Strategies and New Technology for Long Term Preservation of Big Data

Presenter: Sam Fineberg, HP Storage

Co-Authors: Simona Rabinovici-Cohen, IBM Research – Haifa

Roger Cummings, Antesignanus

Mary Baker, HP Labs

Outline

- Introduction
- What makes digital preservation of big data so hard?
- SNIA Long Term Retention technology for big data
  - Self-contained Information Retention Format (SIRF)
- SIRF serializations for cloud and tape
- SIRF status and summary

Big Data…..

- Really is BIG…..
  - 2.5 quintillion (10^{18}) bytes of new data created per day in 2012 (source IBM)
- And the move to the “Internet of Things” is only going to increase this volume
  - 19.8 Billion connected devices by 2020 (source McKinsey)
    - Only 4.2 billion smartphones and tablets, 3.4 billion PCs
- Data analytics is improving all the time
  - Therefore historical information has significant value
    - Apply new techniques and algorithms to gain new insights
    - Need to ensure ALL necessary information is captured to extract full value
- If storing big data is difficult today, how do we preserve it?
  - Over the short term? Over the long term?

SNIA Survey from 2007

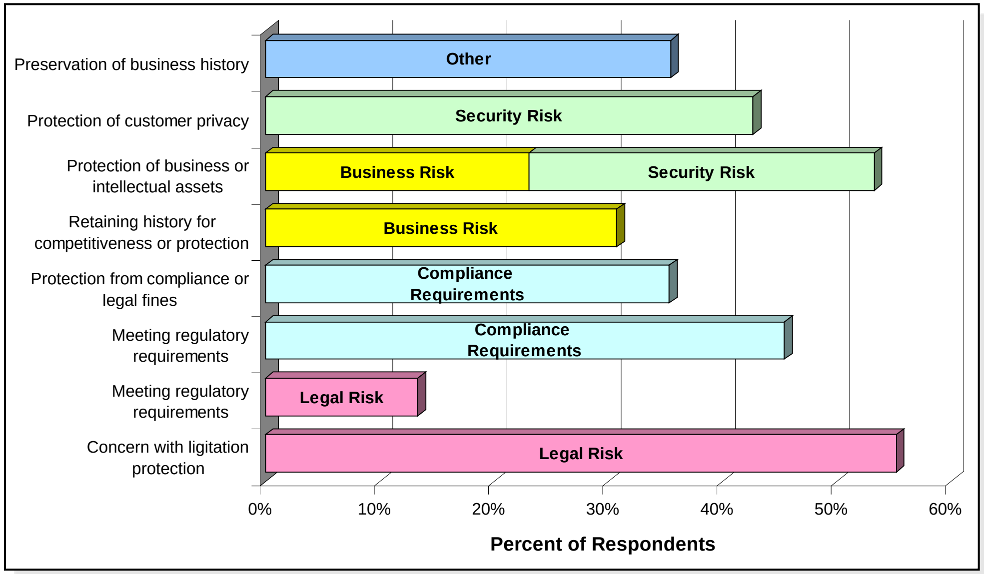

Top External Factors Driving Long-Term Retention Requirements:
Legal Risk, Compliance Regulations,
Business Risk, Security Risk

100 YEAR Archive Task Force

Source: SNIA-100 Year Archive Requirements Survey, January 2007.

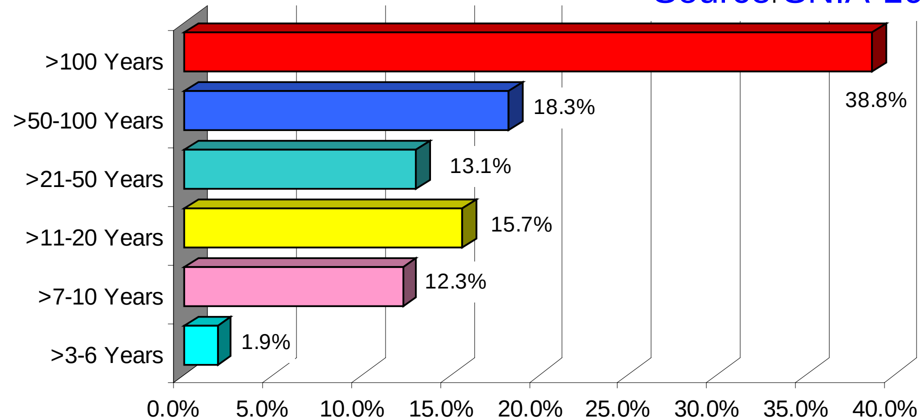

What does Long-Term Mean?
Retention of 20 years or more is required by 70% of responses.

The Need for Digital Preservation of Big Data

- Domains that have Big Data require preservation
- Regulatory compliance and legal issues
  - Sarbanes-Oxley, HIPAA, FRCP, intellectual property litigation
- Emerging web services and applications
  - Email, photo sharing, web site archives, social networks, blogs
- Many other fixed-content repositories
  - Scientific data, intelligence, libraries, movies, music

Scientific and Cultural

Satellite data is kept for ever

We would like to keep digital art for ever

Healthcare

X-rays are often stored for periods of 75 years

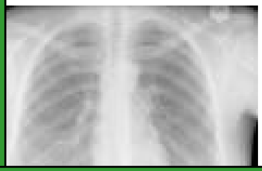

Records of minors are needed until 20 to 43 years of age

M&E

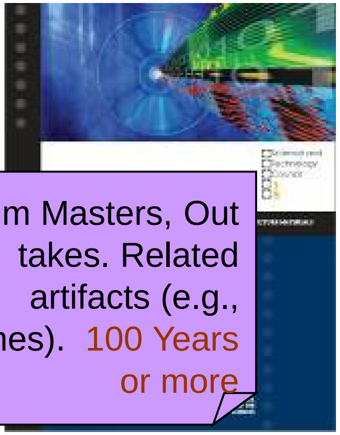

Film Masters, Out takes. Related artifacts (e.g., games). 100 Years or more

Outline

- Introduction
- What makes digital preservation of big data so hard?
- SNIA Long Term Retention technology for big data
  - Self-contained Information Retention Format (SIRF)
- SIRF serializations for cloud and tape
- SIRF status and summary

What makes digital preservation so hard?

Data volume

- Digital content is generated at a much faster rate than analog
- “Big Data” (volume, velocity, variety) makes this even harder

Technology issues

- Media obsolescence
- Media degradation/failure
- Format obsolescence
- Loss of context/metadata

Human issues

- Large-scale disaster
- Human error
- Economic faults
- Attack
- Organizational faults

Large scale & long time periods → Even improbable events will have an effect

Real Life Example Problem

2003

2007

To: roger.cummings@veritas.com
From: fred@nowhere.com
Subject: Something or other

To: roger_cummings@symantec.com
From: sue@somewhere.com
Subject: Something else

Same people?? Could you PROVE it 20 years on?

To: gary.phillips@veritas.com
From: fred@nowhere.com
Subject: Something or other

To: gary_phillips@symantec.com
From: sue@somewhere.com
Subject: Something else

Goals of Digital Preservation

- Digital assets stored now should remain
  - Accessible
  - Undamaged
  - Usable

- For as long as desired – beyond the lifetime of
  - Any particular storage system
  - Any particular storage technology

- And at an *affordable cost*

Preservation Practices must vary by time

- Can't predict what will change – only know it will

- This means processes are key
  - Must be evolvable
    - Current processes get us to the next step
    - At that point we will likely need new processes to take over
  - Must not destroy what we are trying to protect
  - Standards make evolution easier

- A good archive is almost always in motion
  - Digital preservation is not a static activity!

Practices must vary by context

- What do we preserve?
  - Raw data? Analytical results? Context? etc.
  - Depends on customer needs and economics
    - Can we afford to keep the raw data?
    - Is it ever practical to re-analyze?
  - Can’t always predict the eventual use
    - What insights will the data provide in 1 year, 5 years, 20 years, etc?
    - Trade off, based on predicted “business value”
- Many possible preservation techniques and strategies, which could vary over time

Outline

- Introduction
- What makes digital preservation of big data so hard?
- SNIA Long Term Retention technology for big data
  - Self-contained Information Retention Format (SIRF)
- SIRF serializations for cloud and tape
- SIRF status and summary

SIRF: Self-contained Information Retention Format

Being developed by SNIA Long Term Retention (LTR) TWG

Photo courtesy Oregon State Archives

An Analogy

- Standard physical archival box
  - Archivists gather together a group of related items and place them in a physical box container
  - The box is labeled with information about its content e.g., name and reference number, date, contents description, destroy date
- SIRF is the digital equivalent
  - 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

SIRF Properties

- SIRF is a logical data format of a storage container appropriate for long term storage of digital information
  - A storage container may comprise a logical or physical storage area considered as a unit.
  - Examples: a file system, a tape, a block device, a stream device, an object store, a data bucket in cloud storage

- Required Properties
  - 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 Components

A SIRF container includes:

- **Magic object**: identifies SIRF container and its version
- **Preservation objects** that are immutable
- **Catalog** that is
  - Updatable
  - Contains metadata to make container and preservation objects portable into the future without external functions

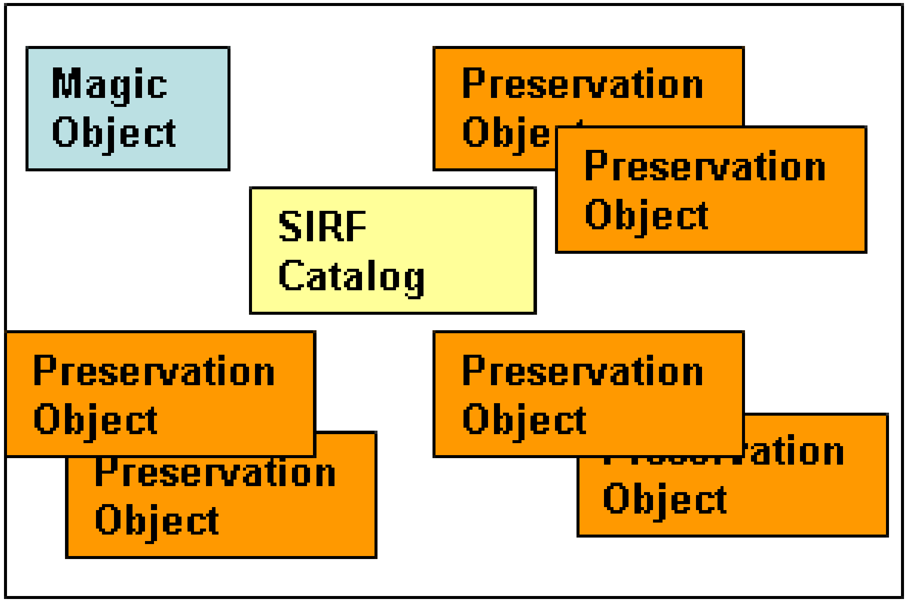

Outline

- Introduction
- What makes digital preservation of big data so hard?
- SNIA Long Term Retention technology for big data
  - Self-contained Information Retention Format (SIRF)
- SIRF serializations for cloud and tape
- SIRF status and summary

SIRF Serializations

- The TWG has developed serializations for CDMI and LTFS
  - specifies how a CDMI container or LTFS Tape also becomes SIRF-compliant
  - XML and JSON schemas for the SIRF catalog
- A SIRF-compliant CDMI container or LTFS Tape enables a future CDMI/LTFS client to “understand” containers created by today’s CDMI/LTFS client
  - The properties of the future client is unknown to us today
  - “understand” means identify the preservation objects in the container, the packaging format of each object, its fixities values, etc. (as defined in the SIRF catalog)

SIRF Serialization for CDMI: Interface

- CDMI API can be used to access the various preservation objects and the catalog object in a SIRF-compliant CDMI container
- Example
  - Assume we have a cloud container named "PatientContainer" that is SIRF-compliant
    - the container has a catalog object
    - each encounter is a preservation object
    - each image is a preservation object

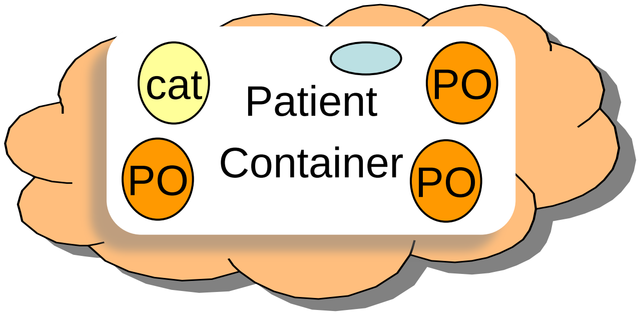

  - We can read the various preservation objects and the catalog object via CDMI REST API as follows:

    GET <root URI>/ PatientContainer>/sirfCatalog

    GET <root URI>/<PatientContainer>/encounterJan2001

    GET <root URI>/<PatientContainer>/chestImage

SIRF Serialization for CDMI

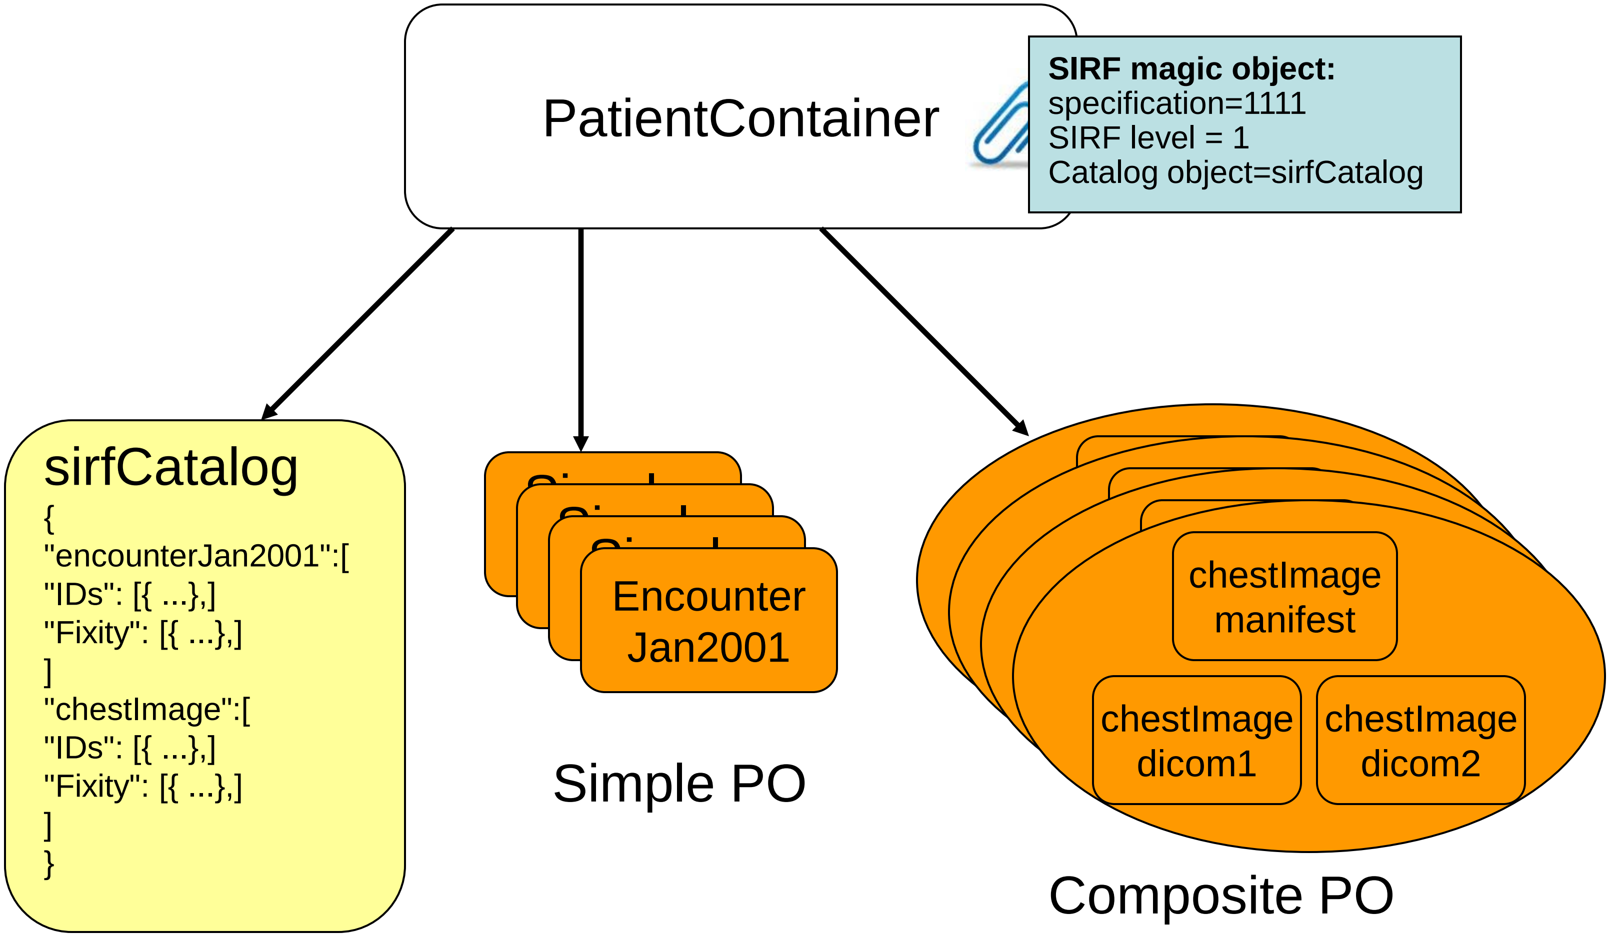

SIRF Serialization for LTFS

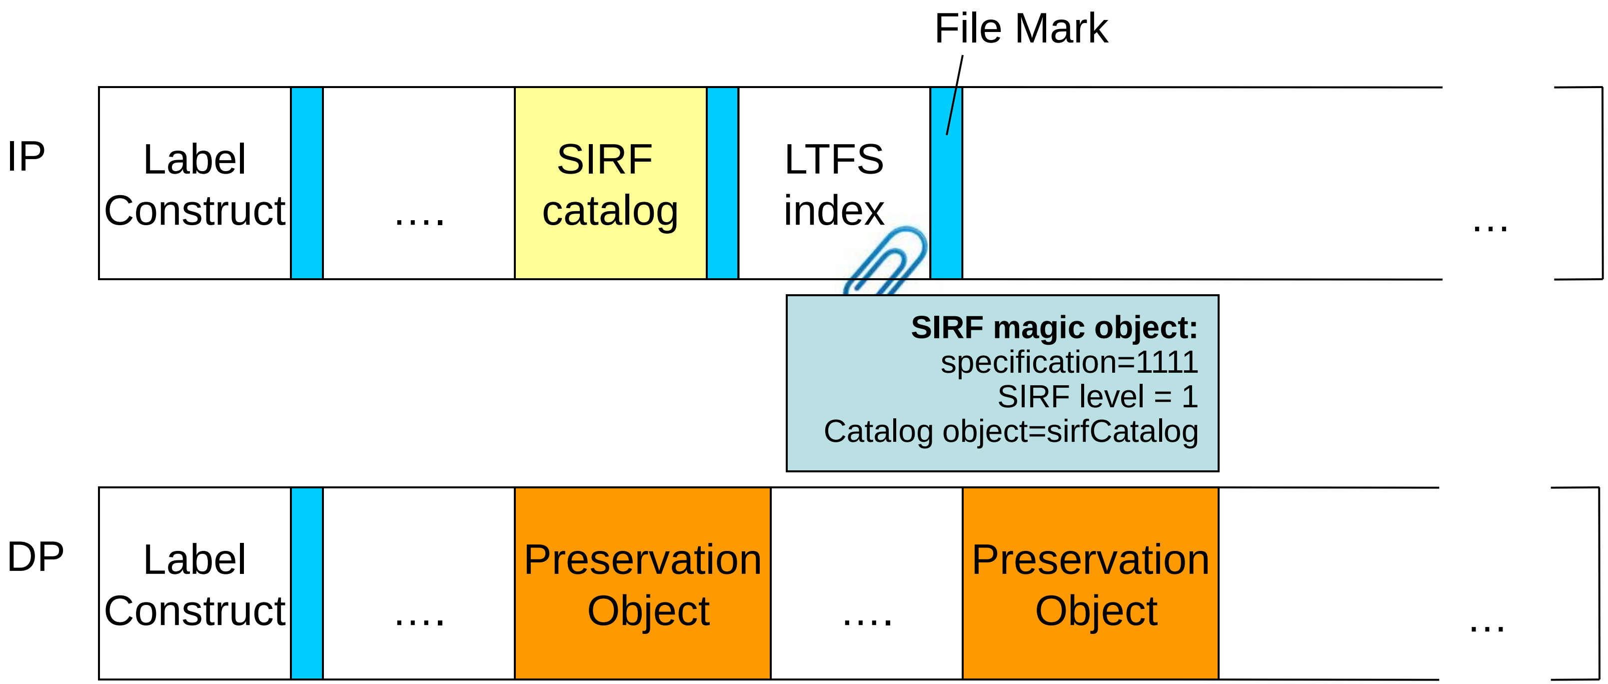

- The index partition of the tape is 2 wraps which is 37.5 GB in LTO-5 and probably larger in LTO-6.
  - The tape index partition is large enough to hold the LTFS index, the SIRF catalog, and even additional information e.g. thumbnails of images

SIRF Serialization for LTFS: General

- A LTFS Tape can also be a SIRF Container when:
  - The SIRF magic object is mapped to extended attributes of the “LTFS index” root directory
    - The magic object includes, for example, specification ID and version, SIRF level, reference to SIRF catalog
  - The SIRF catalog resides in the index partition and formatted in XML
  - A SIRF preservation object (PO) that is a simple object (contains one element) is mapped to a LTFS file
  - A SIRF PO that is a composite object (contains several elements) is mapped to:
    - a set of LTFS files (one for each element) and a manifest file that its content includes the IDs and fixities of the element data objects

Outline

- Introduction
- What makes digital preservation of big data so hard?
- SNIA Long Term Retention technology for big data
  - Self-contained Information Retention Format (SIRF)
- SIRF serializations for cloud and tape
- SIRF status and summary

SIRF Status

- More information is available at snia.org/ltr
- A SIRF draft is in development, and will be available for public review soon
- A reference implementation of SIRF is in development by the TWG

Summary

- Long Term Retention is hard
  - Volume, Velocity and Variety make it even harder
  - Need to retain not only information of immediate interest but ALL other information needed to make it fully usable in future
- SIRF Helps, providing a “digital box” for data and metadata
  - Includes self describing metadata to help “understand” the contents of the container in the future
- No single technology will be usable over the timespans mandated by current digital preservation needs
  - SNIA CDMI and LTFS technologies are among best current choices
    - Are good for perhaps 10-20 years
  - SIRF provides a vehicle for collecting all of the information that will be needed to transition to new technologies in the future