Document Distribution Note
Release to Members: Immediate
Release to Public: 11/07/2012

PASIG Conference

Austin, Texas, 11th–13th Jan 2011

About the event

PASIG – the Preservation and Archiving Special Interest Group – was set up in 2007 under the auspices of Sun Microsystems to provide a bridge between hardware providers and the digital preservation community.

WK represented the DPC and other DPC members were present in their own right – Leslie Johnston, (LoC), David Giaretts (STFC / APA) and Jonathan Tedds (on behalf of JISC).

These notes are intended to provide an informal briefing for members of the DPC not able to attend the event. For an authoritative and comprehensive repost readers are encouraged to contact the organisers of the event or speakers directly.

Presentations and discussion

Ray Clarke, Oracle – Preservation and Cloud Infrastructure

Business drivers for cloud infrastructure remain strong – rationalisation, green ICT, power and economies of scale and scope. Tape remains the most reliable media in terms of bit errors. Cloud computing continues to evolve and 'as a service' is now a common place for things like 'software as a service' and 'platform as a service'. Other elements are now coming, such as 'process as a service' or 'test as a service'. Priorities emerging in the cloud domain are for privacy and security concerns. Must haves are: security, ILM, availability, monitoring, governance, sustainability and enterprise management. A lot of the provenance, authenticity issues of preservation are implied or explicitly embedded within the development of security. CDMI service profiles also include a range of tools like de-duplication and retention periods which facilitate preservation.

Leah Schoeb, SNIA – Green storage initiatives

Research programme on power and cooling issues with storage administrators for data centres. The average storage system has lots of moving parts and external requirements in terms of heat and humidity, and bandwidth has different requirements than database operation: so it's hard to be clear about how understand and therefore optimise the distribution of energy around a storage system. The result is a SNIA power efficiency measurement specification which makes it possible to do proper comparison between systems with test metrics for idling and active use of systems. Alongside there is a user guide to measuring energy outputs. Some considerations include de-duplication and compression of data, higher systems tolerances for heat and humidity (thus reducing need for aircon). Active and idle measurements of power efficiency now exist which in turn depend on the use of recommended tools such as types of power metre and temperature gauge. There are different technologies which can be used to give known and measurable improvements in energy

efficiency. www.sniaemerald.com allows you to see existing reports from manufacturers and resellers who have already run their tests.

Michael Petersen, LTDP – A Vision for Digital Preservation

Some future views for digital preservation – here are the key paradoxes that we work on: data will be lost, migration does not scale, access and use models change, cost overwhelms everything complexity does not, we get lost in the detail and lose the objective. So why are we here: there are typically silver-bullet mentalities that we need to deal with. We also need to keep in mind our objective – the human purposes which are not inside the technology. In truth we are still practising until we move to virtual solutions. The solution needs to include the commercial world and not just government services and academic research. Preservation is a service, not a thing or place or tool and therefore we need to talk about this differently. Retention services are needed – we can't do DP without it – but it's not enough. Trustworthy repositories are not the objective in and of themselves. If the goal is 'no dependencies' then the dependence on a trustworthy repository is a paradox. If digital preservation means being 'litigation ready' then the task of preservation begins earlier. Fixing preservation after the fact is not scalable; utility like transparency is the goal. Virtual services allow a more scalable approach to some of the key issues like migration, digital auditing. Self-healing information and systems is the goal. It 'changes the rules', by for example allowing us to reduce the number of copies for example. Interoperability of information objects is the key however. DP needs to shift from repository to service.

Chris Wood (Oracle) – the ~~billion~~ trillion file problem

2008 predictions about electro-mechanical storage were wrong: the 8tb disc drive is now on the horizon as a single disk. But the predictions about SSD were right except that they have been exceeded 1.6TB eMLC SSDs are around the corner. There is a theoretical limit to capacity of disk and tape called the para-magnetic effect rule, which derives from the nature of the device. We thought we had reached that level but we're not close to reaching it yet. Tape is not going away but there is a need for a better discovery on tape to cope with the growing capacity. Optical disk seems to have stopped development through lack of commercial deployment. The current products are enough for the classic use for movies, so the commercial drivers for development are limited. Memory is continuing to develop but the prospect of phase change memory to allow longer lives. Radically new technologies are not panning out as expected. File systems need to be designed at the 'trillions' of files level rather than the thousands where we currently work.

Gary Francis (Oracle) – Data Storage trends and directions

There's a sense that long-term-digital-preservation is an oxymoron. We know that the data will be effected by bit rot and there will become a cost in the long term to prevent migration if you don't get onto it reasonably quickly. 60 years of storage innovation – Remington Univac tape was invented in 1951 and the IBM RAMAC disk in 1956. There's been growth since then! Costs are continuing to drop and the costs for tape are dropping more quickly than the costs of disk. There was a real hiccup because of the floods in Thailand because suppliers and manufacturers could not meet demand. In fact the costs of disks has been going up lately. The improvements are not all the

same - sometimes to do with the surface of the media, or the quality of the readers, or the write technology. So there are normally a range of engineering solutions which allow researchers to work round the thresholds we might predict. Linear Tape File Systems describe the content and layout of a tape at a fixed point which therefore enables enhanced searching and recovery across a tape therefore allowing the density of tape with the flexibility of the disc. Tape therefore remains the most viable media for archives: there's lots of room for growth and new architectures make it easier to use. Disk densities are not growing as quickly and they are actually becoming more like tape in how they perform.

Bram van der Werf (OPF) – SCAPE and other current OPF project

Bram introduced three projects which OPF are working on and which are designed to create something of a community which in turn will support the development of a community. SPRUCE and AQUA have supported the deployment of hack-a-thons to help us develop and refine requirements and develop working solutions. It allows 'live' testing and prototyping, with appropriate community building. Open source software is not a solution on its own, and there is a need for on-going deployment and development. This leads to a view of 'micro-services' which are large enough to be useful but small enough to be slotted in and slotted out of an architecture, and also maintained with some care. The number of bug reports is a better test of software than the number of downloads: active users reporting errors show that the tool is actually being used and are likely to produce better tools too. OPF hosts code and an issue list which makes bug-fixing transparent. Requirements gathering has been the real problem with working with libraries so OPF has spent quite a lot of effort to improve the way that requirements are gathered. It also helps map tools and users.

Cherie Ekholm (Microsoft) – The obsolescence issue for file formats

Standards enable interoperability and transparency and also to widen engagement. But they take forever to develop. Microsoft is currently working on three main areas of file format standards and putting them into their forward plans – ODF, Open Office and PDF. PDF/A3 is expected in early 2013. Move to the cloud adds complexity to the interoperability and standardisation but the core issues remains – where is my document and how is it sorted; how will I access it in 100 years from now.

Don Post (Saving the Digital World) – Fending off the Digital Dark Age?

A substantial amount of data which should have been retained will not be stored in trusted repositories in 20years time – perhaps as much as 50%.

Jay Boisseau (Texas Advanced Computing Centre) – Perspectives on Data in Science

Data intensive computing has grown in many disciplines, and expertise in storage has of necessity grown with it. TACC is coming round to digital preservation, but from a distinctive angle rather than digital preservation first. Computational science is new – several decades old at most – though experimental and theoretical sciences have been around for thousands and hundreds of years. But science has been one of the drivers of computing in order to solve the mathematical problems behind much research. Big problems require big computing and big visualisation and big storage.

It's not just modelling and simulation. Data-enabled science is enabled by digital data collection, often without formal mathematical laws. Text mining and statistics are showing that data is allowing the creation of first principle laws rather than the other way round. Advanced computing technologies enable discoveries that advance science and society. TACC has made a lot of progress towards supporting research – primarily as simulation and visualisation. TACC has lost simulation data in the past which was bad news but not disastrous because the data could be reconstituted. But more serious was a loss of underlying data from sensors which were backed up. They realised that they needed to rehearse data loss and how they would bring things back from deep store. Genomic data is the biggest usage right now. TACC has created specific posts for archivists to help make sense of the data – persistent data management expertise to ensure value of the data. The archival store – called RANCH – has large scale storage and is based on high density tape storage with a mix of disk storage offering a range of performance and density within the architecture.

Jason Piers (Familysearch) – petabytes of storage on tape

A lot of 'fieldwork' data gathering which makes its way into the preservation system. The pipeline is pretty well developed from source records to ingest to process to enhancement to distribution. Web access alone required 2 petabytes of spinning disks. They use SDB to do core manahgement which designs workflows, characterisation, provenance, AIP tracking reporting and such. Parallel tere is a tape storage management system which provides a media inventory, physical locations, media errors. The hardware is based on Dell Servers with an Oracle Storagetek tape library with capacity for 10,000 tapes. Software is built around SDB bit there is a family of java tape library tools running under linux to manage the tape. Tape has come out as the preferred mechanism for storage after an analysis. Designing a tape library is also subtle: for example ability to update content needs to be examined. 400million Jpeg images in 2 copies: that's but and it creates an enormous migration challenge. There is a 7 part workflow for putting things into the tape library and is reasonably well worked – transfer tape, tape ingest reader, SDB ingest, staging disk, tape service queue, tape service to archive tape. Images are compressed and chunked before being transported. Packages are not split across tapes, and 3-4% space is left on each tape to allow updates if required.

David Minor (Univesity of California in San Diego) – Research Campus infrastructure project

UCSD has written a Blue Print for an E-university. They are working on co-ordination of storage and computing infrastructure for the university which includes actual physical network and servers, and in addition there is now a two year pilot project on data curation which resides on top of this storage and digital asset management infrastructure. Examples include a 'brain observatory' which looks at 3-d scanning of brains, an Open Topography facility which is using DoI's an archaeology group in the Levantine Archaeology Lab which works in the field out of range of the network(!), Scripps Oceanography Laboratory and the Laboratory for Computational Astrophysics. Data management plans are helpful and a local tool has been developed to encourage submissions – very popular and they've had to turn down people who have wanted to do some training.

Reagan Moore (University of North Carolina) – IRODS technology at University of North Carolina

IRODS offers a policy based data environment. The reasons for preservation have to be assembled and implemented based on a formally expressed data management plan that includes a variety of rules Success is measured against policy. It's not simply about preservation because it manages distribution, retention, deletion and such. But it's pretty sophisticated system because it splits clients – which can be distributed – from storage and independent policies and services. Micro-services are particularly well deployed in this environment.

Dan Stanzione (Texas Advanced Computing Centre, TACC) – Digital Preservation at TACC

TACC has three classes of data – local user systems around 15PB; archive systems with around 100pb and global storage of true collections which are around 1pb. There is limited back up for users and so users need to actually put data in the archive: they are clear that local 'scratch' systems are not intended for long term use. 15,000 users – this means that there are 15000 different people responsible for indexing the data and TACC doesn't have to attempt a master index of billions of files in 100PB of storage. In reality the legacy data is small in comparison to the incoming data which is where the growth is. So over a 5 year period a large data set becomes noise, so after 5 years large data sets become low cost to manage. So TACC promises 5 year deals to its users and after five years it can decide what to do next, most likely 'throw it in for free'. They don't make permanent arrangements. But the trillions of files that are coming emphasises the need for machine readable data which in turn means emphasis on ontologies and standards. And if we don't decide, the market place will decide (badly?). Looking at the costs, the cost of compute is much much greater than the cost of storage and this needs to inform the architecture.

Olivier Rouchon (CINES) – Digital preservation and CINES

CINES was created for high performance computing, funded by government to provide a national centre. High performance computing means a large input area for 'scratch data' – currently 700tb and similarly large output area, also about 700tb produced out of these processes every year. Every year they need to decide whether to keep or dispose of these data sets. This includes a large number (though low volume) of documents as well as data sets which have distinctive issues. Researchers are asked to consider whether they want to keep the data for more than a year – and currently about 30% tick the box that says they want the data to be preserved. This is spread evenly across the different research types but is neutral on the size of the data. From 2012 CINES will offer a medium term preservation service for research data. This will provide a platform at national and European scope which in turn would improve the scope and scale of what CINES can offer. Current challenges are to involve more users, encourage more useful metadata, develop more clearly cut business cases and to engage more effectively in a pan European infrastructure programme.

Dusan Kantuscak (National Library of Slovakia) – Digitisation of the National Library

Plan to digitise Slovak heritage, especially books published on acid paper and provide access in digital form. This will be new type of access but also protect the books.

Jonathan Markow (Duraspace) – Duraspace to DuraCloud

Duraspace has been talking to institutions about their requirements and expectations for cloud services. The top five priorities for data storage which emerged from a workshop of Duraspace users were … 1. Connect the operational and archival phases of the data management lifecycle 2. Create simple workflows across data management lifecycle that automatically capture metadata and provenance and creation incentives for additional metadata creation 3. ensure confidentiality, security, privacy and predictability for data in the cloud (trust and control) 4. Automate metadata 5. Create interoperability of operational systems and archiving solutions. Some other principles – open source enterprise software solutions, capture data close to the source, don't interfere with researchers processes, provide incentives added value for metadata and create simple workflows. In response to the workshop they have created a high level set of user stories to help with the functional specifications

Natasha Milic-Frayling (Microsoft Research) – Services for long term access to digital content

What are the principles of quality assurance in content preservation? What do we know about re-use and how can we make claims about the future value? The long term preservation of data almost always depends on being clear about the value and the use of data. You can't box it up: you need to open it up for use. SCAPE project is looking at some of these issues by examining the performance of migration tools. The available conversion tools can be assessed: and mostly these tools exist in the market place because there is a current need for them – not a theoretical need in the future. There needs to be an explicit set of measurement for conversion – that's what's missing. Comparing the tools will then become easier. Digital objects only exist fully when the file is being instantiated and rendered for a user. It's like sheet music which needs an audience, instrument and player.

Mark Evans (Tessella) - Proof of concept cloud based solution

Tessella have been investigating how to put their technology stack into the cloud and what the implications are. SDB is now in its 4th generation and the whole application is built into a set of frameworks which in turn co-ordinate (micro)services which can be aligned to workflows that are defined for a specific instance or policy framework. It is neutral on storage and APIs are published in order to assist with integration into different systems. The problem for preservation has been that individual instances are expensive, and this is a barrier to participation. There are different ways around this barrier, such as open source or shared infrastructure, though both have their own problems and require an initial (capital) investment. A hosted service is an obvious alternative for people who can't get the capital investment and that leads to consideration of the cloud. But there are different ways to do this and it's not clear whether this does indeed lower the barrier. Three different possibilities: multi-tenancy of a local single deployment; storage in the cloud with the application stack residing locally; and fully hosted solution where data and application are in the stack. The 2nd option turns out to be a problem because the storage providers charge for every time you want to recover the data – so there's a bill for every fixity check and this ends up more expensive. So Tessella are now working with Amazon to develop an example service that they think costs something like 3000USD a year for compute (ie not including storage). That changes the market really very much. There are lots of policy issues to resolve of course but these need to be resolved on a case-by-case basis.

Richard Moore (San Diego Supercomputing Centre) – SDSC Cloud Storage Services

Risks to digital preservation are actually about human error and acts of god rather than technology failure and the continuity of funding is problematic. Traditional storage systems have home user systems, a high bandwidth parallel file system and a tape based archival system. The limitations are that it is slow to access and has poor interfaces. All too often the archived data, particularly HPC simulation, is 'write-once-read-never' but not inherently because of the value of the data, but because of the difficulty of access. That's not good economics. SDSC cloud is a paradigm shift in storage. It allows easy access and collaboration and that means data is more likely to be used. Data is 'always there and instant access'. It allows multi-users, simple interfaces to data and data management, encryption and security, transaction and provenance logging, interfaces to commercial and open source products. It's scalable in bandwidth and storage. Applications can be put in front of data and it can provide web based interfaces for external users. SDSC has chosen to build this in Openstack Swiftcloud software which is open source and used by a number of communities. Scale is not the issue – it's the ease of use that matters. It's not built for free storage but is a cost recovery and the cost is metered.

Mark Carlson, SNIA – Storage and data management in a post-file world

File systems are running out of oomph as content is highly fragmented and highly distributed. Local storage is being lost in return for the speed of flash memory. We need a global namespace, support for rich metadata, http access, data services for transparent users, we need to support large objects. SNIA's Cloud data management interface (CDMI) standard is intended to meet these new requirements. These requirements are effectively metadata and they express data requirements as metadata. CDMI is shortly to become an ISO standard

About this document

Version 1	Written at conference	11-13/01/2012	WK
Version 2	Distributed	16/01/2012	DPC members, PASIG group