Report from Dagstuhl Seminar 12472

Is the Future of Preservation Cloudy?

Edited by

Erik Elmroth[1], Michael Factor[2], Ethan Miller[3], and Margo Seltzer[4]

1 **University of Umeå, SE,** `elmroth@cs.umu.se`
2 **IBM Research – Haifa, IL,** `factor@il.ibm.com`
3 **University of California – Santa Cruz, US,** `elm@cs.ucsc.edu`
4 **Harvard University, US,** `margo@eecs.harvard.edu`

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 12472 "Is the Future of Preservation Cloudy?". Our seminar was composed of a series of panels structured as a series of brief presentations followed by an open discussion. The seminar started with a session introducing key concepts and definitions and illuminating the vast array of perspectives from which attendees were addressing issues of cloud and preservation. We them proceeded into a discussion of requirements from different types of communities and a subsequent discussion on how to protect the data and ensure its integrity and reliability. We next considered issues related to cloud infrastructure, in particular related to management of the bits and logical obsolescence. We also considered the economics of preservation and the ability to reuse knowledge. In addition to these pre-planned panels, we had three breakout sessions that were identified by the participants: automated appraisal, design for forgetting, and PaaS/SaaS for data preservation. After the executive summary, we present summaries of the panels and reports on the breakout sessions, followed by brief abstracts from a majority of the seminar participants describing the material they presented in the panels.

Seminar 19.–21. November, 2012 – www.dagstuhl.de/12472
1998 ACM Subject Classification D.4.3 File Systems Management, D.4.5 Reliability, D.4.6 Security and Protection, E.2 Data Storage Representations, E.5 Files, H.3.2 Information Storage, H.3.4 Systems and Software, H.3.6 Library Automation, H.3.7 Digital Libraries, K.6.5 Security and Protection
Keywords and phrases long-term preservation, cloud storage, provenance, obsolescence, data access, storage systems
Digital Object Identifier 10.4230/DagRep.2.11.102

1 Executive Summary

Erik Elmroth
Michael Factor
Ethan Miller
Margo Seltzer

Two significant trends in data management are emerging: data is moving to cloud infrastructures and an increasing fraction of data produced is born digital. We risk losing all record of born digital data if we do not take explicit steps to ensure its longevity. While each of these trends raises its own set of questions, our seminar began with two fundamental questions at the intersection of these trends: What role should the cloud play in preservation? What steps should we be taking now to preserve the future of today's digital artifacts?

We addressed these two questions by bringing together a diverse cohort of approximately thirty participants. Our participants consisted of researchers from both academia and industry, representatives from cloud providers, and archivists and librarians from memory institutions. Every participant was responsible for some aspect of the program, and the workshop was characterized by lively debate. There were four primary outcomes of the workshop:

1. We identified key functional requirements that are critical if cloud infrastructures are to be used for long-term digital preservation.
2. We identified topics where we were unable to reach agreement; since we are trying to look into the future, while not satisfactory, it seems likely we will need to wait until the future to resolve these debates.
3. We identified several specific problems requiring further work and brought together groups of people interested in pursuing those areas.
4. We identified several areas that we were not able to address, either because we lacked the expertise in the room or we ran out of time; these areas represent opportunities for subsequent workshops.

Perhaps the most pressing issue with respect to existing cloud infrastructures is the lack of standardized APIs. If data are to outlive any particular organization, then it is crucial that archives span organizational boundaries; standardized APIs make this dramatically easier and more robust. There was also agreement that some form of automated appraisal was important, but there were no concrete ideas about how to do it.

We had lively debate around the long term cost of cloud storage, in particular public clouds; since this debate depended upon assumptions of future costs, the future will ultimately resolve the debate. We also had much discussion around the importance of logical preservation and whether the modern world, with readily available open source viewers has made the need for logical preservation obsolete.

Several small working groups coalesced around the areas of: archival exit (how do you get data out of an archive), the technical design of preservation-as-a-service (PaaS), technologies for ensuring that data is "forgotten", and searching distributed archives. We are hoping to see these small groups evolve into productive collaborations that continue the work begun at the seminar.

Finally, there were a number of areas related to using the cloud as a preservation service that we were unable to address. For example, what legal issues arise if companies undertake digital archival initiatives? Is there a legal definition of "deletion" of data, and is it practical? Where does "record management" end and "archival" begin? Who is the customer for long term preservation? Is it the data provider? Or perhaps it's the data consumer? What happens to archived data if payment cannot be made? What is the economic model behind long term archival? These and other questions provide ample opportunity for further workshops on this topic.

Organization

The workshop was organized around a series of 90-minute sessions, each of which began with one or more short presentations followed by a moderated discussion. We had one person scribe each session and the session moderators produced the session summaries that appear in this report documenting each session. We also devoted one session to smaller breakout groups, who reported back in our closing session.

2 Table of Contents

Executive Summary

Panel Discussions and Session Summaries

Working Groups

Overview of Talks

3 Panel Discussions and Session Summaries

3.1 Opening Session

Mary Baker (HP Labs – Palo Alto, US)

One of the lovely and remarkable features of this Dagstuhl seminar was the diversity of disciplines represented by the participants. Participants introduced themselves as coming from a wide variety of scientific, industrial, government, cultural, and academic institutions. Their areas of expertise included digital preservation technologies, storage and database products, huge cloud storage applications, memory institutions, digital curation, provenance of digital content, medical records and their associated policies, trace archives of distributed systems, the economics of digital preservation, scientific computing, supercomputing, and so forth. To illuminate the different issues faced by participants, we asked everyone to describe what digital preservation means to their communities and what preservation problems the cloud solves and does not solve for them.

This diversity of participants also posed a challenge for us: finding a common vocabulary and set of concepts for digital preservation so we could avoid confusion and make forward progress. We therefore used the first session to address the problem. We introduced the goals of digital preservation by claiming that digital assets stored now should remain accessible, usable and undamaged for as long as desired – beyond the lifetime of any particular storage system, storage technology, or storage vendor, and that this must be done affordably. The main discussion of these goals centered around the meaning of undamaged, since it has different meanings depending on the kind of asset being preserved and the preservation purpose. For some digital assets undamaged means the bits must not change. For others the bits may change but the meaning must remain the same. For yet other assets, the contents need to remain usable.

The terminology and concepts we presented included:

- “physical” or “ bit preservation” (and why it is still a challenge to do affordably),
- “logical preservation” (and the Performance Model used by the National Archives of Australia to illustrate the problem),
- “metadata” (which is too broad a term according to some of the participants), and
- what different communities mean by “preservation.”

When we preserve an asset – what are we preserving? For instance, for a book do we just save the text? How about images of the pages? What about saving the political and cultural context in which it was published? For applications, do we save the entire ecosystem in which they run? Or are screen shots of the various user interface activities sufficient?

We deliberately avoided defining “the cloud” and this came back to haunt us later in the seminar!

3.2 Domain Specific Needs

Erik Elmroth (University of Umeå, SE)

This panel, focusing on domain specific preservation needs, included Ian F. Adams (University of California, Santa Cruz, CA, USA), Dirk Nitschke (Oracle, Hamburg, Germany), and Gillian Oliver Victoria University of Wellington, New Zealand) and was moderated by Erik Elmroth (Umeå University, Sweden).

The goal of our panel was to provide a concrete examples, from a variety of domain, about what types of information need to be preserved in the cloud, for how long must they be preserved, for whom, by whom and in what type of cloud? We wanted to move beyond common preservation requirements and focus on requirements specific to one or more domains. For example, when discussing memory institutions, forgetting became important; it is equally important to intentionally and thoughtfully decide to not remember (preserve) things as it is to select things to be remembered.

A second topic that created much lively discussion concerned the risks of storing information in public clouds. This flowed seamlessly into a discussion of existing laws and the challenges of jurisdiction – how do cloud providers and customers come together when legal requirements differ for each of them? Someone described the Megaupload case where customer content was seized, because some customers had uploaded data without appropriate copyrights, as an example of the legal complexities that arise.

The legal discussion then flowed naturally into a discussion concerning the difference between selective archival and censorship and whether and whether it is at all feasible for archivists to decide what, from the massive data stored in clouds, should be preserved. For example, the Internet archive uses a statistical, or perhaps random, approach to archiving, not a human-centered manual one.

When the conversation moved to scientific data archiving and high-performance computing applications with very large data sets from climate, particle colliders, etc., we began trying to distinguish between data and information. In cases where data can be reproduced (which was agreed to be a fundamental concept), should archives leverage this capability to reduce capacity needs. For example, if a person's DNA sequence has been computed, need we save that sequence or can we discard it and then resurgence it later? Of course, once we begin discussing data reproduction, the challenge of preserving software and execution environments becomes critical.

When taking the business perspective, the discussion touched upon cases where data are truly mission critical and data loss must be avoided "at all costs." In fact, the requirements may be even more stringent – in some cases, it is not only necessary to make it possible to obtain archived data, it might need to be always available relatively quickly. This generates both bandwidth and latency requirements. As with memory institutions, the issue of forgetting also came up in the business context. For example, regulations may require retaining financial data for a defined period, but beyond that period the data can become a liability. Unlike memory institutions, businesses do not have archivists, so end-users, untrained in archival and digital preservation, are responsible for identifying data to archive. A long discussion on incentives led to the conclusion that finding a way to make money out of archived data was probably a more effective incentive than a legal frameworks.

After we covered the domains individually and had a more encompassing discussion, someone observed that there is no accidental digital preservation, while there is accidental physical preservation. Such accidental physical preservation, e.g., finding a cache of lost letters in an attic, has turned out to be crucial for historical investigations. The discussion on whether we can ensure that something is not preserved led to our distinguishing between "digital preservation" and "digital archaeology", e.g., the ability to recreate old computer games when the physical media has become obsolete.

3.3 Protecting the Data

Margo Seltzer (Harvard University, US)

Our panel, comprised of Jean Bacon (University of Cambridge), Ewnetu Bayuh Lakew (Umea University), Peter Pietzuch (Imperial College London), and Ken Moody (University of Cambridge), and moderated by Margo Seltzer (Harvard University) represented a diverse set of opinions on what protection meant, how it applied to preservation, and what about the cloud made it different. The presentations addressed issues ranging from the legal challenge when data moves across geopolitical borders, to the privacy challenge when the data being stored was health-related, to the technical challenges of describing and translating security policies.

There were two distinct parts of the session – first there was discussion about some of the key aspects of protecting data and second, it became clear that the attendees needed to come together to agree upon terminology surrounding clouds. After much discussion on this second point, Alexandr Iosup (Delft) provided the NIST definition of a cloud, "cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." This provided a specific definition that was used for the rest of the workshop.

The technical discussion ranged over three main topics: the relationship between provenance and preservation, the role of the cloud in preservation, and the differences in preserving private versus public data. Some participants questioned what provenance had to do with preservation, but there was consensus that documenting ingest processes, migration efforts, and emulation requirements was critical to ensuring meaningful access to digital objects in the long term.

Unsurprisingly, the group spent a great deal of time discussing the role of the cloud in preservation – asking and discussing questions such as whether the cloud is simply a technology, a delivery mechanism, a platform, a business model, etc. Today, clouds are clearly only part of a solution – they are not fundamentally preservation systems. However, if they are to be part of a long term preservation solution then there are important avenues for research and development. For example, while many existing infrastructures claim to support "standard" APIs (e.g., Amazon S3), experience with LOCKSS and other systems suggests that they are not; you cannot write an application or service to a single API and have it run across multiple providers. Therefore, vendor lock-in, is a huge issue when contemplating long term preservation. More fundamentally, the group asked whether commercial providers were the right parties to provide preservation platforms, or whether governments or memory institutions were better suited to the task. No clear consensus emerged from this part of the discussion.

Finally, the group turned to questions surrounding the preservation of both public and private data. Private data requires adequate security, but what can be considered adequate if we are trying to preserve objects for tens or even hundreds of years? With today's computational resources, it's possible to break a lot of encryption with a hundred years of compute time; with future breakthroughs, what can we consider truly protected? In many ways, public data presents a much easier problem, because archives need only concern themselves with ensuring authenticity. However, it was pointed out that some public data starts out as private data, e.g., census data. Therefore, it is not sufficient to simply focus on public data and ignore the challenge surrounding private data. There was no group agreement on what kinds of guarantees for long term security were sufficient.

3.4 Reliability and Integrity

Lawrence You (Google Inc, US)

The panel on Reliability and Integrity had four presenters: Gerhard Schneider, University of Freiburg, Germany; Nikos Chondros, University of Athens, Greece; André Brinkmann, University of Mainz, Germany; and Lawrence You, Google Inc, USA.

Our four panelists each made short presentations on different facets on cloud storage reliability and integrity, approached as a user (Schneider), verifier (Chondros), exploiting provider diversity (Brinkmann), and as a provider (You).

Common themes were that cloud services vary, due to different product offerings and commitments. There was consensus during the presentations and questions that data preservationists must use multiple cloud providers to ensure long-term integrity and reliability if they are to use them at all. Presenters raised many points on numerous challenges that still exist with cloud storage services: that storage cloud services alone are insufficient for integrity when access also requires software; that verification requires access and re-fetching copies, and costs must be factored in; that cheap storage is unreliable; and that consumer versus enterprise storage pricing meet different business needs.

The discussion following the presentations produced a lot of questions, particularly in regard to measurement of reliability and types of failures to guard against. Cloud as a model for preserving data has some requirements that differ due to time scales that are long, and reliability requirements where failure modes of error rates (for example, the number of "9"s of reliability) are independent of trust in the cloud providers, which are businesses. Diversity is important, but so is cost of the service and cost of failure.

We wrapped up the discussion with the basic question:

Do we think reliability and integrity are a solved problem?

There was some disagreement, but there were a number of questions/comments from the wrap-up:

We have (meaning people here use) a solution using a combination of techniques. It's not possible to say it's solved, because not all information is classified, or it is classified incorrectly. Different cloud pricing and business models differentiate service/reliability. Cloud providers have a core competency in bit preservation, while archivists do not. Will the cloud cost too much? What is the real cost of reliability (for preservationists)? Including cost of failure?

Failure comes in many forms: natural disasters, insider abuse, black swans. A moral hazard: cloud providers want to claim 11 nines and the buyer wants to believe such claims, but will everyone be gone when the claims are tested? Is reduced reliability and more providers a better solution?

3.5 Preservation Storage / Cloud Services Issues

Hillel Kolodner (IBM Research – Haifa, IL)

The panel had four presenters: Joanne Syben (Google) (lead), Alexandru Iosup (Delft University of Technology), Sam Fineberg (HP), and Hillel Kolodner (IBM).

Joanne Syben spoke about technical aspects of current commercial storage clouds and issues regarding user-owned and user-derived data. In current storage clouds, there are several interesting combinations of durability and availability. For example, Amazon currently offers Reduced Redundancy Storage for low durability and high availability, Standard S3 for high durability and high availability, and Glacier for high durability and low availability. Joanne mentioned life cycle issues that need to be automated: migration to new hardware, migration to cheaper hardware for cold data, and deletion (including high volume deletion). She also discussed the appropriateness of some hardware options including tape and shingled disks.

Joanne also discussed the distinction between user-owned and user-derived data and the resulting technical and legal issues. A significant technical issue around user-owned data is how should it be curated? Methods include relevance ranking, time based ranking and explicit annotations. Regarding legal issues, given that data can be considered an inheritable asset, what rights should heirs have to curate it? What about jointly owned data in the event of a divorce? Search quality depends on user-derived data. User-derived data also needs to be curated, yet it is hard to predict what user-derived data will be useful in the future. There are also privacy issues.

Following Joanne's presentation there was a discussion about the difficulty of deleting data; this discussion led to one of our breakout sessions. There are both legal and technical issues. Legally there are regulations, e.g., a legal definition of secure deletion. Technically, there may be many copies of a data item that need to be deleted, e.g., multiple copies on-line and multiple copies in backups, and it might not be possible to find all of the copies.

Alexandru Iosup presented the idea of a distributed systems memex, i.e., logging and preserving the entire history of a distributed system; he also presented several experiments that he and his team have done in this regard. These include the Grid Workloads Archive, containing six online traces, the Failure Trace Archive, containing 25 online traces, and the Game Trace Archive. Alexandru went on to discuss the benefits of using clouds for the storage of the traces, e.g., simplification of management, reduced cost for the infrastructure, and the availability of computation close to the data. However, there are also challenges: processing close to the data leads to vendor lock-in, and the challenges of migration and security. Finally, Alexandru presented several experiments that he and his team have done regarding the performance and performance variability of AWS services (EC2 and S3) and DropBox. This presentation was followed by a discussion on long term funding for archives and curation of archives; although, no conclusions were reached.

Sam Fineberg presented some advantages and drawbacks of using clouds for preservation, and then briefly introduced SIRF (Self-contained Information Retention Format). Advantages of using cloud storage for preservation include the simplification of migration and economy of scale. Drawbacks include the elimination of transparency and control. Sam argued that preservation should be considered SaaS (software as a service) rather than IaaS (infrastructure as a service).

There was a discussion about what would be the SaaS services; possibilities include logical migration, physical migration, and verification services. David R. pointed out that there are already preservation-as-a-service offerings available today, but why should we trust such a service. One service called out in particular was DuraCloud which has the benefit of being open source. There was also a discussion on whether we should be striping the preserved data over multiple such services.

Sam also presented SIRF, which is Storage Networking Industry Association effort to define the logical equivalent to storage boxes which can be seen in a physical archive. It needs to be self-describing, self-contained, and extensible. One concern that was raised was the scoping of metadata to a logical container in SIRF; is it too easy with such scoping to have broken links, e.g., if single data set doesn't map cleanly to a container.

Hillel Kolodner raised issues regarding the requirements from a storage cloud to support a preservation system. For example, should a preservation system provide deep support for preservation, e.g., very high reliability and end-to-end security. Or should storage clouds provide basic support (e.g., lower levels of reliability) and preservation systems ensure high reliability themselves, e.g., by storing data on multiple clouds. In this case, what is the basic support that is needed in each cloud and how should the preservation system leverage it? Issues include the costs for replication, integrity checking and provenance. For example, we would not want three way replication per cloud when replicating across three clouds – this would mean keeping nine copies. Integrity checking would also be occurring repeatedly in each of the clouds and then also by the preservation system. Provenance would also be hard because it would be necessary to do full provenance on each cloud and also in the preservation system.

Hillel also raised the question whether there are advanced features that storage clouds could provide that facilitate preservation systems. For example VISION Cloud provides support for rich metadata and allows objects to be found based on their metadata values. It also supports safe and secure computation in the storage system, e.g., which can be used for integrity checking. Finally, storage clouds can provide support for the secure handling of data, e.g., secure isolation between the data of tenants and users, geographic constraints on the placement of data and secure delete.

Following Hillel's presentation there was a discussion about some of the issues raised. A cloud provider can really go out of business. One solution that was suggested is that the data owners can physically own their own disks. However, this raises several issues. It is less elastic. It is harder to achieve economies of scale. And it could raise problems when necessary to migrate data to new physical media.

3.6 Handling Logical Obsolescence

Liuba Shrira (Brandeis Univ. Waltham, US)

The panel had four presenters, Michael Factor (IBM), Natasa Milic-Frayling (Microsoft), Matthias Grawinkel (Universität Mainz), and Liuba Shrira (Brandeis) (lead).

Matthias Grawinkel surveyed the factors governing the longevity of archived objects, considered the constraints on the storage media imposed by different longevity time scales. He then discussed the importance of techniques for preserving the object interpretation environment, describing the extreme case of the strong versioning approach adopted by NASA, which allows reproduction of the exact processing environment.

Natasa Milic-Frayling focused on the computational nature of the digital artifacts. She defined digital preservation as "enabling digital artifacts to be instantiated in a contemporary environment," and argued that the key to ensuring long-term preservation is providing efficient software development environments for developing "bridging components" such as format translators and virtual machine adapters. She then described an Azure based service for format migration designed using this approach, as part of the SCAPE project, that is extensible in both formats and data storage.

Michael Factor shared lessons from his work on the EU funded ENSURE project, focusing on three points,

1. supporting requirements-based preservation plans that protect different data in different ways;
2. preservation of metadata, including ensuring that meaningful cost effective metadata is available for all entities at ingest, since it may be impossible to reconstruct metadata afterwards;
3. automation of transformation and verification, achieved in ENSURE by a combination of a workflow engine, virtual appliances for short term preservation, and computational storage supporting transformation and verification for long term preservation.

Liuba Shrira described an efficient just-in-time transformation service for handling logical obsolescence in a cloud-based preservation system. Such a migration service defers transformation until the object is used, avoiding transforming work when a new object format becomes available. The challenge for just-in-time migration is how to avoid introducing transformer dependencies on future versions. Shrira described a framework, developed in her research on database upgrades, that supports efficient just-in-time transformation and eliminates transformer dependencies on all but a single predecessor version. She raised the question of what would it take to achieve analogous transformation system properties at different levels in the cloud-based preservation system software stack.

The follow-up discussion centered around two key issues, the implications of managing logical obsolescence at different levels in the software/hardware stack and the challenges that arise when the digital artifacts to be preserved and the techniques of preserving them have a computational component.

The discussion started by reviewing the accepted roles of migration and emulation in handling logical obsolescence, (migration being the standard tool with a downside of losing information, and emulation being the fall-back for when migration fails), and then discussing the open problems with both approaches.

An issue for emulation is at what level to emulate. VMs are a standard emulation level, because the represent a "slow moving layer". This works well in the short term but poses difficulties for long term preservation. VMs have not been designed with migration in mind, they are complex and have dependencies that need to be handled when eventually migration takes place. Could a different layer work better? An interesting example of a different layer is IBM's successful move to a different hardware architecture (S400) at an internal OS level rather than machine level. A standardized API for emulation would have a dramatic impact, but few participants were optimistic this would happen any time soon.

An issue for migration is that the layer above may become obsolete as well, making it hard to resort to emulation at the point where migration does not work anymore. For example, it is hard to rebuild renderers after they are gone without a good enough "live" artifact that preserves the experience. Of course, emulation does not have to be complete, but it is hard to know what path will be useful in the future, especially, if we want to ask old data new questions, in addition to old questions.

Another issue for migration is that transformers are programs too and need an environment in which to run. Could the cloud help with keeping transformers alive by maintaining libraries?

Logical obsolescence of digital artifacts manifested (and consumed) using computation rises several issues. On the one hand, programs may have assertions, and executable specifications, and these can be used to assemble a testable verification, e.g., to verify whether rendering is correct. A cloud may even help with high-fidelity recursive emulation at multiple levels that results in substantial computational overhead. On the other hand, the more complex the digital experience, the harder we may need to think what obsolescence might mean? For example, the gaming industry has economic incentives in preserving gaming experiences, re-issuing them in the future, just like Hollywood is preserving old movies. What does it mean to preserve a multi-player game experience? Similarly, what does it mean to preserve a spreadsheet combining live results from multiple distributed data streams? The participants agreed that a new principled approach must be developed for preservation of "digital experiences".

Throughout the discussion the participants considered how cloud infrastructure could impact the handling of logical obsolescence, we had more questions than answers, but the consensus was that by further separating the owner of a digital artifact from the infrastructure that handles obsolescence, the cloud makes it harder to exploit native approaches. On the other hand, the cloud could provide the advantage of more plentiful computational resources or have available a wider choice of preserved standard environments.

3.7 Knowledge Re-use

David Giaretta (APA, GB)

URL https://dl.dropbox.com/u/6959356/Knowledge use notes.docx

This session had three presenters: David Giaretta (APA), Christoph Becker (TU Wien), and Milena Dobreva (University of Malta). It covered the importance of knowledge re-use in preservation.

The session began with David Giaretta's talk "Knowledge, Value and Services for preservation with some thoughts on clouds". Drawing on quotes from Neelie Kroes and a report from the High-Level Group on Scientific Data, David illustrates the desire that results of publicly funded research provide valuable assets and claims that preservation, by its nature of facilitating reuse, makes assets more valuable. He goes on to claim that while traditional archival focuses on documents, we should focus more on data, because data is more challenging – data value increases when we know the semantics of the data and can combine it with other, potentially massive, pieces of data.

This leads to the challenge David puts before us: how do we make the stuff we preserve more valuable? David proposed that a way to look at what is needed for preservation was to look at threats arising from changes in technology (hardware and software), environment, e.g. e.g. name resolvers, and tacit knowledge of users

To counter these threats the mantra one tends to hear from libraries is: "Emulate or migrate." However, in line with the argument above one can see that emulation works well with data only in special cases, because one can repeat what was done in the past rather than doing new things.

Turning to what kinds of semantics – or, more generally, knowledge – is required, the simplest are things such as units of measure, etc. Next comes "Representation Information," the OAIS term for what many might call metadata (a term David chastises us all for using, because it is ill-defied). Noting that emulators are also a type of Representation Information one can restate the mantra as "Add Representation Information or Transform. Or move to another repository" Another complexity which one must deal with when one wishes to keep data valuable by keeping it usable is to recognise that any piece of Representation Information must itself be usable. This introduces a potentially problematical recursion, which OAIS resolves through the concept of a "Designated Community." The advantage is that this requires ways to figure out how much Representation Information needs to be provided for some community to be able to understand and use the data. This is precisely what is needed to add value to data by making it more widely usable and for data from many sources to be combined easily.

The CASPAR and its successor the SCIDIP-ES projects are addressing the preservation and use of digitally encoded information by providing tools and services to supplement what a repository does by allowing curators to

1. know something has changed
2. identify the implications of that change
3. decide on the best course of action for preservation
4. determine what RepInfo we need to fill the gaps

Michael Factor asked if there is added value by putting the services in the cloud, to which David's reply was that this must be the case – allowing greater resilience and pooling of resources.

Liuba asked for an example of adding value by combining data. David pointed to all the climate change work, the studies combining sociological data, health data with satellite information about temperature changes. Christoph added examples such as combining 60 years of nutritional data with climate change.

The second talk was by Christoph Becker, titled "Some thoughts on clouds, preservation, and knowledge". Christoph focused on what he called Creative Friction between the important factors that interact. Looking in detail at preservation one needed to address bitstream preservation as well as, for example, the logical layer, semantics, costs, etc.

A useful source of information on this is Open Research Challenges in DP, wiki at http://sokrates.ifs.tuwien.ac.at. Another view was of digital preservation as communication with the future, however with several potential twists: at the time of reception (1) the message may no longer exist (2) there may be no sender (3) there may be no easily available encoder to check against and (4) the recipient may not be the original addressee.

Assuming we have the bits in the future we could treat those bits as a block box, but the key question is whether we can get to the knowledge encoded in those bits. Can we do something useful with them? Can we find different knowledge in them?

We are moving towards knowledge organisations with diverse knowledge and correspondingly diverse needs of representing knowledge. Some key initial questions about organisations preserving this knowledge are: (1) what do they have (2) which capabilities does such an organisation require (3) which services do these capabilities require (4) how can both be measured and (5) what are the cost/ risk and value?

The third and final presentation was by Milena Dobreva and looked into the current context of digitisation and preservation in the memory institutions in the EU, raising the possibility that knowledge re-use may be even more cloudy than the future of digital preservation. The presentation also provided examples on capturing intermediaries' requirements for digital preservation systems.

The results of the ENUMERATE survey for 2012 into digital preservation in the EU produced some fascinating statistics and alarming conclusions. The good news is that there is a lot of activity in digital preservation:

- 83% of institutions have digital collection or currently involved in digitization
- 23% have a written DP strategy
- 33% are included in a national preservation strategy
- 30% are included in a national DP infrastructure

The bad news is that both users and curators have serious issues. Users report that "Archival practitioners" are not disciplined, the terminology they use is not consistent, hierarchical representations are unclear, the existing search tools are inadequate, and content visualization is not uniform. The curators report that they face challenges in the diversity between documents (e.g., fonts, multimedia), diversity of collections (e.g., data size for audio/video), the level of metadata (e.g., how much to include, how to collect (cost-effective), extract, select, and predict it), the absence of linguistic support in discovery and finding aids, workflows, and cost.

After these presentations, we had a lively discussion which started with Michael's question: Is there a conflict in re-use/combination (i.e. generating knowledge) vs preservation? David G had an unequivocal "no." Andre pointed out that preserving data and knowledge is just the first step and gave as an example a recent project that combined weather data with energy generator (solar, wind) data in a new context. Margo added the observation that the earlier you use the data, the more chances you have to improve the accuracy of stored data. Liuba thought that by providing the collective knowledge required for future changes, a cloud could help address the fact we don't know how data will be used in the future. Later on in the session, Ethan expanded on this point describing that the cloud can provide standard migrators.

David Rosenthal raised the concern that since 50% of the cost is paid upfront for the ingest, if one requires "extra stuff" for re-use then cost increases, but budgets are not flexible. Therefore making things re-usable may not be economically feasible. David G argued against this conclusion, because the 50% figure comes from research into rather specialised, small data sets, mostly of documents. When one deals with large volumes of data, with more uniform ingest mechanisms, then the figures will be vastly different

The key points in the remaining discussion were:

- what about preserving non-born digital artifacts?
- where are the representation information repositories and can we depend upon them?
- we need to remember access and not just collecting
- what about cost of computation versus cost of data

3.8 Economics

Ethan Miller (University of California – Santa Cruz, US)

This panel had for presenters: Ethan Miller (University of California – Santa Cruz), David S. H. Rosenthal (Stanford University Libraries), Ross King (Austrian Institute of Technology) and Raivo Ruusalepp (National Library of Estonia).

This panel discussed the economics of long-term preservation, focusing on two major issues. The first two panelists discussed the problem of forecasting the cost of long-term preservation storage and the impact of different factors on this cost. The second two panelists then discussed how these costs might be borne by users and potential funding models for long-term preservation.

The first half of the session featured two talks that described the use of economic modeling to study the costs of long-term storage, particularly as device types and costs shift and systems shift to a cloud-based model. The first panelist, David Rosenthal, examined the costs of long-term storage in the cloud, as exemplified by Amazon and others. He explained that the long-term cost of storage is dominated by the rate at which storage gets cheaper, and that this rate may be slowing. While Amazon and other large cloud providers have a large advantage in that they can "smooth" demand for long term storage across many consumers (organizations), large cloud providers also have a captive market allowing them to somewhat artificially keep prices higher. The alternative—building a private cloud-like system—can provide a return on investment within three years, given current pricing. Ethan Miller, the next panelist, explored trends in long-term storage and discussed the interplay between changes in storage cost and desire for reliability. He noted that, as storage growth rates slow, reliability becomes increasingly important because it is more worthwhile to retain storage for longer. This has a big impact on cloud storage because it can give an advantage to devices such as solid-state storage that are currently unfeasible for long-term preservation. He suggested that architects of cloud preservation systems encourage storage designers to improve reliability, even if it means slightly higher storage costs.

The second half of the session examined the problem of economics from the preservationists' point of view. Ross King discussed the commonly-used "endowment" funding model, and noted that it was similar to a pension or Ponzi scheme in which new users pay the costs for existing data. He expected that collections would grow, allowing new data to dominate old data in size. However, since this growth may slow in the future, he suggested that we need to stop archiving everything and increasingly turn to automated appraisal to limit the amount of data we preserve. In response to a question, he added that preserved data will expand to fill available capacity, further motivating the need for automated appraisal. Other comments included suggestions to better monetize stored data, and to perhaps re-examine archived data after a period of time to see if it's still worth preserving. The final panelist, Raivo Ruusalepp, explored economics from a users' perspective, noting that digital preservation is an unfunded mandate and that few organizations even know their annual unit cost for digital preservation.

Ruusalepp then presented statistics from digital libraries, based on a study available from http://dp4lib.langzeitarchivierung.de/. This study found that staff costs are the dominant factor for the three phases of preservation: ingest, curation, and access. He then described several DCH and GRID initiatives, including a DC-Net project survey (http://www.dc-net.net) that surveyed state-of-the-art digital preservation services in 2012 (survey at http://www.dc-net.org/getFile.php?id=467). He also discussed the Indicate project (http://www.indicate-project.org) and e-Culture Science Gateway (COMETA). He concluded with a discussion of how the cloud fits into the life-cycle costs of a digital object, noting that the cloud makes it difficult to "click and forget" because of the need to monitor cloud providers and pay for storage on a per-gigabyte basis.

The session concluded with questions from other workshop attendees. The first question asked what costs went into the different phases of preservation (ingest, curation, access). David Rosenthal stated that storage costs were only relevant for a few years, and could then be ignored because of the growth in storage density. Michael Factor countered David's argument, saying that ingest, migration, and other factors make the problem worse than presented. Margo Seltzer added that cloud preservation changed the game somewhat because it allowed an organization to pay by the gigabyte-month rather than paying for an archive up front. This was followed by a discussion of who actually pays the costs for preservation services, the user who stores the data, or the one who accesses it? If they are the same person (or organization), this is a moot point, but it's often the case that they are different.

Overall, the session was successful in providing a broad overview of the economics of digital preservation in the cloud, both from the perspective of architects modeling long-term storage costs and from the perspective of users and consumers finding the funding to pay for it.

3.9 Preservation Storage & Cloud Issues

Joanne Syben (Google Inc. – Mountain View, US)

Common sense and intuition suggest that a simple approach of treating the value of data more or less equivalent to its freshness, is a good, default model for how easily accessible the data should be. Older data can automatically migrate to progressively colder and colder storage based on its age. Tradeoffs may be made between durability, availability and reliability. The current 'coldest' storage is tape, however its viability for servicing retrieval of objects from multi Petabyte or even Exabyte volumes is questionable.

The laws surrounding Data Protection Authority for many countries may confound this simple algorithm. If a user wishes to delete her account, the action needs to stretch across possibly many platforms of storage. What is the scope of the data associated with a user given the complex graphs of social media? Should the distribution of user data across different storage media be controlled by the provider or curated by the user? If curated by the user, how can this be made simple and intuitive, as well as practical? There are also many concerns about the heritability of data. Different storage media can be used to reflect the value of data, but establishing what the value is in any systematic way apart from age and possibly 'last time accessed' is an open question.

4 Working Groups

4.1 Automated Appraisal

Michael Factor (IBM Research – Haifa, IL)

The participants in the Automated Appraisal breakout session were: Ian F. Adams, André Brinkmann, Michael Factor, Ross King, Ken Moody, Gillian Oliver, and Joanne Syben.

The discussion started with level setting on the definition, the tools and the use cases. Automated appraisal refers to decisions about what to ingest and what not to ingest, priority, time to live (TTL), and semantic tagging. Tools for automated appraisal include rule-based systems, image analysis (face recognition), (near) de-duplication, natural language processing (NLP) (e.g., cross-referencing), machine learning (approximate the archivist). The use cases include regulation enforcement (e.g., health records), corporate assets, and receiving a box of "disks".

In the breakout session, the team reached the following conclusions:

- automated appraisal is hard for memory institutions, but may be easier for other domains
- it is very domain specific
- it implies new roles for curators and archivists, e.g., define rules, tweak algorithms, train machines

We concluded that a cloud can be quite beneficial for automated appraisal since it involves tasks that are bursty and computationally expensive.

In thinking about automated appraisal we turned to a discussion on the distinction between archival preservation and records management. For instance, is the storage of health records really archival or is it just record management?

Perhaps this distinction is a remnant of the paper world, but is being re-thought. It's not clear – does integrating record management and archival make things easier or harder? Which problems should be solved early in the life cycle and which problems should be solved later in the life cycle. Based upon the discussion, it seems that the answer might be domains specific. Perhaps the answer comes down to life time? Something that has to live "long enough" requires archival "processes" while things that don't have to live "sufficiently long" perhaps don't?

In the discussion with the entire group, David R. suggested that we focus only on born-digital media, so we do not have to continually come back to the decision about the relative costs of digitization and storage. This decision on what to digitize has a significant effect on appraisal.

We then discussed the following hypothesis brought up by Michael Factor: Automated appraisal is a hard problem for memory institutions but may be much simpler for other areas (e.g., health care)? This was controversial. Liuba thought that it was exactly the opposite – it is easier for memory institutions, because they already have procedures. Christoph thought it was just too broad; you can always find an exception. Mary looked at the problem from the aspect of finding the data to appraise, which is the hard problem for large organizations; once the items are found, appraisal is easier. Finally, Raivo argued that the challenge is articulating the appraisal policy. This point was quite controversial! Mary and Michael totally disagreed, claiming that the actual point of difficulty is defining or implementing the procedures.

As a whole this session brought up several points of debate and pointed out that in the digital world there is a lack of clarity on the border between archiving and records management. Further, given the lack of agreement on what constitutes the hard problems, it seems like there is not a shared understanding of the problem definition.

4.2 Design for Forgetting

David Rosenthal (Stanford University Libraries, US)

Main reference P. Druschel, M. Backes, R. Tirtea, "The right to be forgotten – between expectations and practice," ENISA, [Deliverable – 2011-10-18], Nov 20, 2012.
URL https://www.enisa.europa.eu/activities/identity-and-trust/library/deliverables/the-right-to-be-forgotten

The participants in the Design for Forgetting breakout session were: Hillel Kolodner, Liuba Shrira, David Rosenthal, Lawrence You, Margo Seltzer, Mary Baker, Matthias Grawinkel, and Gerhard Schneider.

The discussion identified two reasons for forgetting data in a preservation system:

- The system may have a legal or contractual requirement to forget specific data items. For example, under the EU's "right to be forgotten", all data about an individual must be able to be removed.
- The system may need to forget data that meets specified criteria in order to meet budget or other constraints by reducing storage consumption. For example, all data more than Y years old with importance attribute less than I.

The key difference between them is the level of proof required that the data matching the criteria are gone. If the goal is to reduce resource consumption an aggregate proof, before and after resource consumption, is adequate. Different systems will need to implement different criteria and parameters for the forgetting decision algorithm; it will reduce costs if these are known at ingest time.

Google and others have found that complying with legal mandates to forget is an expensive and technically difficult process. In at least some jurisdictions it involves searching the entire content to identify the removal candidates and then performing a secure deletion process. Thus the content of the preservation system suffers frequent write updates across its entire extent. These jurisdictions apparently do not regard deleting the keys used to encrypt candidates as adequate.

The group then engaged in a lengthy discussion about the difficulty of proving that information has been destroyed. The group proposed a theorem and a corollary:

- Theorem: It is impossible to prove that information destruction has taken place. This may be a consequence of quantum determinism and reversibility, see the Black Hole Information Paradox.
- Corollary: The best achievable goal is "good enough" forgetting, such as "key tossing" – encrypting and discarding the key. Note that this makes recovering the forgotten data expensive not impossible. The cost will decrease through time.

Thus, legal interpretations surrounding "forgetting" need to be revised and clarified. They presently impose requirements that may simply be unimplementable. See for example the ENISA report.

The group was very reluctant to include forgetting as a basic requirement of preservation systems for two main reasons:

- Systems that make forgetting as hard as possible, preferably at least as hard as the printed paper library system, are important to resist censorship and protect society's heritage. See, for example, the US government's attempt to suppress Volume XXVI of Foreign Relations of the United States.
- Adding the capability to forget to a system decreases its ability to fulfill its primary mission, to preserve information. The capability's implementation may have bugs that cause unintended forgetting. Even if it is implemented perfectly, it may permit insider abuse or external attack to cause data loss that in its absence would not have occurred.

4.3 PaaS/SaaS for Data Preservation

Alexandru Iosup (TU Delft, NL)

The participants in the panel PaaS/SaaS for Data Preservation were: Jean Bacon, Christoph Becker, Nikos Chondros, Erik Elmroth, Sam Fineberg, David Giaretta, Alexandru Iosup, Natasa Milic-Frayling, Ethan Miller, Dirk Nitschke, Peter R. Pietzuch, and Raivo Ruusalepp.

How can data preservation become an application domain of cloud computing? Although cloud computing can be generically defined as a useful IT service, which type of IT service would be useful for data preservation (that is, Infrastructure-, Platform-, Software as a Service)? We propose here a set of data preservation services that together form an interface to a Platform as a Service (PaaS) cloud for data preservation. To enable both bit and logical data preservation, our proposed services store multi-layered data at the bit and logical layer. To enable long-term preservation, our services are designed to store not only raw data, but possibly also other elements in the interpretation stack, such as libraries, the OS, and even the machine model. Our approach promises to enable later re-enactment (emulation) of the recorded performance, even for complex multi-media artifacts such as online game playing and spectating. The PaaS we propose includes ten core primitive operations for data preservation, grouped into operations for ingestion, curation, and access; it also includes four primitive operations that are orthogonal to data preservation operations, such as reporting. We have validated our proposed PaaS by mapping five use cases to it: digital photography, a small company preserving financial and business records without the help of an archivist, a public museum running a digital service for the general public, a scientific setting with proprietary lab equipment and data, and digital game and game performance preservation.

5 Overview of Talks

5.1 Scientific Data Archiving Requirements

Ian F. Adams (University of California – Santa Cruz, US)

Joint work of Adams, Ian F. ; Storer, Mark W. ; Miller, Ethan L.;
Main reference I. Adams, M.W. Storer, E.L. Miller, "Analysis of Workload Behavior in Scientific and Historical Long-Term Data Repositories," ACM Trans. on Storage, Vol. 8, No. 2, ACM, 2012.
URL http://dx.doi.org/10.1145/2180905.2180907

Scientific data archiving is one of many important areas within the growing field of archival storage and has several defining characteristics to note. Its data sets may be quite large in private or controlled data, but tends to be more modest in size for public or web accessible data. In all cases however, we find that automated processes such as integrity checking, indexers, and file migration make up the majority of activity. We find that the old adage of "Write-Once, Read-Maybe" archival data should not be relied upon. Updates to data are infrequent compared to enterprise data, but hardly rare. We also found that while individual files are often not appreciably more popular than any other file or record in a corpus, users often show strong locality of access in their activities. This locality may be leveraged to improve the performance and efficiency of both local and remote storage.

From the perspective of utilizing public cloud storage, we have concerns on several fronts. First, large scientific data sets may be extremely expensive to store on a "public" cloud due to their large size. Second, some scientific data may be sensitive in nature and being stored on a shared infrastructure may be risky, particularly as cloud services often have limited, if any, liability. Third, as many providers charge on a per-access basis, we see strong disincentives for useful activities, such as web indexing and remote integrity checking.

5.2 Security technologies for cloud service provision.

Jean Bacon (University of Cambridge, GB)

Individuals and organisations wish to have security guarantees before they store their data on cloud services, short-term or long-term. Ideally, cloud service providers should state clearly what their guarantees are. But current contracts that must be accepted by cloud tenants before using cloud services explicitly avoid cloud-service providers' responsibility for potential failures. Even if this responsibility were included, it could be unenforceable because of the international scope of the jurisdiction.

Security technologies should be used as appropriate for the cloud and updated over time: authentication, access control, encryption, information flow control, data anonymisation and partitioning. I have worked with Ken Moody on role-based access control (RBAC) policy specification and enforcement; with Peter Pietzuch on Information Flow Control (IFC); on health and lifestyle monitoring in the PAL project. The focus has been on systems spanning multiple administrative domains but under a single national jurisdiction i.e., access control policy is specified nationally at a coarse grain and within individual administrative domains for fine-grain local detail. We have not worked under an assumption of international jurisdiction.

Issues from my previous work that are relevant to cloud-service provision:

- Quantification of risk: The consequences of loss or leakage of data should be quantified and used.
- Data: (1) UK cancer records have been gathered by law since 1971. Large fines are imposed if data is leaked. Cloud storage is being considered. A private cloud with known geographical locations is probably most appropriate.
- Data: (2) A large amount of personal health and lifestyle monitoring data is being gathered. How should this be summarised and stored?
- Data of types 1 and 2 tends to be append-mostly except for disambiguation and correction. Read access is needed for (1) long-term statistics about cancer and (2) analysis of personal health.
- Integrity is an aspect of data protection and must be ensured by security technology.
- Audit: Accesses to data should be logged and made available to the owner.
- Dynamic, run-time data-flow monitoring using IFC seems highly appropriate for use in the cloud.

5.3 Challenges for information longevity in the cloud

Christoph Becker (TU Wien, AT)

Searching for creative friction in the intersection of cloud technologies, delivery models, and digital longevity, we can build on the metaphor that digital preservation can be viewed both as interoperability over time and communication with the future. Keeping the bits safely stored is a necessary, but not sufficient condition for preserving information. The preservation field is increasingly moving forward from mere storage and preservation of the bits to an awareness of the value chain of information and knowledge. Preservation in this sense can be seen as the processes required to ensure that an artifact remains connected with the contemporary computing ecosystem. This is much more challenging than intuitively recognised, owing to the "black box phenomenon" of digital artifacts, which only become meaningful through a computed interpretation. While clouds introduce interesting opportunities of scale and flexibility, they also raise questions of control, transparency, and trust. I shortly discuss a few easily overlooked issues surrounding the obsolescence debate and raise a few opportunities and challenges:

1. How can we standardise and automate the underlying key processes of information preservation to leverage the flexible scalability of cloud technologies and use emerging delivery models for addressing the long tails of content artifacts, users, and access environments?
2. What are the basic factors contributing to life expectancy of digital artifacts? Can we predict life expectancy and life cycle costs for digital information in dynamic environments?
3. How can we integrate the concerns of digital longevity and information preservation over time into emerging computing paradigms? How can we establish longevity as a valid design concern in the information systems life cycle from the very beginning?

Taking the example of a specific "knowledge organisation", we see that on many conceptual levels, it does not matter much if we preserve images, documents, research data, health records, or videos – the underlying fundamental computing and communication principles, as well as many of the organisational questions, apply equally across content types and scenarios. A key challenge for both preservation and cloud computing over the next decade will be the question of systematic assessment of information artifacts, processes, systems, and organisational capabilities across scenarios and artifacts.

5.4 Reliability and Integrity of Cloud Storage

Andre Brinkmann (Universität Mainz, DE)

Cheap Cloud storage can become an interesting alternative to in-house archiving and preservation. Nevertheless, trust in the reliability, integrity, and security of Cloud storage still has to be built. Many of the design challenges for Cloud-based preservation have already been investigated in previous work on distributed storage environments such as LOCKSS or OceanStore. The storage should be assumed to be unreliable, intrusion detection should be integrated, and third-party reputation as well as long-term secrets should be avoided. Interestingly, the reliability concerns are only partly due to the quality of the provider's backend storage, but also based on the availability of the Internet connection.

Integrity is one of the key requirements of preservation, but could also conflict with the demand of Cloud providers to implement the storage backend as cheaply as possible. It is therefore important to regularly check the integrity of stored data, as data losses are otherwise silent for the data owner. Unfortunately, Cloud storage is charged based on accesses and transfer volume, therefore scanning the complete archive becomes too expensive. Efficient techniques from secure auditing could be used to uncover larger losses, small incidents are more difficult to find without reading huge parts of the archive.

Ensuring integrity and reliability clearly includes conflicting demands. Lots of copies make stuff expensive and using Cloud storage should help reduce costs. Furthermore, you should not trust a single provider, but unlike in previous work on distributed storage, the number of cloud providers is rather small. The protocols therefore have to work on this small set, requiring us to rethink the underlying assumptions of many distributed storage protocols. Furthermore, techniques such as secure auditing help detect data losses, but do not prevent them. Recovering from data losses therefore requires multiple sites to be involved, and these sites have to be coordinated, at best, based on interfaces agreed to by all Cloud providers.

5.5 Content verification

Nikos Chondros (University of Athens, GR)

The integrity of an archive needs verification to protect against both bit-rot and malicious modification. Due to the latter, even when doing this locally by storing hashes of the content and later verifying them, it necessitates going outside the digital preservation system and asking external nodes to be witnesses to the archive's integrity. For example, these witnesses might store the complete archive, a portion of the complete archive, or perhaps a versioned digest of a hash-tree that summarizes all content checksums. The use of external witness nodes results in the formation of a distributed system, adding to the complexity of the solution. If the digital preservation system is stored in the cloud, the new challenge is whether these external nodes are leased from the same cloud provider or not, questioning the single provider approach. Our current research focuses on implementing an efficient solution based on a persistent hash tree (RBB-Tree, Petros Maniatis), with the aim of minimizing nodes' storage requirements.

5.6 Thoughts on (Preventing) Logical Obsolescence

Michael Factor (IBM Research – Haifa, IL)

URL http://ensure-fp7-plone.fe.up.pt/site

Logical obsolescence occurs when one is no longer able to interpret a digital object. This problem exists whether or not a cloud is part of the preservation infrastructure; however, the use of a cloud can most definitely exacerbate the problem by further separating the data owner from the preservation infrastructure. Based upon experience in the EU funded ENSURE project, there are several points we should consider. First, sometimes it is best to do nothing – what we do needs to depend upon the value of the data as well as the cost of any action. In this context, it is important to define a preservation plan based upon requirements, evaluating cost, risks and value and protecting different data in different ways, including stopping investment in some data after a period of time. Second, don't forget the metadata; it may be the most important thing. We need to leave ourselves *inexpensive* breadcrumbs when the information is stored. We should do this by building on OAIS, e.g., mapping OAIS APIs to objects, e.g., via the Cloud Data Management Interface (CDMI), storing OAIS metadata as object attributes and ensuring we have meaningful, cost effective, metadata for all entities at ingest. And finally, automation of action and verification is essential. In ENSURE, we identified the following as important approaches: 1) use a workflow engine, such as jBPM, to manage flow of all actions, 2), use virtual appliances which encapsulate the software used for the data to provide shorter term preservation and 3) use some form of computational storage, such as storlets, for transformations and quality verifications to support longer term preservation.

References

1 O. Edelstein, M. Factor, R. King, T. Risse, E. Salant and P. Taylor, "Evolving Domains, Problems and Solutions for Long Term Digital Preservation", in *Proceedings iPRES 2011 – 8th International Conference on Preservation of Digital Objects*, Singapore, 2011.

5.7 The Self-contained Information Retention Format (SIRF)

Sam Fineberg (HP Storage CT Office – Fremont, US)

Joint work of Fineberg, Sam;Cohen, Simona;Cummings, Roger
URL https://sites.google.com/site/snialtr/self-contained-informarion-retention-format-sirf

The SNIA Long Term Retention Technical Working Group (LTR TWG) was created to address the "grand technical challenges" of long term digital information retention & preservation, namely both physical ("bit") and logical preservation. A major component of the TWG's Program of Work is the creation of a logical container format, named the Self-contained Information Retention Format (SIRF), for the long-term storage of digital information.

Key aspects of a long term storage container:

- Self-describing – can be interpreted by different systems
- Self-contained – all data needed for the interpretation is in the container
- Extensible – so it can meet future needs

SIRF is a logical data format intended to be the digital equivalent of an archivist's box

- Logical container for a set of (digital) preservation objects and a catalog
- The SIRF catalog contains metadata related to the entire contents of the container as well as to the individual objects
- SIRF standardizes the information in the catalog

The LTR TWG is currently creating bindings of SIRF for Tape (Linear Tape File System) and the cloud (Cloud Data Management Interface). We see SIRF as a key part of a sustainable cloud preservation store.

5.8 Logical obsolescence

Matthias Grawinkel (Universität Mainz, DE)

The term obsolescence opens a diversity of questions on toddy's storage systems. When is data forgotten so that information cannot be accessed anymore?

Data can be deleted on demand, but how do we guarantee that no copies are left? Metadata, links or encryption keys may be lost, so that an orphaned file cannot be interpreted anymore or can only be interpreted at high overhead and cost. While a storage provider can prove that it has stored particular data, can it also prove that it does no longer stores any copies of the data?

Managing expiry dates and encryption keys solve only one part of the problem. Data can be spread to multiple data centers and to different storage media, so that deleting each byte that make up the file is impossible. Instead, encryption keys can be deleted, but legal issues remain.

5.9 Towards Logging and Preserving the Entire History of Distributed Systems

Alexandru Iosup (TU Delft, NL)

The implications of archiving large amounts of daily information for science and society are clear since at least the 1940s, when Vannevar Bush defined the concept of the personal memex as an individual's device for storing and accessing all information and communication involving that individual. Among these benefits are learning about and eradicating humankind diseases, enabling human beings more creative and thought-related time by eliminating tasks that can be automated, etc. Similarly, we posit that archiving large amounts of operational traces collected from the many distributed systems that currently underpin societies across the world would be beneficial for tuning today's systems and designing better systems in the future. What is the Distributed Systems Memex? How can such a Memex be designed and implemented? To address these and related questions, we have taken a bottom-up approach, in which we focus on the archival needs of specific application areas in which distributed systems are prominent and hope to gain sufficient understanding for the future.

In the mid 1990s, the grid computing community promised the "compute power grid," a utility computing infrastructure for scientists and engineers. Since then, a variety of grids have been built worldwide, for academic purposes, specific application domains, and general production work. The Grid Workloads Archive (GWA) [4][1] is a workload data exchange and a meeting point for the grid community. We have defined the requirements for building a workload archive and have described the approach taken to meet these requirements with the GWA. We have introduced a format for sharing grid workload information and the tools associated with this format. Using these tools, we have collected, ingested in the archive, and processed data from over fifteen well-known grid environments, covering thousands of scientists submitting millions of jobs over a period of over twenty operational years, and with working environments spanning hundreds of operational sites comprised of tens of thousands of processing machines.

Resource failures are currently common in distributed systems, likely as a side-effect of the increasing complexity and scale of these systems in real- life deployments. To facilitate the design, validation, and comparison of fault-tolerant models and algorithms, we have created the Failure Trace Archive (FTA) [2][2] as an online public repository of availability traces taken from diverse parallel and distributed systems. We have designed a new data format and a toolbox that facilitates automated analysis of trace data sets. We have collected, ingested in the archive, and processed data from over twenty-five well-known distributed systems, covering millions of users over a period of over twenty operational years, and with working environments spanning the entire world and millions of computers.

Peer-to-Peer (P2P) systems have gained a phenomenal popularity in the past few years; among them, BitTorrent alone serves daily tens of millions of people and generates an important fraction of the Internet traffic. Measurement data collected from real P2P systems are fundamental for gaining solid knowledge of the usage patterns and the characteristics of these systems and can improve the modeling, the design, and the evaluation of P2P

[1] http://gwa.ewi .tudelft.nl/
[2] http://fta.scem.uws.ed u.au/

systems. We have created the P2P Trace Archive (P2PTA) [3][3], an archive that facilitates the collection and exchange of P2P traces. Currently, the P2PTA hosts over twenty traces of various P2P applications (from file-sharing to VoIP to video-streaming), with over 60 million sessions, tens of millions of content items, and multiple years of system operation.

Spurred by the rapid development of the gaming industry and the expansion of Online Meta-Gaming Networks (OMGNs), we have designed the Game Trace Archive (GTA) [1][4] to be a virtual meeting space for the game community. We have proposed a unified format for game traces and introduced a number of tools associated with the format. With these tools, we have collected, processed, and analyzed 9 traces of both games and OMGNs. We have already collected in the GTA traces corresponding to more than 8 million real players and more than 200 million information items, spanning over 14 operational years. We also show that the GTA can be extended to include a variety of real-game trace types.

Our work in designing and building a Distributed Systems Memex has only just begun.

References

1 Yong Guo, Alexandru Iosup: The Game Trace Archive. NetGames 2012: 1-6

2 Derrick Kondo, Bahman Javadi, Alexandru Iosup, Dick H. J. Epema: The Failure Trace Archive: Enabling Comparative Analysis of Failures in Diverse Distributed Systems. CCGRID 2010: 398-407. Best paper award.

3 Boxun Zhang, Alexandru Iosup, Johan Pouwelse, and Dick Epema. 2010. The peer-to-peer trace archive: design and comparative trace analysis. In Proceedings of the ACM CoNEXT Student Workshop (CoNEXT '10 Student Workshop). ACM, New York, NY, USA, Article 21.

4 Alexandru Iosup, Hui Li, Mathieu Jan, Shanny Anoep, Catalin Dumitrescu, Lex Wolters, Dick H. J. Epema: The Grid Workloads Archive. Future Generation Comp. Syst. 24(7): 672-686 (2008)

5.10 Endowment Models and Archival Institutions

Ross King (Austrian Institute of Technology – Wien, AT)

A growing number of archival institutions are turning towards POSE (Pay Once, Store Eternally) or Endowment models for funding their long-term digital archiving and preservation activities. The endowment model has a number of seductive advantages. First, it fits in nicely with project-oriented digitisation efforts, as the endowment costs can be included in a project budget and do not have to be added to annual running budgets. Endowment models also allow simple budget calculations based on total storage volume, which in turn support business models based on archival services. As archival institutions face pressure to become self-sustaining, such business models are in great demand.

However, there may be a number of dangerous assumptions behind simple endowment models. There is a pervading view that data centers with endowment models are like pension plans, in which incoming endowments (workers) will pay for old data (retirees). This analogy is clearly false because, unlike unfortunate pensioners, old data never dies. The reply to this

[3] http://p2pta.ewi.tudelft.nl /

[4] http://gta.st.ewi.tudelf t.nl/

is usually, "but the old data is so much smaller than the new data." This is true, but the only way in which an endowment model can handle the ever-increasing volumes of new data is by basing the cost model on a careful analysis of storage costs, such as detailed in [1]. Too many endowment models simply assume that per volume storage costs will continue to decrease (Kryder's law) forever, which simply cannot be the case. Rather, the storage capacity per unit cost, which is presently in an exponential growth phase, will eventually reach a stationary phase and level off, just as every other exponential growth scenario in nature. It is incumbent upon an endowment model to at least attempt to predict when this stationary phase will occur and at what rate storage capacity will continue to grow.

Commercial storage providers such as Google and Amazon are well-aware of these difficulties, and offer business models that are highly advantageous to themselves as a result (the decreases in service costs offered over the past five years are still much higher than the real decrease in storage costs). The question is, are libraries, archives, and other data centers equally aware? Endowment models are complex and probably more expensive than we think. The inevitable conclusion is that we can no longer afford to archive everything.

References

1 David S.H. Rosenthal, Daniel Rosenthal, Ethan L. Miller, Ian Adams, Mark W. Storer, Erez Zadok, "The Economics of Long-Term Digital Storage", *The Memory of the World in the Digital Age: Digitization and Preservation*, September 2012.

5.11 Implementing a Preservation System over a Storage Cloud

Hillel Kolodner (IBM Research – Haifa, IL)

There are many issues and trade offs to be considered in building preservation systems over storage clouds. Typically a storage cloud provides high levels of security and reliability, and in addition, advanced features to differentiate itself from competitors. Yet, paradoxically these features may be overkill for a preservation system. In particular, a preservation system cannot incur the risk of depending on a single cloud provider, who may go out of business, and also may not trust the provider to keeps its content secure. So, given that the preservation system is itself built over multiple, possibly untrustworthy clouds, it must take on the responsibility for high reliability and end-to-end security and can likely make due with clouds that provide a lower level of reliability and security. This is similar to the way that cloud systems, themselves, are typically built, where it does not pay to invest in high reliability at the low level, e.g., in hardware, since the high level cloud software needs to be able to deal with failure in any case.

On the other hand cloud storage systems are providing increasingly more sophisticated features that can simplify the task of building a preservation system. If preservation systems are built over multiple clouds, how will they leverage these features? For example, consider some of the advanced features provide by VISION Cloud, an FP7 project: support for rich metadata, computation in the storage system and support for the secure handling of data.

Rich metadata is supported as an integral part of an object by the storage system. The storage system is responsible for the integrity of the metadata as well as the data. Whereas object data cannot be updated in place, metadata can be updated. Furthermore, objects can be found/searched based on the values of their metadata fields.

The storage system provides a way to run computations in a safe and secure way. This can be more efficient as it can avoid network bandwidth. Furthermore, it can be more secure since the data does not need leave the storage system.

Support for the secure handling of data includes secure isolation of tenant/user data, encryption of data, and geographic constraints on the placement of data.

References

1 E. K. Kolodner, S. Tal, D. Kyriazis, D. Naor, M. Allalouf, L. Bonelli, P. Brand, A. Eckert, E. Elmroth, S. V. Gogouvitis, F. Harnik, D.and Hernandez, M. C. Jaeger, E. B. Lakew, J. M. Lopez, M. Lorenz, A. Messina, A. Shulman-Peleg, R. Talyansky, A. Voulodimos, and Y. Wolfsthal. A cloud environment for data-intensive storage services. In *Proceedings of the 2011 IEEE Third International Conference on Cloud Computing Technology and Science*, CLOUDCOM '11, pages 357–366, Washington, DC, USA, 2011. IEEE Computer Society.

5.12 The Life of Digital and the Cloud

Natasa Milic-Frayling (Microsoft Research UK – Cambridge, GB)

URL http://www.scape-project.eu/

Digital obsolescence is an ecosystem problem. Finding a solution requires understanding the nature of digital technologies and the structure of the Computing Technology Ecosystem (CTE) that has emerged to support the creation, application, and sustainability of digital technologies as well as the production and the reuse of digital assets.

Products of digital technologies are digital artifacts that are created through computation and can be consumed in its digital form only when the computation can be executed. A digital artifact cannot be stored – it is, at the most basic level, a sensory experience that is enabled by computation through adequate electronic hardware and lasts only while the computation lasts.

One important aspect of digital is reuse. We persist both the computer programs and the results of computation so that they can be re-instantiated again. We ensure that we have all the programs required to realize the digital artifact in the form we can consume. In many instances that involves a number of applications: the application for writing software, i.e., the original program, the application that can instantiate the results of the program, i.e., data files or document files. Finally, we need the complete stack of software to run the program as well as the hardware that enables us to sense, i.e., perceive, the digital artifact through viewing, hearing, or through our tactile sensors.

Preserving data files and program files is therefore not sufficient for preserving our digital assets. Persisted encoding of programs and data is necessary but far from sufficient. The existence of digital is about computation. Only through computation can it be experienced.

Thus, we define digital preservation as enabling digital artifacts to be instantiated in the contemporary computing environment. That may happen in different ways, but the key is always computation. We can create a virtual machine (VM), emulating the hardware and providing the required software stack. Alternatively we can port the software to the new environment and utilize the data files as they are. Finally, we can identify a contemporary application with required functionality and transform the data file format into the format that can be consumed by that application. This approach is referred to as format migration.

The cloud naturally arises as an environment in which access to digital content can be enabled over a long period of time. First, it brings together the persistent aspects of digital assets, i.e., data and program file storage, and the computation. To illustrate that point, we have created, as part of the SCAPE project, an Azure based service for format migration that is extendable in both directions – the data storage and transformation of formats. The key in ensuring access to digital artifacts is to provide an efficient software development layer that enables developers to create 'bridging' components. Bridging software can, for example, be a format translator that converts an old format to a contemporary format. It can be a virtual machine that enables hosting of a program and computing on the data files. If the results of computation need to be reused outside the virtual machine, then we need either format translators that run within VM or software to capture digital artifacts directly from the presentation layer and then convert them into a form that can be used with other applications outside the VM.

The cloud paradigm may ease the pain of digital obsolescence but only if we put in place the standards that ensure interoperability among cloud platforms and demand that they are designed with the longevity of digital in mind.

5.13 The Economics of Devices Over the Long Term

Ethan Miller (University of California – Santa Cruz, US)

Joint work of Miller, Ethan L.; Rosenthal, Daniel; Rosenthal, David S. H.; Adams, Ian F.; Zadok, Erez

The economics of long-term storage are different from those of "normal" storage: in long-lived systems, different metrics become important. The cost of long-term storage depends on complex, time-dependent interactions between metrics such as device lifetime, density growth rates, and long-term capital expenses, as well as costs of switching to media.

We are investigating the impact of several factors on the cost of providing long-term storage, both for single organizations and for cloud providers, since these costs must be paid either initially when content is first stored or over time to maintain content. For example, we find that the historical rapid growth in storage density makes it worthwhile to replace devices before they fail; as density growth rates slow, it becomes worthwhile to make devices more reliable. Similarly, keeping devices for longer reduces data migration costs and integrity verification costs, further motivating more reliable devices. Further, by moving preservation storage into a cloud environment, we can better spread capital costs over time and provide better device and software diversity.

5.14 Domain Specific Needs: Archives are becoming mission critical

Dirk Nitschke (Oracle – Herndon, US)

The use of archives is changing. A common perception is that (classical) archives can be *slow* and it's not a problem when they are *closed.* However, digital information changes the game in several ways:

More and more companies provide 24 × 7 services and create centralized archives instead of multiple isolated solutions. Consequently, multiple departments or applications depend on the centralized archive. This implies that a digital archive must always be up and running.

Companies store their most valuable assets in these archives. The digital assets are created by a variety of different applications, and very often they are the source material for other products. Data loss is not an option, and the archive must integrate into your enterprise applications.

The end user expectation in regards of access latency changes. Today, we are all used to getting answers to our questions instantaneously from our favorite Internet search engine, and we can download data at a reasonable bandwidth. The same is expected from archives today.

Data is ingested into and retrieved from company archives by end users, not by specialized personal, so they must be easy to use.

Nobody claims that this is easy to accomplish, nor is there a one size fits all solution.

If you are involved in an archiving project, make sure to talk to the right people. Archiving is an organizational task and not an IT task. Typically, the IT department has no knowledge about the data that they have to store. Define who wants to archive what, why, and how you intend to re-use the data in the future. Afterwards, do a data classification. Find out what you have, evaluate the value of your data over time and make some decisions. Making decisions can be the hardest part because you have to decide what *not* to archive.

Then you should think about your time scale, how long to keep the data, the projected amount of data and objects and your migration strategy.

Last but not least, select the right tools and do not over-engineer. Someone has to run the system for a long period of time.

5.15 Cloud computing and future memory

Gillian Oliver (Victoria University – Wellington, NZ)

The objectives of cultural heritage institutions are concerned with individual, organisational and societal memory, but it is important to note that memory consists of remembering and forgetting, plus has long term and short term dimensions. Accordingly, Information professionals in workplaces and memory institutions have developed tools and techniques to prioritise and target preservation actions (including destruction). The advent of cloud computing, and the consequent outsourcing of memory, poses significant challenges to our future memory. Future memory could be a massive accumulation of digital information sludge, with remnants of what should be forgotten remaining, and only occasional glimpses of the important and truly significant. Consequently archival authorities around the world have identified risks associated with cloud computing and are actively working to raise awareness of the issues involved.

5.16 Long-Term Cloud Storage needs Data-Centric Security

Peter R. Pietzuch (Imperial College London, GB)

Joint work of Pietzuch, Peter R.; Bacon, Jean; Moody, Ken

Security considerations are a major issue holding back the widespread adoption of cloud storage: many organisations are concerned about the confidentiality and integrity of their users' data when hosted in third-party public clouds. Today's cloud storage providers struggle to give strong security guarantees that user data belonging to cloud tenants will be protected "end-to-end", i.e. across its entire life cycle. Therefore security engineering must be integrated with all stages of data storage in clouds. We want cloud providers to isolate data of each of their clients. This is crucial for cloud infrastructures, in which the stored data have different owners whose interests are not aligned (and may even be in competition).

We propose a principled approach to designing and deploying *end-to-end secure data storage* in the cloud by means of thorough tagging of the security meaning of data, analogous to what is already done for data types [1]. The aim is that such a *data-centric* security approach using Information Flow Control (IFC) techniques can ensure that—above a small trusted code base—data cannot be leaked by buggy or malicious software. End-to-end information flow control thus preempts worries about security and privacy violations. The cloud storage infrastructure enforces data flow policies through multiple layers of security mechanisms following a *defense-in-depth* strategy: based on policies, it creates *data compartments* that isolate user data. A small privileged kernel, which is part of the cloud infrastructure, constitutes a trusted computing base (TCB), and tracks the flow of data between compartments, preventing data flows that would violate policies. Due to its minimal size and reliance on hardware protection mechanisms, such an approach can strengthen the security of a cloud storage infrastructure against internal and external security attacks.

References

1 Jean Bacon, David Evans, David M. Eyers, Matteo Migliavacca, Peter Pietzuch, and Brian Shand, "Enforcing End-to-end Application Security in the Cloud", ACM/IFIP/USENIX 11th International Middleware Conference (Middleware'10), Bangalore, India, November 2010.

5.17 Using Provenance to Protect Data

Margo Seltzer (Harvard University, US)

When we discuss protecting the data, we typically make the assumption that we are protecting the data from adversarial attack and that the only thing of import is the "integrity" of the data, for some definition of integrity. However, I believe that the question is broader. There are (at least) three different constituencies to whom we can offer protection: the owner of the data, the user of the data, and the provider of the data. Provenance or lineage, which is the complete history of how the data came to be in its current form and at its particular location, is critical for all parties.

For the data provider, data provenance documents the authenticity of the data and ensures that the provider has the appropriate rights to store and serve the data. For the user, authenticity is a key concern, but so too are details of the processing and transformation that have been applied to the data. A user may want to know specific versions of software used to analyze data or the particular system on which a data set was produced. Finally, a data owner uses provenance to establish or maintain reputation as well as to track the data as it is disseminated. All three constituencies have a vested interest in maintaining complete and reliable provenance.

Maintaining reliable provenance requires cooperation between all the systems that participate in data creation, transmission, and storage. While such cooperation would seem to imply that we are paralyzed without well-established standards, I claim that we cannot let ourselves be paralyzed. Data is being produced, manipulated, transmitted, stored, copied, transformed, and destroyed constantly. We need to start documenting that now – we should accept the fact that different data will have different provenance and establish simple ways to communicate the provenance from users to systems, systems to systems, and systems to users.

From an archive's point of view, it must be able to accept provenance from an external source, add provenance to record all archival acts, integrate seamlessly with external and internal sources, and not require that all systems agree on a single, standard format/representation.

5.18 Modularity and incrementalilty in handling logical obsolescence

Liuba Shrira (Brandeis Univ. Waltham, US)

The software stack in a cloud-based preservation service will have different layers. Each layer will need to handle the logical obsolescence problem. We hypothesize that similar properties will be desirable in different layers, ideally allowing us to use similar techniques to achieve them.

With this goal in mind, this talk puts forward an abstract framework for handling logical obsolescence for long-lived interlinked stateful objects described by a schema, borrowed from our work on a system for automatic data store object upgrades. In such a system objects may need to be migrated in two directions, from older version to newer version to allow us to ask new questions about old data, and from newer version to old version to allow us to ask old questions about new data.

We consider two desirable upgrade properties, modularity, which makes it easy to write code that automatically transforms objects from one version to another, and incrementality, which enables low-cost on-demand object migration between versions, and discuss the difficulties of achieving them, hoping to examine with the group how these concerns apply to the obsolescence handling across the different layers.

Participants

- Ian F. Adams
University of California – Santa Cruz, US
- Jean Bacon
University of Cambridge, GB
- Mary Baker
HP Labs – Palo Alto, US
- Christoph Becker
TU Wien, AT
- André Brinkmann
Universität Mainz, DE
- Nikos Chondros
University of Athens, GR
- Milena Dobreva
University of Malta, MT
- Erik Elmroth
University of Umeå, SE
- Michael Factor
IBM – Haifa, IL
- Sam Fineberg
HP Storage CT Office – Fremont, US
- David Giaretta
APA, Dorset, GB
- Matthias Grawinkel
Universität Mainz, DE
- Alexandru Iosup
TU Delft, NL
- Ross King
Austrian Institute of Technology – Wien, AT
- Hillel Kolodner
IBM – Haifa, IL
- Ewnetu Bayuh Lakew
University of Umeå, SE
- Natasa Milic-Frayling
Microsoft Research UK – Cambridge, GB
- Ethan Miller
University of California – Santa Cruz, US
- Dirk Nitschke
Oracle – Herndon, US
- Gillian Oliver
Victoria Univ. – Wellington, NZ
- Peter R. Pietzuch
Imperial College London, GB
- David S. H. Rosenthal
Stanford University Libraries, US
- Raivo Ruusalepp
National Library of Estonia – Tallinn, EE
- Gerhard Schneider
Universität Freiburg, DE
- Margo Seltzer
Harvard University, US
- Liuba Shrira
Brandeis Univ. Waltham, US
- Joanne Syben
Google Inc. – Mountain View, US
- Lawrence You
Google Inc. – Mountain View, US