Designing Storage Architectures for Digital Preservation
Churchill Hotel, Washington, DC
September 22-23, 2009

The following provides brief highlights of these sessions and some of the participant comments for the convenience of interested parties. Full presentations are available on this site.

Overview: The purpose of the meeting was to bring together technical industry experts, IT professionals, digital collections and strategic planning staff, government specialists with an interest in preservation and recognized authorities and practitioners of Digital Preservation to identify common areas of interest to inform decision-making in the future. The agenda, a background reading list and copies of all presentations are available at http://www.digitalpreservation.gov/news/events/other_meetings/storage09/index.html

Action Items:
Library of Congress will:

- Make presentations available on the web.
- Compile notes.
- Work with others who are interested in developing common vocabularies in these areas (such as the Storage Networking Industry Association (SNIA) to see if we can talk about vocabularies across the community.

Meeting introduction: Martha Anderson, the Director of Program Management in the National Digital Information Infrastructure and Preservation Program (NDIIPP) at the Library of Congress (LC), set the tone for the meeting with a brief introduction, noting that a main goal of the meeting was to set up a conversation between those who take care of digital content and those who provide the platforms and tools to support the work.

Tuesday September 22, 2009

Storage for Digital Preservation: Panel of Case Studies (Module 1)

Library of Congress: Overview and Framing: Thomas Youkel, Library of Congress
Overall challenges facing the Library over the coming years:

- Data Integrity - The Library's main challenge is assuring the integrity of items on both disk and tape. How can the Library be sure that what it received is what it still has?
- Workflow - Automating the process from ingest through preservation.
- Content management - A challenge, and there are many unknowns on what will really be needed over time.
- Migration - Especially technical migration. Somewhere in the next 5-7 years, almost every storage job could be a technical migration job. The Library acknowledges that the technology is going to change, and technical migration will be inevitable. In ten years the Library will have 50 – 100 PB, and eventually it will need to rewrite every tape.

Chronopolis: Present and Future Storage Environments: David Minor, San Diego Supercomputer Center
The Chronopolis project is in its second year of existence. It is a three-node federated data grid involving connections at UCSD, UMIACS, and NCAR. There is 50 terabytes of capacity per site, and the infrastructure is based on the Storage Resource Broker (SRB) tool. (http://www.sdsc.edu/srb/index.php/Main_Page). Chronopolis adds layers on top of SRB for verification, replication monitoring and integrity checking with the Audit Control Environment tool

(https://wiki.umiacs.umd.edu/adapt/index.php/Ace). The data comes in to SDSC and is then copied out to the other sites, ultimately put into the tape archive at UCSD. Chronopolis currently manages four collections including data from the Scripps Research Institute, the California Digital Library, ICPSR and geospatial data from the North Carolina NCGDAP project.

MetaArchive and Cloud Computing: Bill Robbins, MetaArchive Cooperative
MetaArchive is exploring cloud computing within the confines of a private LOCKSS network.
Their content is on multiple geographic caches, linked together by the MetaArchive Properties Server.
MetaArchive is exploring the use of a cloud environment to store and manage the common Properties Server function.

Cloud computing is a chance to take open source and move it to an easier level to work with. You can build an entire package in the cloud (database, web server, etc.) that is ready to run right out of the box.

Requirements for a Digital Motion Picture Preservation System: Andy Maltz, Academy of Motion Picture Arts and Sciences
The Academy report, the Digital Dilemma (http://www.oscars.org/science-technology/council/projects/digitaldilemma/index.html), has been widely disseminated (including translations into three languages).
What were the requirements identified in The Digital Dilemma?:

- Access should be guaranteed for 100 years
- Assets will survive periods of benign neglect
- Any digital system should at least be as good as the analog systems it replaces (picture quality, etc.)

The Academy is involved in an NDIIPP-supported project (http://www.digitalpreservation.gov/partners/ampas/ampas.html). The 5 point program of NDIIPP project:

- Digital preservation case study: StEM-standardized test and evaluation material
- Report: what about the independent filmmakers and small archives?
- File formats & metadata
- Research - distributed storage, low cost/low power storage, data compression
- Education and Outreach

The project is about generating an interim digital preservation program. It is looking at bit- and essence-level preservation.

A Storage and Content Management Use Case: Laura Graham
Web archiving began at the Library of Congress in 2000, originally done by the Internet Archive on contract. The Library uses two Internet Archive tools: the heritrix tool for web crawling and the wayback machine for access. Heritrix writes 2 file formats: ARCs and WARCs. WARC is a recent revision of ARC, implements several record types that are important for digital preservation, and it is also an ISO standard (http://www.digitalpreservation.gov/formats/fdd/fdd000236.shtm

The Library uses the Bagit content transfer specification to move files within the Library and from external partners such as the Internet Archive.
(http://www.digitalpreservation.gov/library/resources/tools/docs/bagitspec.pdf)

Requirements for tools:

- Easy and efficient interoperability between machines
- Automation of copying of content

- Reduction in requirements for copying of content.
- Future statistical auditing of content; should not need to verify items at every step.
- Universal tracking through the workflow - tracking everywhere the content might be, not just in storage at rest.

HathiTrust: Building the Universal Collection: John Wilkin and Cory Snavely, HathiTrust
The mission of the HathiTrust is as an enterprise devoted to preserving digital copies of the public record. There are currently 25 partners, including major academic institutions. This was originally developed out of activities surrounding the Google and IA digitization projects. Based on the OAIS reference model.

The Trust is currently at about 4 million volumes, with a big growth trajectory, only limited by the speed of digitization workflows. They should have 5-6 million volumes by the end of the year, 10 million next year. They see content type growth in both retrospective conversion and born-digital publishing.

In terms of technical capacity, they needed to have systems that would operate at scale well, but still be able to run on the margins otherwise the cost of operations can easily get out of hand. Archives that operate at this kind of scale need to be able to trust storage more, with integrity checking, virtualization, levels of abstraction, and the ability to do bit-level preservation within the storage system itself. The ingest processes they are doing now are more highly automated than anything they've done before.

DuraCloud: Managing durable data in the cloud: Jane Mandelbaum, Library of Congress
Duraspace is a collaboration between dSpace and Fedora. The Duracloud project (http://www.digitalpreservation.gov/partners/duracloud/duracloud.html) is an NDIIPP-supported effort to explore cloud computing, an elastic web-based infrastructure for storage and computing. Duracloud will act as an intermediary between libraries and archives and the cloud computing providers, hoping to bring digital preservation efforts into the cloud.

The advantage of cloud services:
- Flexibility
- Scalability
- Pay for use
- Easy to implement
- Cost

Issues:
- Stability
- Transparency
- Data lock in
- Service Level Agreements
- Trust

The project will explore replicating pilot data to various storage providers in different geographic areas and then monitoring and auditing the process.

Petabyte-Scale Processing with Hadoop: An Open Source Perspective: Jimmy Lin, University of Maryland and Cloudera
Lin is on the faculty at the University of Maryland, and is heading an IBM Google Cloud project at the university to explore cloud computing issues. Cloudera is a private company that hopes to provide commercial support for Hadoop. Hadoop is an open source implementation of MapReduce, a tool designed for deep analysis and transformation of very large data sets.

The open source framework underlying Hadoop allows users to leverage commodity hardware, share nothing architecture, move code and not data. The ecosystem also includes other tools such as Hive, Pig Hbase, and Zookeeper. Yahoo and Facebook both run Hadoop to manage terabytes of data per day. Lin proposed that a Hadoop cluster will run about 1/3 of the cost of other systems with more computing capacity.

Storage Products and Future Trends: Vendor Perspectives (Module 2)

Digital Library, Repository, and Preservation Trends: Art Pasquinelli, Sun Microsystems, Inc.
Pasquinelli is leading the Preservation and Archive Special Interest Group (PASIG).
The academic community is now paying more attention to the Storage Networking Industry Association (SNIA) (http://www.snia.org/home), and SNIA paying more attention to PASIG members. Both sides are getting tied into open standards.. Open computing and open storage are requirements for this community.

Everyone is now requiring upfront data archiving plans for grants. Research and content are morphing the IT focus at universities. People mainly go to the Library to ask for curatorial help because they are not getting storage guidance from IT. Are the IT departments impediments? Many are starting to turn to cloud computing.
The data long tail is important: you don't know what is going to be useful in the future.

Storage Products and Future Trends - IBM: Michael R. Mott, IBM
In storage, people are now talking about "cores" and not "frequencies." To raise the frequency you have to raise the voltage, and we lost. The price/performance curve is over. (Moore's law is over).

Architecture paradigm is now more like a "river." The time component of what you do today, the river shapes the landscape and the land tells the river how to move. We don't know what the river is going to look like in 20 years. We worry about the properties we want of preservation over time, and that will be the property of our system over time.

Hard Drive Directions: Dave Anderson, Seagate
Where is capacity going?

Key areas:

- Capacity growth on disk drives may not increase at the same rate, due to hardware factors.
- Solid state pricing vs. power issues.
- Interfaces
- Expected form factor changes.
- Impact of consumer PC market on availability and type of drives available for the enterprise market.

Storage Products and Future Trends -EMC: Timothy K. Harder, EMC
Over the past three years they've heard a change in "asks." Organizations have been asking for storage and computing units to be located in different locations, but to be seen as a single resource, and to do it in a price band where EMC has not traditionally been active. We are looking at commodity components and creating storage arrays out of commodity components.
Want to enable end-user self-provision, get smarter about metering and billing and charge-back data.

Storage Products and Future Trends -Isilon: Paul Rutherford, Isilon Systems
Storage market trends, where are they going? Flash is expensive, but if it helps the business move faster people will pay for it. SATA as an interface goes away as the need for block-level access goes

away.
File System Trends. Scalability: can't run on a standard UFS file system and scale it up. Single namespaces are about simplifying the world you work in.

Storage Products and Future Trends – Cisco: Kevin Ryan, Cisco
In the evolution of the data center, virtualization is the next proof point. We can't do automation without virtualization. Beyond virtualizing the server, virtualize the network and the storage as well. Data Center Evolution: A unified fabric, with single or dual-connections coming out of each server. There is nothing proprietary about a “unified fabric.”

Storage Products and Future Trends – Sun: Raymond A. Clarke (Sun Microsystems)
We are moving to a utility-type model (pay as you go) with different classes of services and business models that include both public and private domains.

There is some information you should have more control over, while providing public access for less sensitive data. Platform as a service and infrastructure as a service: both starting to see some acceptance within the industry.
SNIA is promoting the adoption of cloud storage through the Cloud Storage Initiative (http://www.snia.org/forums/csi/) and the Cloud Storage Reference Model (http://community.stortoc.org/index.php?autocom=downloads&showfile=10).

Moderator John Warren suggested that each participant post a response to one of three questions. The responses would be compiled and shared the next day:

- What's something that needs to be focused on (R&D) in terms of long-term preservation? Something not done now that should be considered?
- Is there anything that we've not heard today? Is there a missing set of people? Missing organization? A missing stream of discussion? Missing ____?
- What's going to happen in 2015 that's not here now?

Designing Storage Architectures for Digital Preservation

Wednesday September 23, 2009

Data Integrity: How Do We Do This? (Module 3)

Data Integrity: Framing the Session: Thomas Youkel, Library of Congress
Everything we do is about data integrity. What are the inherent imperfections of a system – engineered or human—and what are we going to do to mitigate those imperfections.

Data Integrity: Henry Newman, Instrumental, Inc.
Using the cloud for preservation is the antithesis of what's happening in the larger community/market. Tape storage is crucial to the digital preservation market because of cost and power issues. Vendors won't make products that can only be sold in small market, and preservation archives need extreme reliability not required in most markets. Disk density has increased but reliability has not. We've gone from 9 GB drives to 600 GB drives, but haven't increased reliability.

There is a mismatch between preservation requirements and market trends. Some of the requirements, such as the per-file checksum, hardware cluster approach, trade cost of power for cheap hardware, throw a lot of power and hardware at things. Take cheap stuff and replicate it. If you want complete reliability for 50 PB, what would that require?

Integrity in the Cloud: David Rosenthal, LOCKSS
The idea of "magic media" is a myth. You can't trust the storage layer more. As Reagan said: trust but verify. This presentation is about the verification.

Clouds don't guarantee to keep data. Cloud providers make money on the movement of data, which means that preservation will always be a small user in the market, which means that cloud technology will never be our answer. Availability is the proportion of requests that get an answer. Reliability is the proportion of requests that get the right answer. The cloud provider says that's your problem. Each copy in the cloud will be unreliable. Preservation needs copies in multiple storage providers, needs systems that detect and repair damage to each copy and must be able to do the audits in the provider's compute service. But we don't trust the providers and perhaps we shouldn't: the providers try to cover up data loss. 3rd party auditing is necessary!

Data Integrity: Raymond A. Clarke, Sun Microsystems
There is a difference between "back up" and "archive," and they involve very different practices, especially regarding the kinds of techniques used. Man gets increasingly loquacious throughout history – we have more to archive and preserve than ever before. Most data remains untouched – very little is used regularly – lots of data is not used or rarely used. Tape continues to make sense; among all the other issues the energy needed to run tape is lower. Need to put into place a system that helps prevents human error: migrate automatically, policies that drive the systems, keep data fool proof.

Data Integrity: Michael R. Mott, IBM
We're going to lose stuff. We need to determine what is an acceptable level of loss for utility for different formats and media. How much can we lose and still keep the utility? You can lose 1 bit and still be able to see the movie. But if you lose half you won't. What is the acceptable number?

Data Integrity: Timothy K. Harder, EMC
The cloud is not useful for heritage preservation copies, but we are solving that problem with existing technologies today. For 281 exabytes of data, if you drop 1 "nine" in the Bit Error Rate it's 281TB of

data. Sampling has its challenges, and you need to add proof of retrievability to sampling. Can do this by adding additional encoding information natively to file itself which will make files resilient to small errors. Limit amount of data that needs to be moved.

Trust but validate: bring it through and advance it through a community of peers. Don't just trust the back-end provider but also use it as a validator for the backend. There will be other monetization techniques around this – not all data assets will have the same useful lifetime costs and those could be isolated and applied through metadata tags.

Data Integrity: Paul Rutherford, Isilon Systems
In real life there are disk drive failures, controller node failures and user errors. The lesson is don't trust anything: verify data at all layers. Don't trust yourself either. Do cross node protection with multiple copies and isolate failures to individual nodes. A cloud site is a node. What happens if that site goes down? You don't want to be the largest consumer of storage in a cloud provider. You need to segment the data and isolate it into failure units. Time to recovery is the most important thing to ask your storage providers. Reliability of disk in field can be anything from vendor-provided numbers to upwards of 3% failure in the field, which means a mean time to data loss of about 10 years Data-loss estimate is complex. RAID is dead. Have to get something different to go forward such as cross-node data striping. Have three storage nodes and add a layer across that spans all three.

Moving and Distributing Data (Module 4)

Typical Workflow: Thomas Youkel, Library of Congress
Includes the producer, ingest, process, presentation copy and preservation copy. Library is looking at all the movement and the potential for data loss.

Moving Distributing Data: Brian Bailey, Signiant
Signiant is a spin-out of Nortel that does software that manages the transfer of files. Allows you to aggregate and distribute media assets, and you can use Signiant to process files in between. Signiant automates the movement of digital assets into the Library, according to policies. You set the policies once, and then automate the process. Then overwrite the policy if necessary. You can automate according to a schedule or an event. Policies include bagging, re-bagging, QC, registering with a DAM system, others. Simple user interface to get the files in but with the same user controls after it arrives.

Moving and Distributing Data: Michael R. Mott, IBM
How will societal change impact access and distribution? There is a new collaborative generation coming. How are we to deal with rich media documents?

Moving and Distributing Data: Raymond A. Clarke, Sun Microsystems
SNIA DMF 100 year archive task force. XAM-extensible access method API, normalizes the semantics across a variety of applications and storage subsystems. Self-contained information retention format (SIRF). Open Solaris -->ZFS (Hybrid Storage Pool), a new kind of file system than enables the movement of data, etc.

Moving and Distributing Data: Kevin Ryan, Cisco
How long does it take you to transfer a petabyte of information? Load up the tractor trailer.

Technical Migration (Module 5)

Technology Migration: Paul Rutherford, Isilon Systems
If I add new technology, how do I get old technology out? Similar to node failure.

Technology Migration: Dave Anderson, Seagate
Seagate did a study that posited an environment where they could solve those problems to create longer-lived disk drives. This is technically feasible, but the capacity of the media is going to change, and too much of the marketplace gets too much benefit out of regular transitions. For example, if you had purchased a 100-year 30 GB in drive in 2000, you'd need 300 drives to equal what you could put on a 1 TB drive in 2010.

Technology Migration: Mike Mott , IBM
Talking about some of the system concerns arising from the long-term preservation of digital information. You need to keep at least 2 generations of technology and overprovision for migration.

Technology Migration: Raymond A. Clarke, Sun Microsystems
SNIA 100 year task force addressed requirements for information over the long-term, looking at both logical and physical migration. Physical migration is moving from one generation of technology to the next. Logical is how we will understand the data. What does long-term mean, in terms of physical technology? Requirements for the data layer format: media agnostic, self-describing, support for self-contained data.

Closing Comments: Martha Anderson, Library of Congress
Our strategies are based on our experiences with our technical environments, and there's usually something that we've learned that influences our thinking about our practice. This data is not as inactive as you might think. We are a knowledge-seeking community, and we appreciate the kind of architectural view that looks at the relationship between the community and technical environments. . "The river shapes the land and the land influences the flow of the river."

BRAINSTORMING SESSION – DAY 1

What's Missing

- The folks involved in making bandwidth do more.
- Standards (universal application)
- Single biggest limiter is power to large scale repository management. What is being done to address efficiencies at the processor, etc. layers to address this above & beyond this?
- Software standards bodies
- Camera manufacturers-to explain why they won't standardize raw file formats
- Microsoft
- DARPA
- High performance computing community
- Need better interfaces between standards organizations, communities of practice
- How do we persistently associate necessary semantic description to stored bits over archival time spans?
- How can we better engage potential LC stakeholders to help identify the bits we want to save ex ante? Pre-appraisal?
- Consumer digital data preservation will be important when they figure out there's no "shoebox of pictures and negatives in the attic"
- Open source toolkit to transform data types
- Define challenges more often
- How fast will our storage needs grow in 100 years (given our poor predictions so far)
- Self-healing file formats
- Federal storage standards
- Cloud computing and the developing countries
- Add health sector—MRI, imaging, patient records, etc.
- Media production & distribution standards bodies
- How to do reasonable-time backups and data migrations to new storage media for large amounts of data when needed
- Critical mass of lead institutions: Dutch KB (National Library), British Library. The BNF
- Missing people in charge of moving data between systems (e.g., networking)
- Preservation requirement definitions
- Policy-based preservation-how enforce compliance with assessment criteria
- Need to hear from the medical records/biomedical community
- Case studies from libraries that <u>have</u> successful storage systems
- "dynamic metadata with static" digital preservation files
- differentiate between bit preservation and information (content) preservation
- DOJ
- DHS
- A discussion of the relative trends of data storage and movement. Too much to move? Copy? Use?
- Scenarios to guide decisions
- Standards end to end
- How are we training or re-skilling workforce to support & provide innovative solution in the field of digital preservation
- Economics of technology; life cycles
- What do we need to do to make migration work?
- Management framework
- What's wrong with relying on the cloud?
- Selection criteria expiration date on data

- Historian/anthropologist or other academics to explain the value of preservation to the historian of the future. What could be most important to preserve.

R&D – Where should this be focused?

- permanent storage system: hardware, software system
- How do we predict, detect, and correct errors in data before it is too late?
- consumer systems that capture essential characteristics of data at point of creation
- We don't understand failure modes/models/impacts on data
- Since bandwidth is so important, what is being done to pack more data or to make the pipes fatter?
- Presentation methods for archive content
- User interfaces to simplify/speed management of archival storage systems
- Why not focus on storage as a whole rather than just archive?
- Wide area wireless transmission networks. Better, faster, and cheaper than today's wired internet connections.
- Data & models to answer question "how many copies to get given reliability"
- What % of long-term archived content ever actually gets revised?
- Self-monitoring preservation services/products
- Concept of "eternal memory"
- Much more reliable media
- What's the DOD already solved
- Long-term (over 100 years) high-capacity storage
- Interoperability particularly on distributed/federated/synchronization-type features
- Not reliable hardware but statistical analysis reliability of strategies for viability
- Where's the intelligence in the system
- Vendor agnostic standards
- International and industry-wide standards
- How to prove or demonstrate reliability
- Evaluation & analysis—what gets saved
- Mobility of large amounts of data (physical or digital?)
- Digitization capture to long term preservation analysis. Cost/benefit of preservation
- Archive governance & analysis modeling for decision-making
- "greening" digital preservation—support Obama Admin priorities
- R&D on an open workflow to put any type of product into mix & match architecture
- What might we learn from natural systems (biomimicry), the "preservation ecosystem"?

2015 Predictions

- Will we be able to store our work on data crystals?
- Everything will be virtualized
- Advancements in capacity (not necessarily performance) will come from somewhere (faith?)
- Storage cubes of flash or solid state
- Generation of standard policy sets-computer-actionable rules
- There will be a preservation network offering specific preservation services based on commercial technology
- New models more (much more) data
- More sophisticated data management techniques for determining what to throw out
- De-duplication techniques for archiving long-term
- Drowning in data
- Time-sharing will be back
- There will be a major loss of important digital data

- DOD
- Tape might be nearing the end
- Optical tape will come & go again
- Object metadata will be stored within most file formats
- Archives & libraries shift to proactive selection of information; industries owns limitations
- How many really big reliable data storage utilities will exist in 2015?
- By 2015 we will have very primitive quantum computing.
- We could be at the early days of optical processing. Today's ghz processors would become dinosaurs almost overnight.
- Will have undergone 10 data migrations at PB scale to new physical infrastructure access layer services
- Use of phase change memory

Other

- Technical test beds examining managed evolution of storage architecture
- Understanding implication of media requirements
- For Hathi-what is the copyright scenario?
- For DuraSpace-is JPEG.2000 a good archival format; would JPEG XR be better?
- Why will the move to USB3 produce such a dramatic change?
- Advantages of pre-configured packages of open source
- Can we leverage "high replica" models to test exotic media at lower cost & lower risk
- Life of the carrier vs. ability to restore bits—seems like all vendors want SaaS

- SOME PARTICIPANT COMMENTS FROM OPEN DISCUSSIONS:

  - Real costs are not just hardware/software. Also have network, maintenance, cooling etc.
  - You can spend a lot of money on hardware/power, but people are the most costly.
  - Power costs are significant when you get to the petabyte level.
  - Deploy open source by packaging it up as server images in the cloud.
  - Need to look at how to “engineer-out” the end-users.
  - Scaling up vs. scaling out.
  - You want to preserve a petabyte of data for a century – how will you test a level of reliability? You cannot afford to do the experiments that will guarantee that level of reliability.
  - Physical materials deteriorate. How much do you want to pay to monitor the media?
  - Users don't want to hear that you won't get stuff back. This is always true in any physical process. We have to engineer strategies that are actually true in all cases.
  - Trust issue. People who work in digital materials don't trust digital materials, but the other folks who don't understand the technology are actually racing to get rid of their analog materials.
  - What does it take to verify? We need low-cost ways of monitoring the degradation of the material. Just because they are not met in the software today doesn't mean that we don't need it.
  - Moore's law is the economic engine we're used to, and it's ending. Twice the performance for half the effort is disappearing. With Moore's law coming to the end, is there any emerging data storage media that would provide 100 year storage capability that would require no power?
  - The Japanese are taking “eternal storage” very seriously. How seriously are vendors looking at eternal storage products?
  - Where are the standards activities? SNIA is actively engaged in developing two efforts in this area. The XAM-Interface is an application-level interface coordinating information metadata between applications and storage systems (http://www.snia.org/forums/dmf/programs/xam/). It is addressing the logical migration issue. Also the Self-contained Information Retention Format (SIRF).
  - What need is there for security on the drive, irrespective of the data that goes on it? Drive can hold a lot of data, no sense encrypting data that doesn't need it.
  - The cost of IT is rising faster than GDPs. Need to turn IT into a value center. How can we monetize the benefit to the business.
  - Media might last 10-30 years, but the interfaces and software drivers only last 5. We have a model that is broken for the interface.
  - Need standards for interface.
  - The archivists are always going to be leveraging somebody else's business. Nobody is going to build hardware this cheap because the archiving is a tiny fraction of the industry. Archivists have to build their stuff out of other customers’ business.
  - Nobody is going to do the R&D to figure out how to make super-reliable storage because it is very expensive. We have to give up the idea of the magic media.
  - If you're in the media preservation business, what do you preserve? The environment? The canonical form of the data?
  - Can we get service level agreements about moving data?
  - Location needs to be considered as part of mean time to data loss.
  - If we assume that we'll never have super-reliable storage, is there a different question we should be asking?
  - Maybe we need an amalgamation of organizations including LC, science, government, etc? What would it take as a market for vendors to provide a product if specified by an appropriate group?
  - Items are preserved over time based on the interest in those objects over time. Has anyone considered that there is a ton of information that we are preserving that no one will care about?
  - In some environments (e.g., the movie industry) you never know which data someone will want

in the future.

- .Digital technologies require that we think up front. It's a complete change in thinking. We don't have the information or the perspective of time on what will be needed in the future.
- The storage market is a churn market; vendors make money by turning over technology.
- The issue is that the vendors don't control the end-to-end. There's a whole infrastructure that needs to be created. If you had to get buy-in from every step of the process, it's not going to happen.
- Most important cause of data loss is operator error. People are reluctant to talk about this. A government agency has faced this problem and come up with a solution: An anonymous reporting system, which won't lead to the need for liability lawyers. This is also being pushed this for the medical industry.
- On the question on how much people are interested in this, there's an anthropological researcher at Microsoft named Cathy Marshall who is looking at the average consumer view.
- The end-to-end perspective includes content production and related error correction options. For example, there is the issues related to forward error correction: how do you put enough extra bits in to get accurate reproduction. Are forward error correction codes in conflict with error correction codes? Are the error correction codes at the data layer going to conflict at the software level? Moving images don't have built-in error correction codes. If you have to make changes to the essence then make changes. Within each compression scheme there are forward error correction schemes. You are choosing to recreate different collections of bits. In forward error correction you choose which bits to reproduce. Compressed or uncompressed video or audio data files could put in place with a different toolkit or set of rules. Lossy compression causes the loss of bits. Each layer would handle the data according to the policies set by the layer.
- With archives at large scale, they need to have more intelligence so that we can trust them with more tasks. Virtualization will also help make migration a regular part of systems.
- A perspective change: technology migration shouldn't be something that looms in the future that we fear, but a part of our regular work. Things don't have to be done all at once. Piecemeal content transformation.
- The concept of archival capacity should include most of what you can read in before the media deteriorates.
- Integrity: how do we build a storage substrate that gives us high enough bit accuracy rates to do what we want? Need to look at a higher level than that. Recognize that no matter what we do we're going to have bit failures. Need a body of engineering that talks about how to bound and minimize the scope of loss. If you lose a block, how much do you lose? Close in on this from 2 sides. Set of mechanisms on how to reduce the error rate on the substrate. Complement with discussion on data structures to bound loss. How we might promote research and the sharing of experiences on that side of that fence
- If we're going to start measuring success, we need to have a much more nuanced approach to what we're talking about. For example, review the significant properties of file formats. Some bits are more important than others. What are you most concerned about in preserving those formats and information?
- Metrics for failure- how to figure out our failure rates. In high-end computing, many folks have failure statistics but can't talk about them. As a community we're reluctant to discuss these things.
- There are bodies that have put together audits. If you're putting together a repository, look at some of these: DRAMBORA and the TRAC checklist, which is being proposed as an ISO standard. See the following URL's:
- http://wiki.digitalrepositoryauditandcertification.org/bin/view
- http://www.repositoryaudit.eu/
- Focus should be on ingest, access and preservation requirements. Focus on architecture based on these three requirements.

- There are other communities that have aligned interest. Not exactly the same but similar. Very a provocative piece in *Nature* about the tragedy of scientific data loss.
- Big research organizations could collaborate, and establish requirements for SLA's
- There is a Federal government effort specifically devoted to resiliency in high-performance computing that is focusing on an integrated multi-disciplinary interagency approach. See: http://institute.lanl.gov/resilience/
- There is a Federal government interagency group on digital data that is working on developing guidelines
- New expertise that is needed for the people who will be dealing with this in the archives. Where are we going to get the people from? Who will train the trainers? Whole structure of libraries and archives is going to change.