Digital Preservation Technology

Neil Jefferies
Bodleian Libraries
University of Oxford
@NeilSJefferies

ILIDE 2018

Session Structure

- Risks and Strategies
- Hardware
  - Keeping the bits safe
- Software
  - Making the bits useful/usable
- Useful design patterns
- Cloud-based preservation

Risks and Strategies

Putting technology into context

ILIDE 2018

The Digital Preservation Environment

- Ongoing access to digital information involves…
  - Systems – Hardware & Software
  - People & organisations
  - Processes & procedures
- All of which are subject to obsolescence and change

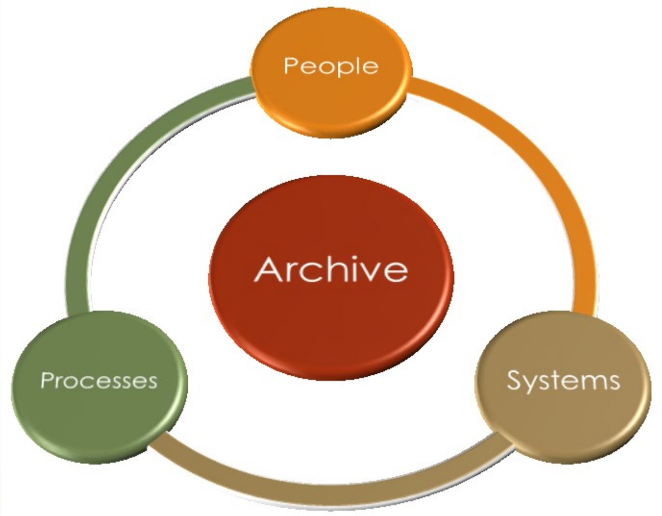

Bodleian Libraries University of Oxford
pasig preserving and archiving special interest group
Digital Preservation Coalition

TRAC Requirements for Digital Archives

- Mandate and Commitment to Digital Object Maintenance
- Organizational Fitness
- Legal and Regulatory Fitness
- Efficient & Effective Policies
- Adequate Technical Infrastructure
- Acquisition and Ingest
- Preservation of Digital Object Integrity, Authenticity & Usability
- Metadata Management & Audit Trails
- Dissemination
- Preservation Planning and Action

Bodleian Libraries UNIVERSITY OF OXFORD pasig preserving and archiving special interest group dpc Digital Preservation Coalition

“Traditional” Preservation

The mission of the Preservation Directorate at the Library of Congress is to assure long-term, uninterrupted access to the intellectual content of the Library's collections.

http://www.loc.gov/preservation/about/org.html

- Effective Strategies for Traditional Materials
  - Materials science: stable media
  - Physical Protection
    - Stable environmental conditions
    - Controlled use
  - Conservation, Repair and Reformatting

Bodleian Libraries UNIVERSITY OF OXFORD

pasig preserving and archiving special interest group

DigitalPreservationCoalition

Degradation behaviour…

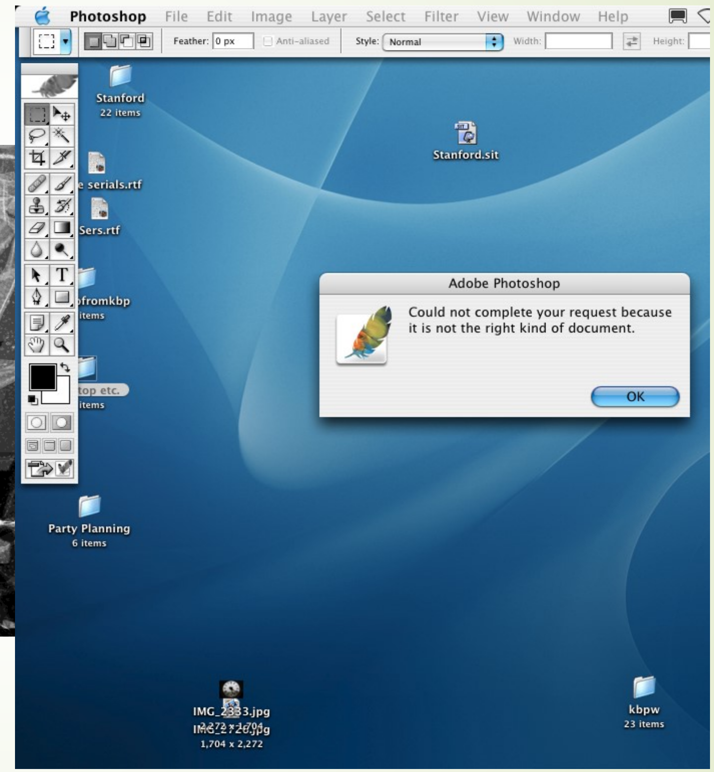

Examples…

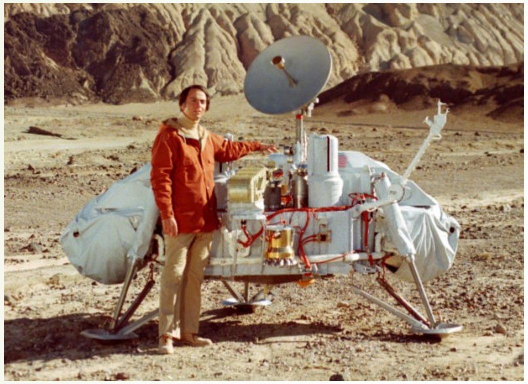

- Viking Mars landings 1976
- 1988: Extracting 3000 images from tapes took 2 years of reverse engineering effort
- 2003: Biological researchers looking for data couldn't read the raw data format. Instead they hired students to rekey the data from printouts.

- The 1986 BBC Domesday project commemorated 900 years of the William the Conqueror's survey
- ~1M participants submitted images, maps, video, statistics and stories
- Published on two 300MB per side 12" Laser-discs, requires a Philips VP415 "Domesday Player"
- 1999-2003 – Two recovery attempts using emulation and migration
- 2090 – Final date copyright expires

Bodleian Libraries University of Oxford
pasig preserving and archiving special interest group
DigitalPreservationCoalition

Quantifying the Risks

- Bit rot (not really)
  - ~~Media decay~~
  - Corruption
  - Transmission errors
- Obsolescence
  - File Format
  - Software
  - Hardware
  - Media
- Technology Failure
  - Software
  - Hardware
  - Media
- Loss of context
  - Data but no codebook
- Ambiguous IP State
  - Copyright
  - Licensing
- Disasters
  - Natural disasters
  - War
- Organizational failure
  - Loss of will
  - Human error
  - Sabotage
  - Economic failure

High Level Strategies

- Redundancy
  - Replication
  - Heterogeneity
  - Location
  - Organisation
- Emulation
  - Hardware
  - Software
  - "Appliance"
- Encapsulation
  - I prefer "Locality"
- Succession Planning
  - People
  - Technical
    - File formats
    - Media
    - Hardware
    - Software
- Capability

Digital Preservation (long-term access) is realized as a series of relays over time.

Preservation and Archiving is **not**…

- Backup
- Disaster Recovery
- Business Continuity
- Document Management Systems
  - Compliance systems
- Enterprise Content Management Systems
  - Document, Records, Web, Email
- Digital Asset Management
  - Images, Audio, Video
- Hierarchical Storage Management (HSM)
  - or Tiered Storage

- However, any preservation & archiving system can and probably will…
  - interact with such services to obtain or disseminate content
  - use such services to deliver certain preservation requirements

Hardware

Keeping the bits safe

ILIDE 2018

Trends: Disk vs Tape

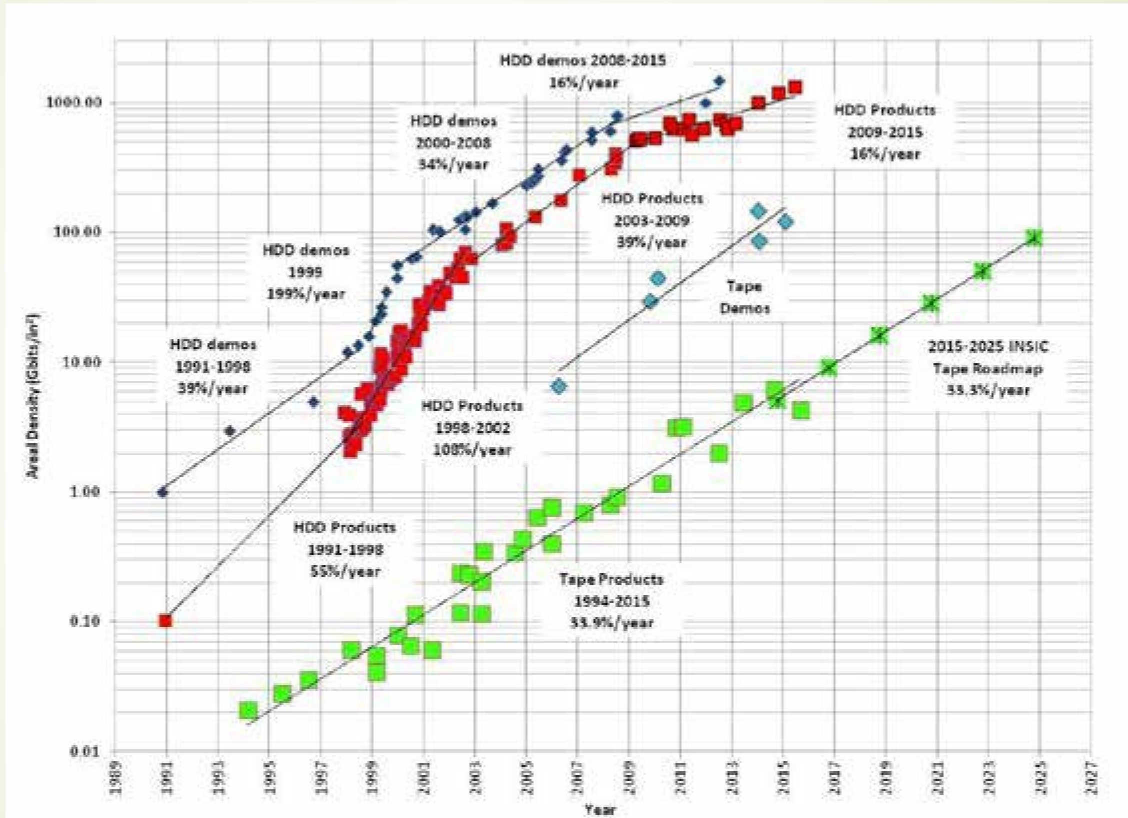

Bodleian Libraries UNIVERSITY OF OXFORD
pasig preserving and archiving special interest group
Digital Preservation Coalition

Hard Disk

- Properties
  - Latency 5-15ms (this has not changed significantly for years)
  - Sustained data transfer rate 200MB/s
  - Capacity per unit (2017) 12TB
  - Cost per TB (2017) $56
  - Requires power (power cycling not recommended)
- Lifetime
  - 5 year warranties (MTBF figures are meaningless)
  - Interface longevity: SATA 2003, SAS 2004, FC (ANSI) 1994, Ethernet (802.3ab Gigabit) 1999
- Systemic Risk
  - 3 Manufacturers (Seagate, HGST and Toshiba)
  - Consumer market squeezed by PC substitutes (phones & tablets with flash)
  - Enterprise market squeezed by flash
  - Cloud enables higher utilisation by sharing -> lower unit shipments

Hard Drive Technology

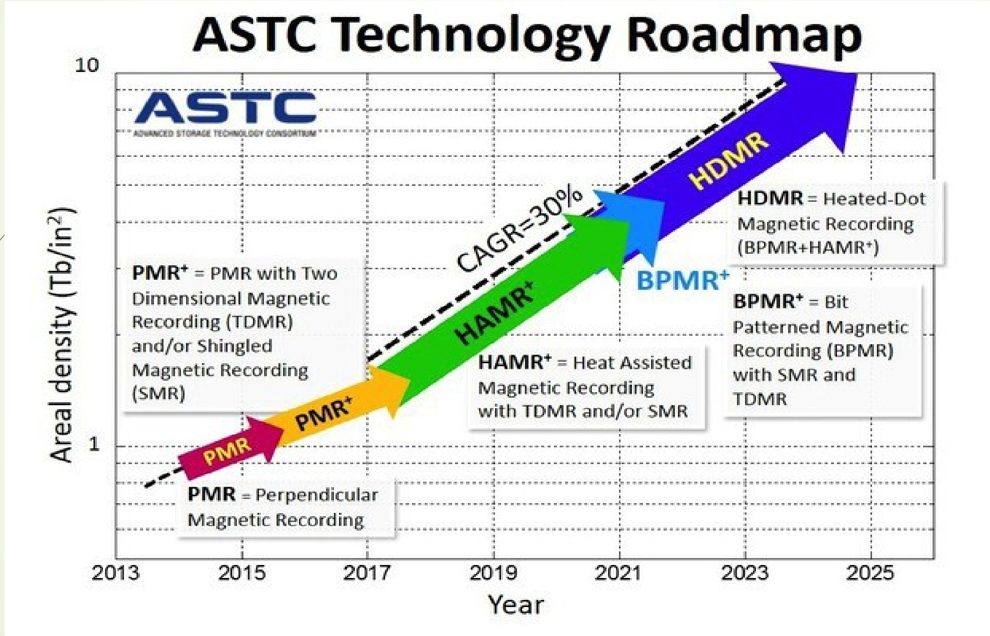

(c) 2017 ASTC (Advanced Storage Technology Consortium)

- Hard Disc Technology is at a critical juncture
- He-filled drives allow more platters, density increases require new technology

Tape

- Properties
  - Latency 100s (load from a robotic library), robot speeds gradually increasing
  - Sustained data transfer rate 360MB/s (faster than HDD once loaded!)
  - Capacity per unit (2017) 12TB
  - Cost per TB (2017) $21 (including library)
  - Media is unpowered, robot still needs power
- Lifetime
  - 2 Formats: IBM Magstar and LTO (Oracle T10K frozen in 2017)
  - 30 year media life (media warranty typically 1 year, though)
  - Drives typically can read back two generations (generations typically 2-3 years for LTO)
  - Drive warranties typically 5 years -> probably safe to keep media 10 years
    - IBM allows formatting older media at higher capacity (new with LTO-8, too)
- Systemic Risk
  - IBM: 1 drive manufacturer
  - LTO: 3 drive manufacturers (HPE, Quantum, IBM)

Flash

- Properties
  - Latency 10us (decreasing rapidly)
  - Sustained data transfer rate 2000MB
  - Capacity per unit (2017) 60TB (highe
  - Cost per TB (2017) $250 (decreasing rapidly)
  - Needs power! (but typically less than a hard drive)
- Lifetime
  - Enterprise SSD guaranteed retention 40 days (Consumer: 1 year, USB: indefinite)
    - In practice, retention is much longer
  - Warrantied according to total bytes written
    - Writing data is primary media degradation mechanism
  - Interface longevity similar to hard disks
- Low Systemic Risk
  - Many manufacturers (>10)

Long Term Total Cost of Ownership

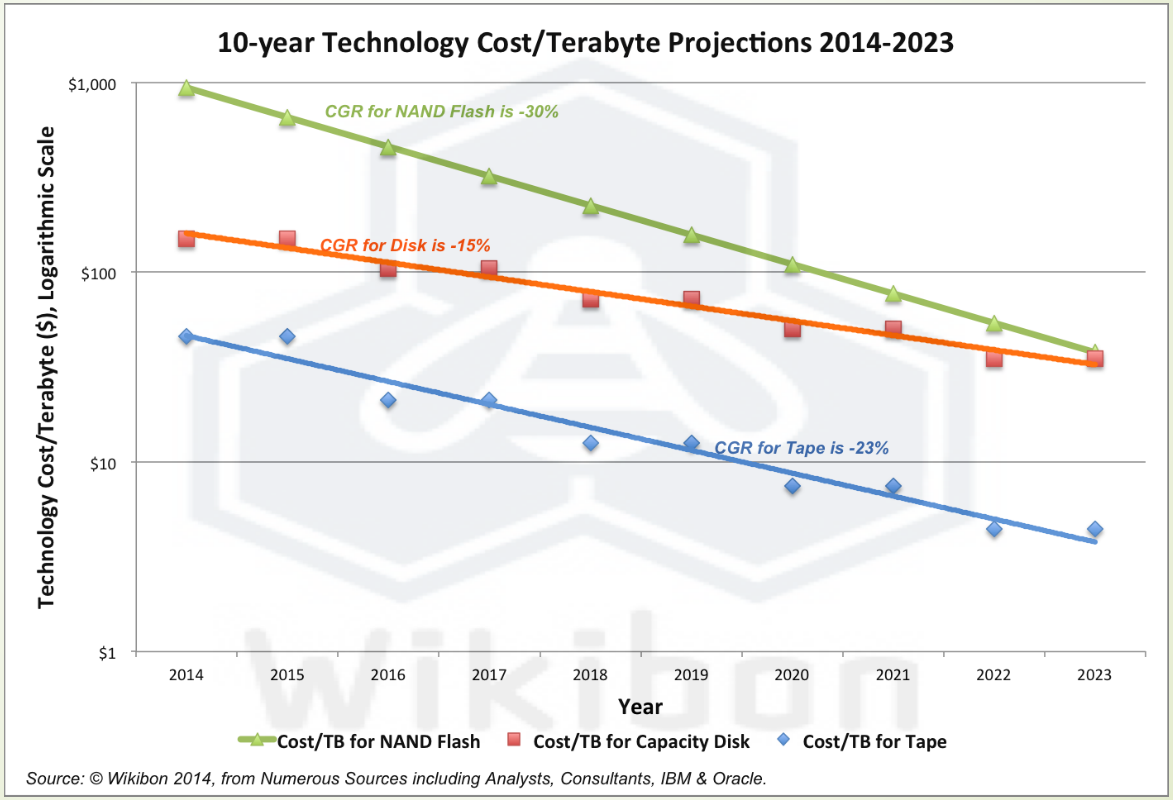

Bodleian Libraries UNIVERSITY OF OXFORD

pasig preserving and archiving special interest group

Digital Preservation Coalition

ILIDE 2018

Archive Optical

- Properties
  - Performance figures are scarce
  - Sustained data transfer rate 40MB/s (similar to Blu-Ray)
  - Capacity per unit (2017) 3.3TB (actually a cartridge of 9 disks)
  - Cost per TB (2017) $100
    - Enterprise SSD guaranteed retention 40 days (Consumer: 1 year, USB: indefinite)
    - In practice, retention is much longer
- Lifetime
  - Claimed 50 years+ for media
  - Drive promise backwards compatibility for all generations (only 2 exist so far)
- High Systemic Risk
  - Archive Optical: 2 Manufacturers (Panasonic, Sony)
    - Interoperability apparently not guaranteed
- M-Disc: 1 Manufacturer (Milleniata, has gone bankrupt once)
  - Variant of CD/DVD/Blue-Ray (LG, Lite-on, Asus produce compatible drives)
  - Has proved very robust in tests but low density (100GB)

Media Summary (2017)

Medium	Latency	Data transfer rate (MB/s)	IO per second	Unit Capacity (TB)	Cost per TB ($)	Suppliers	Data Retention
Offsite Tape (LTO-8)	4h	360	<1	12	11	3	**
Tape Library (LTO-8)	100s	360	<1	12	21	3	**
Hard Drive	10ms	200	200	12	56	3	*
SSD (NVME)	<10us	2000	100,000	60	250	Many	*
Archive Optical	??	40	??	3.3	100	2	***
USB 3.0 Stick	<1ms	200*	1000	2	250	Many	***

Tiering – Cost/Performance Optimisation

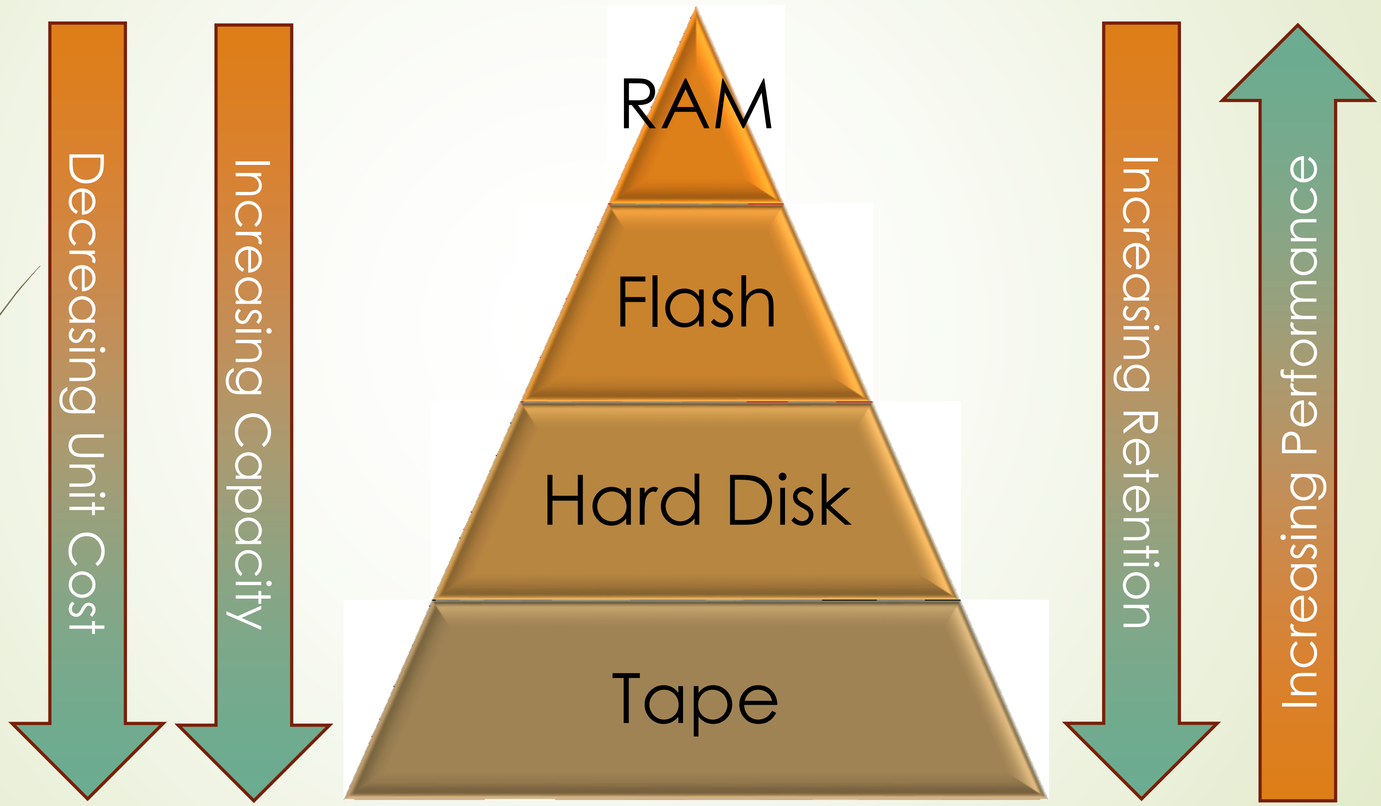

Software

Making the bits useful/usable

Bodleian Libraries
UNIVERSITY OF OXFORD

pasig
preserving and archiving
special interest group

dpc
DigitalPreservationCoalition

ILIDE 2018

Low Level Software

- Hardware controller functions largely replaced by software as complexity increases
- Redundancy and Replication
  - Geographic distribution (power failure etc.)
  - Technology distribution (disk **and** tape, different manufacturers etc.)
    - Don't use replication (or you copy failures/errors)
    - Explicit copy/checking – read tapes back on a different drive!
  - RAID – Redundant Array of Inexpensive disks
    - With large volumes: time to recover > time to next disk failure
    - MAID – Massive Array of Idle Disks, powers disks down to sqve power
  - Erasure coding/clustering (Ceph, IPFS, ScoutFS, ZFS…) – similar technology as media uses
    - Specify a number of fragments and how many needed to recover
    - Faster rebuild times
    - Tunable for fault tolerance/costs balance
- Technology Migration
  - Hardware obsolescence cannot be avoided
  - Multiple copies reduce risk

Workflows

- Archived data *should* be hard to change
  - Versioning with an audit trail can increase storage costs
  - Try to do as much possible beforehand -> Ingest workflows
- Many preservation tools are open source
  - Operations are verifiable and repeatable
- Proprietary and open source preservation systems package these tools with the addition of…
  - Workflow creation and management
  - Reporting
  - Discovery
  - Audit trail generation (PREMIS)
  - Storage integration
- Beware of lock in
  - Always have an exit strategy
  - Using archived material should **not** require the ingest software (or indeed any special software)
- Ingest bottleneck
  - Overly long workflows can lead to loss because material hasn't been ingested yet!
  - Sheer curation – focus on what is necessary

Key Tools

- Fixity
  - All **modern** media have better error checking and correction than SHA etc.
  - Checksum to prevent mistakes/malicious tampering
    - Store these separately, or digitally sign them
    - Check after data is moved
    - Unaudited sysadmin activity can be a significant source of loss – hardware fixes, media rotation, "upgrades"
- Identification and Validation (PRONOM, JHOVE, UDFR)
  - File formats
  - Metadata format and completeness
  - Risk profile for stored content
- Mitigation actions
  - Format migration (with care)
  - Packaging for emulation
  - Metadata extraction
- Specialist Forensic Tools (Forensic Toolkit, BitCurator)
  - Media images
  - Recover deleted information
  - Prune out recognised non-content files

Dissemination

- Beware the "dissemination copy"
  - It is the copy that people will reference, cite and care about
  - It will need to be regularly cross-checked with the archived material
  - It will need to be preserved
  - …so, ideally, generate it on-the-fly from an archival copy and cache it
- Emulation
  - Some formats just cannot be easily migrated or displayed
    - E.g. Macromedia Shockwave, FLASH, Multimedia titles
    - Security concerns with some formats too
  - Possible to emulate most hardware using modern software
    - Able to run older operating systems and software securely
    - "If it can play games then an emulator has almost certainly been written"
    - Most emulators are open source – easy to obtain
    - Long term support is harder – opportunity for DP community
- Discovery
  - Frequently neglected part of re-use
  - Depends on good metdata
  - Incremental curation – expect to add/update metadata over time

Retain Information

- Transformations are rarely 100% accurate
  - Keep the original files and metadata
  - Not always true – e.g. lossless image formats
- Provenance and context are essential
  - Audit trail (needed for certification!)
  - Versioning is better
    - Differential versioning better still
  - Physical provenance applies to a digitised object as well
  - Capture as much information as possible – just store it for later even if it doesn't "fit"
    - RDF/Linked data is a very good mechanism for this
  - Look to digital criminal forensics for guidelines

A digital object should be considered a greater whole comprising several streams of information that can be arbitrarily labelled data or metadata but all of which contribute to the intellectual content of the object

Bodleian Libraries
UNIVERSITY OF OXFORD

pasig
preserving and archiving special interest group

DigitalPreservationCoalition

Helpful Design Patterns

ILIDE 2018

Retain Flexibility

- Avoid making premature decisions
  - Evaluate all decisions from point of view of what it might **prevent** in the future
  - Don't bake decisions into systems at a low level
  - For example:
    - (AIP) Basic archiving of objects need not be file or metadata format specific.
    - (SIP) Formats can be enforced as part of ingest
      - Metadata by the data entry form
      - File format by the characterisation/transformation process
    - (DIP) Display only what is in the DIP
      - And whatever metadata there is (it's probably just XML/XSLT)
      - Keep your designated community as broad as possible
- Accept that you can have multiple ingest/dissemination paths

Layers and Abstraction

- Modularise your systems with well defined interfaces between them
  - Standards – should be simple and easy to implement:
    - REST-ful API's rather than SOAP
    - Text/XML rather than binary formats
      - Not always economic/scalable, so…
    - Open rather than proprietary (documented rather than not!)
  - Easier to replace components
  - Applies to people/processes as much as to technology
    - Applies especially to vendors!
- Look for asynchronous workflows
  - Tasks that can be deferred, decoupled or parallelised help scalability
    - …especially when things aren't going to plan!
  - Break down into simple, small, atomic tasks

Agility

- Technology and technology-driven change move much much faster than archivists are traditionally comfortable with
  - The capacity to create data volumes always seems to slightly exceed the capacity to ingest and archive it
- The window of opportunity to acquire some data is very small
  - Ever tried getting a file off an old tape archive?
- New data types, old media
  - Donated archives now contain digital materials
  - Scholars think up new ways of using digital technologies
    - How do you archive an interactive visualisation
- Funding shifts
  - Library/collection-led bulk digitisation on the wane
  - Research-led, targetted digitisation
    - Multiple collections, sharing an interoperability (RDF again!)

Summary

- Modular – Components can be easily replaced or upgraded
- Flexible – Don't limit your collection arbitrarily
  - Comprehensive – If you hold something, hold as much information as you can about it
- Agile – The drive to change is indicative of usage and interested parties -> sustainability
- Scalable – it can only get bigger
- Asynchronous – removes bottlenecks, helps with scaling
- Resilient – Don't try to avoid failures but plan for and handle them – "Never get down to your last copy"

Cloud-based Preservation

Is nothing new – it is the same hardware and software but with one important new risk factor…

ILIDE 2018

Expect to Migrate!

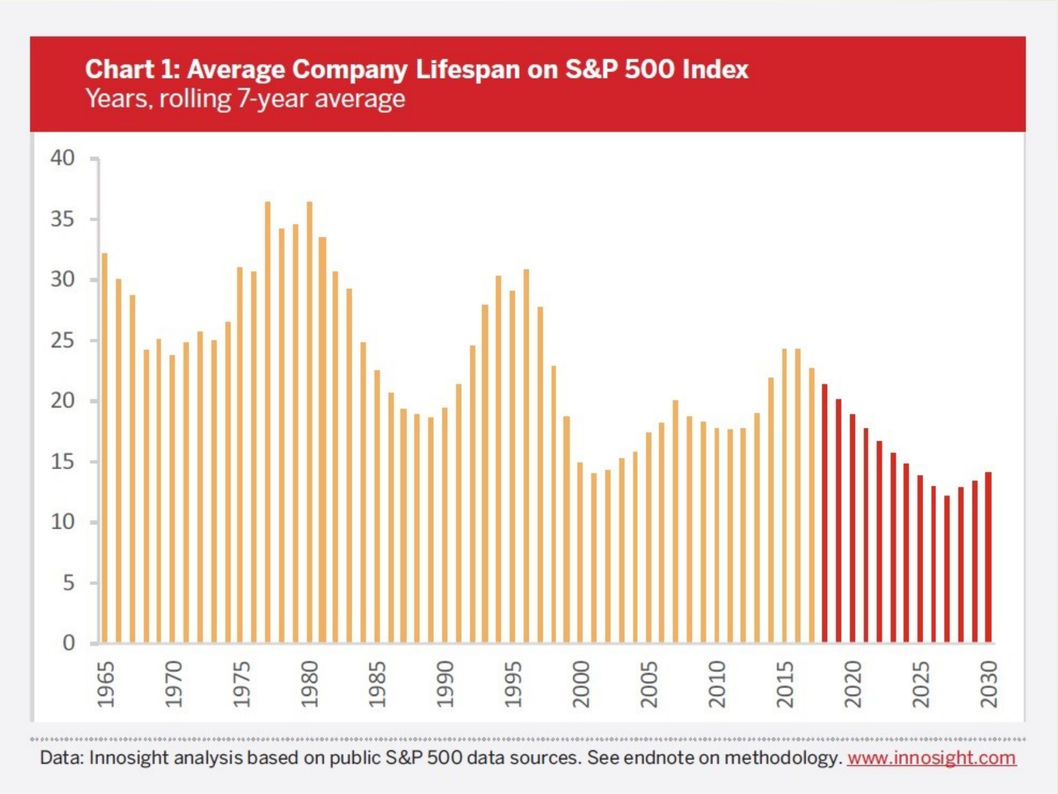

Things will break

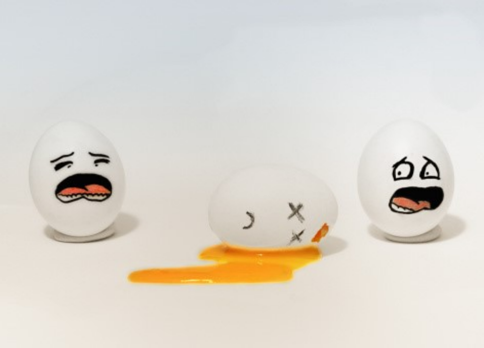

- Preservation is for the long-term
  - In time, improbable events become likely so you need to plan/design for them:
    - Human error (not that improbable!)
    - Technology will fail (in ways that bypass fault tolerant features)
    - Natural (or external) disasters
  - Data will be lost – have a process
    - You may have to deliberately delete stuff
  - Management will change
- Balancing risk vs economics vs reputation

Bodleian Libraries UNIVERSITY OF OXFORD pasig preserving and archiving special interest group Digital Preservation Coalition

Thank you

neil.jefferies@bodleian.ox.ac.uk

Bodleian Libraries UNIVERSITY OF OXFORD

pasig preserving and archiving special interest group

Digital Preservation Coalition