time-consuming challenges mandated by our principle of positive confirmation for deletion. We had to individually track down hundreds of users, some from 20 years previous (before email was common), using web searches and other means to get current contact information. Many users had changed institutions, some had changed names, and some had retired or passed away. It took many person-months for our user support team just to try to reach all users.

Some users responded in a timely way to our email, and followed our requests on a pro-active timeline. Large-scale users were mostly responsive, reviewing their data holdings for potential deletions, and initiating large-scale transfers of important data to other centers. But a not-insignificant fraction of users provided little or no response to our emails and phone calls, and this posed another time-consuming challenge to our user support team. Had we not established the principle of requiring positive confirmation, we could have just adhered to the (generous) one-year deadline and then deleted data. But instead this process dragged out for 2-2.5 years – sending multiple emails, calling people, trying to prompt response and action on the part of users. And as we approached various deadlines, we would sometimes finally hear from procrastinating users – often with last-minute problems related to retention/review of their files that necessitated further delays. This was a frustrating experience for both users and SDSC personnel.

It was 2.5 years from the time the announcement to decommission HPSS was made until it was retired, and two years from the time the SAM-QFS/system decommission decision was made until the equipment was finally turned off. We did not keep separate labor accounting for the decommissioning, but it probably required 1-2 FTE-years of user support personnel interfacing to users, plus the labor and hardware maintenance costs required to maintain the archive hardware/software during this three-year period. Despite heroic efforts, there were a few users where we were unable to get positive confirmation for deletion. However, it was satisfying that there have been no users that contacted us after the decommissioning looking for their data.

## Conclusions and Recommendations

While this paper describes the decommissioning of an archive, the intention is not to help other resource providers or users prepare for that eventuality, but rather to take lessons learned from this process in designing storage facilities, developing storage policies, implementing user agreements, and improving storage allocations procedures. Data storage and policies are a major challenge for NSF and the scientific community, and some of the issues discussed here may contribute to the community's dialog about how to address these challenges, including efforts like the Open Storage Network (www.openstoragenetwork.org).

SDSC operated a data archive, as opposed to a digital library, with a key difference being that the service was provided to users primarily to recover previous computational results, with the assumption that it was cheaper to store and retrieve data than to re-compute it. However over time, the type of data stored in the archive expanded to include experimental data and a variety of digital objects. As a resource, data storage is very different from computing time – while computing time is temporal and both users and resource providers understand its use-it-or-lose-it nature, data is sticky and users expect it to persist as long as they deem it necessary to retain. Furthermore, there are fundamental differences between experimental data, unique digital artifacts, and re-computable simulation data. In creating a storage facility, it is important for the resource provider (and its users) to assess and define its purpose and time-frame – ranging from a simple archive for reproducible computational simulations to a digital library for long-term preservation of data.

Especially as the community considers the FAIR (Findability, Accessibility, Interoperability and Reusability) data principles (Wilkinson *et al.* 2016), the benefit of operating a simple archive to serve just data generators versus the broader scientific community should be reviewed. We recommend that users be required to provide substantive standardized metadata for any longer-term storage (beyond 1-2 years). Not only would this make that data discoverable by others, but it would also greatly assist in the user's maintaining their files over time and their long-term disposition. For many of our users, the daunting task of reviewing poorly-documented decades-old files contributed to procrastination and/or simple binary decisions (keep/delete it all). Plans for the documentation and longer-term disposition of computation-related data could be expected to be included in Data Management Plans for funding agency grants and/or added as an additional section in proposals to request computing time (or storage resources) from resource providers.

From an operator's perspective, maintaining an archive requires continuous, adequate funding. While an operator might be able to weather reduced funding for a limited period of time (e.g. foregoing hardware upgrades), a dormant unfunded archive will soon become unusable. At universities, the libraries may have more stable funding than grant-driven research organizations, but this still is no guarantee of adequate funding.

As a corollary to this, funding continuity must be considered one of the highest risks to data preservation. There is substantial discussion in the data preservation community regarding risks to long-term retention of data, including media errors, hardware/software failures, obsolescence of hardware/software/operating systems, *etc*. But it need only take a brief period of inadequate funding for data stored in an archive to be at-risk, especially for large-scale archives.

There needs to be a structure to maintain current contact information for all users. Not only was it time-consuming for us to individually track down user contact information for ~10-20% of our users, but maintaining regular two-way contact reminds users that they have data in the archive and it would be prudent to routinely maintain their files with respect to retention/deletion.

There should be specific provisions between users and the archival operator regarding disposition of their data over time. Particularly if the archive service is offered for free and supports reproducible computational simulations, this agreement may reasonably be tilted in the operator's favor, particularly in the face of potential funding discontinuities – *e.g.* data can be deleted by the operator after N months' notice to the user's email of record – as an incentive for users to keep their information current and to always have backup plans for critical data. But in practice, this kind of clear policy was not considered when SDSC first established its archive and we felt the right thing to do was to get user permission before deleting any user data.

The users that created the data are almost always in the best position to evaluate the value of their data and what data is most important to retain (for their own research or as community databases). Unless unconstrained amounts of storage can be provided free to users, there needs to be some incentive for users to make disposition decisions over time. For example, the NSF-funded national centers generally offer annual, renewable storage allocations (particularly for disk-based storage), but the peer-review process for storage allocations (and accounting) is immature compared to the well-honed review process for compute time, and nearly all reasonable requests are simply approved or left to the resource provider to accommodate. Another potential approach is that archives charge users for long-term storage, even if the charge is only pennies-on-the-dollar of actual storage costs. Alternatively, traditional peer-reviewed allocations could be made as a combined 'compute&storage' allocation, with a fungible exchange rate based on resource provider costs. (This method would probably require improvements in storage allocations, quota systems and accounting.) These approaches provide an incentive for users to routinely review their files for retention/deletion - and leaves the key issue of data value to them to decide.

Finally, there should be incentives for users to document the retention objectives and long-term plan when data are placed in an archive, to review their archival files over time, and retain only the most valuable files. As stated earlier, SDSC operated the archive largely on the assumption that it was cheaper to store and retrieve simulation data than to re-compute it. But most simulation data become less valuable with time, and keeping all data is like just putting more and more boxes up in the attic. For some data-prodigious HPC calculations, it is already the case that the cost of re-computing is less than the cost of storing massive data files for a few years. (On the other hand, the more traditional paradigm of computational resources generating large amounts of data is being reversed in applications such as artificial intelligence/machine learning, where large input training datasets are required as input for computing algorithms; retaining these input datasets is required for reproducibility of the results.) But this trade-off of compute/storage costs is generally only apparent to the resource provider, not the users that compete fiercely for compute time (and typically get storage for free). As mentioned above, fungible 'compute&storage' allocations, with a cost-based exchange between compute time and long-term storage, would bring this trade-off into the user's hands. Furthermore, while most simulation data become less valuable with time, there is increasing recognition that reproducibility of research simulation data is a priority. If users are incentivized to be conscious of resource provider costs, retaining the means to reproduce the data (software, makefiles, VM images or containers) may be less expensive, especially long-term, and better addresses reproducibility of research simulation data.

## Acknowledgments

We want to thank Craig Stewart of Indiana University for reminding us that the lessons learned from the decommissioning process would be of value to others in the community, even at this late date. And we want to thank the many people at SDSC that worked diligently and patiently to 'do the right thing' when we decommissioned the archive.

## References

M. D. Wilkinson *et al*. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Nature, Scientific Data, Vol. 3, Article number: 160018.

R. L. Moore, J. D'Aoust, R. H. McDonald and D. Minor (2007). Disk and Tape Storage Cost Models. Proceedings of the Archiving 2007 Conference, Society for Imaging Science and Technology (Washington D.C.).