

The large-scale archival
storage of digital objects

The BL DOM Programme for Digital Object Management

Sean Martin
Sean.Martin@bl.uk

Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Brief Introduction to the British Library

What Is The British Library ?

- Created by British Library Act 1972 - commenced 1973
- Merger of British Museum Library (1753), National Reference Library of Science and Invention (1855), National Central Library (1916), and National Lending Library for Science and Technology (1961)
- Subsequent incorporation of British National Bibliography in 1974, India Office Library and Records in 1982, and British Institute of Recorded Sound in 1983
- Flagship building at St Pancras - largest public building project in Great Britain in 20th century - opened in 1998

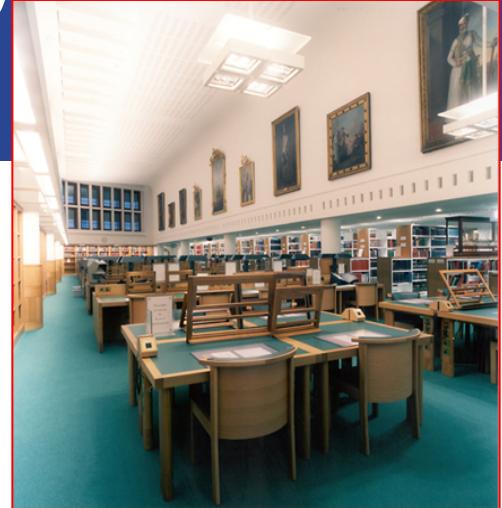


Brief Introduction to the British Library

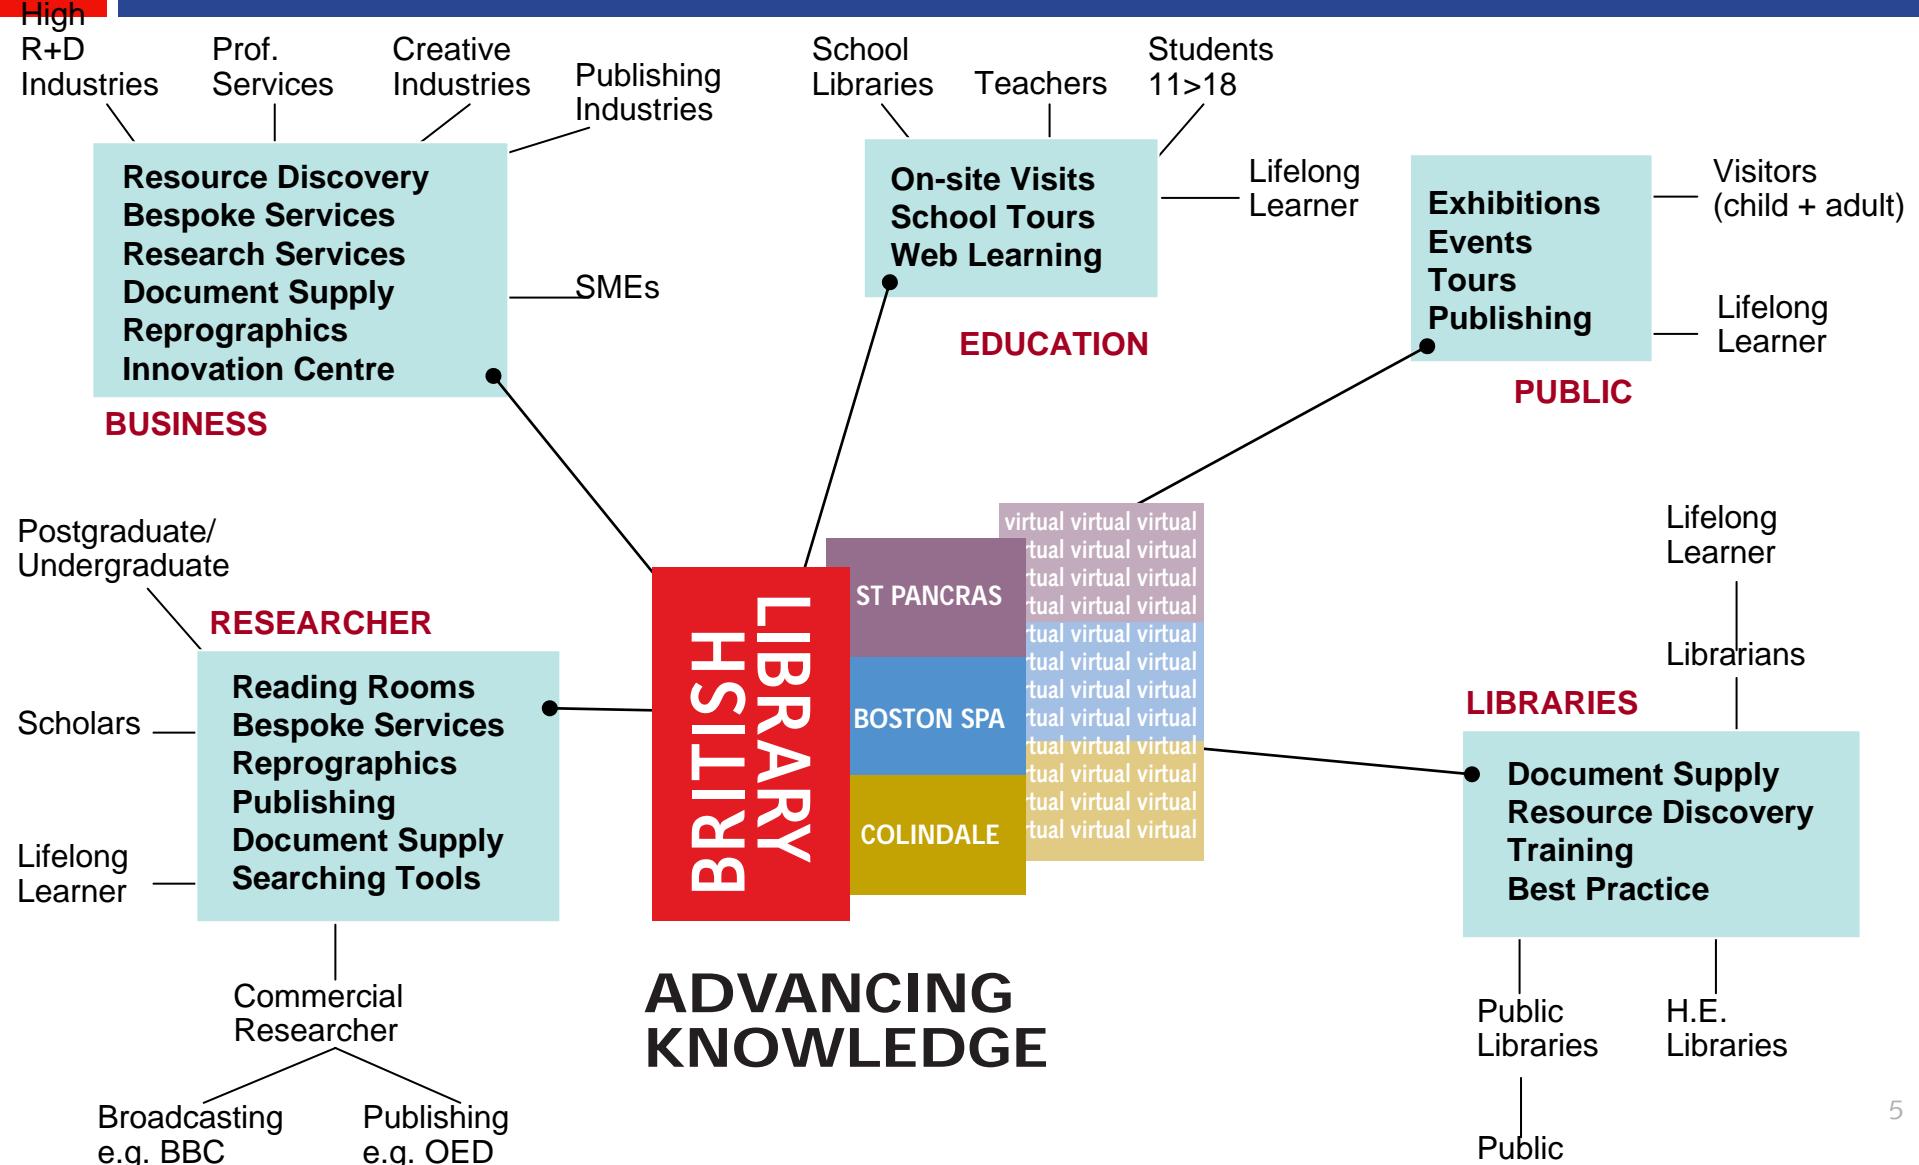
World-Class Research Library

Key Statistics 2003/4

- 150 million physical items
- 42 million catalogue records added 3 million
- 5.3 million items consulted or supplied
- 398,000 reading room visits (1,266 seats)
- 572,000 items received on legal deposit
- 655 km shelf capacity 93% full added 7 km Items
- 24.4M Web Site Hits (www.bl.uk)
- 2,246 staff
- £89.6 million Grant in Aid and £25.4 million trading income



Brief Introduction to the British Library - Customers



Brief Introduction to the British Library – Strategic Priorities

1. **User experience:** improve the user experience by developing an easy-to-use interface and coherent, integrated range of services for users from all disciplines and chosen markets
2. **21st century resource discovery:** deliver world-class resource discovery by streamlining our processes and exploring new approaches
3. **Digital infrastructure:** develop the library's digital infrastructure to collect, preserve and provide sustainable, long-term access to the UK's digital information heritage
4. **Collection development and stewardship:** recalibrate and communicate the Library's collection development policy to ensure alignment with user needs, and continue to ensure good stewardship for our holdings
5. **Financial sustainability:** sustain our commercial and grow our donated income streams; free up resources for key priorities
6. **People and organisation:** align our organisation to maximise our effectiveness in delivering our strategy
7. **Strategic development:** Develop our strategies regarding the Library's next set of critical challenges and opportunities

Agenda



- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Introduction - Programme Mission and Vision

Our mission is to enable the United Kingdom to preserve and use its digital output forever

Our vision is to create a management system for digital objects that will:

- store and preserve any type of digital material in perpetuity
- provide access to this material to users with appropriate permissions
- ensure that the material is easy to find
- ensure that users can view the material with contemporary applications
- ensure that users can, where possible, experience material with the original look-and-feel

Introduction - Programme Scope

We need a generic and cost-effective approach

- to take in material coming from many sources
- to take in material of any and all types
- to store it securely for the long term
- to allow controlled access
- to be enduring

Introduction - Life cycle scope

DOM is concerned at present with the familiar processes of (in conventional library terms):

Collection

- Selection
- Acquisition
- Accession
- Description

Retention

- Storage
- Preservation

Access

- Resource discovery
- Delivery
- Rendering

A complete life cycle would also include

- creation
- deletion

Introduction - Content Drivers

- Legal deposit legislation for non-print material: royal assent in October 2003 but still needs secondary legislation to bring it into force
- Existing voluntary deposit scheme operational since 2000
- Digitised versions of BL material from early '90s onwards
- New digitisation initiatives: newspapers, sound, etc
- Electronic journals
- Sound Archive's 15TB of material per year (with 50 year collection)
- Web archiving
- Cartography and datasets
- &c &c

Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- ■ Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Principles - DOM Programme

- Our approach is to be incremental, not 'Big Bang'
- Prototype so as to learn, understand, reduce risk and uncertainty, and demonstrate the basis of a good solution
- Use standard industry tools (e.g. Microsoft Message Queue)
- Aim for 3 releases per year
- A principal goal is to define an overall long term "logical architecture"
 - Within this, there will be successive generations of physical architectures
- We will use our knowledge of the storage marketplace to manage storage procurement
 - We are certain that we will need very large amounts of storage, but we are uncertain when – so we need flexible and scalable procurement

Principles - System Design

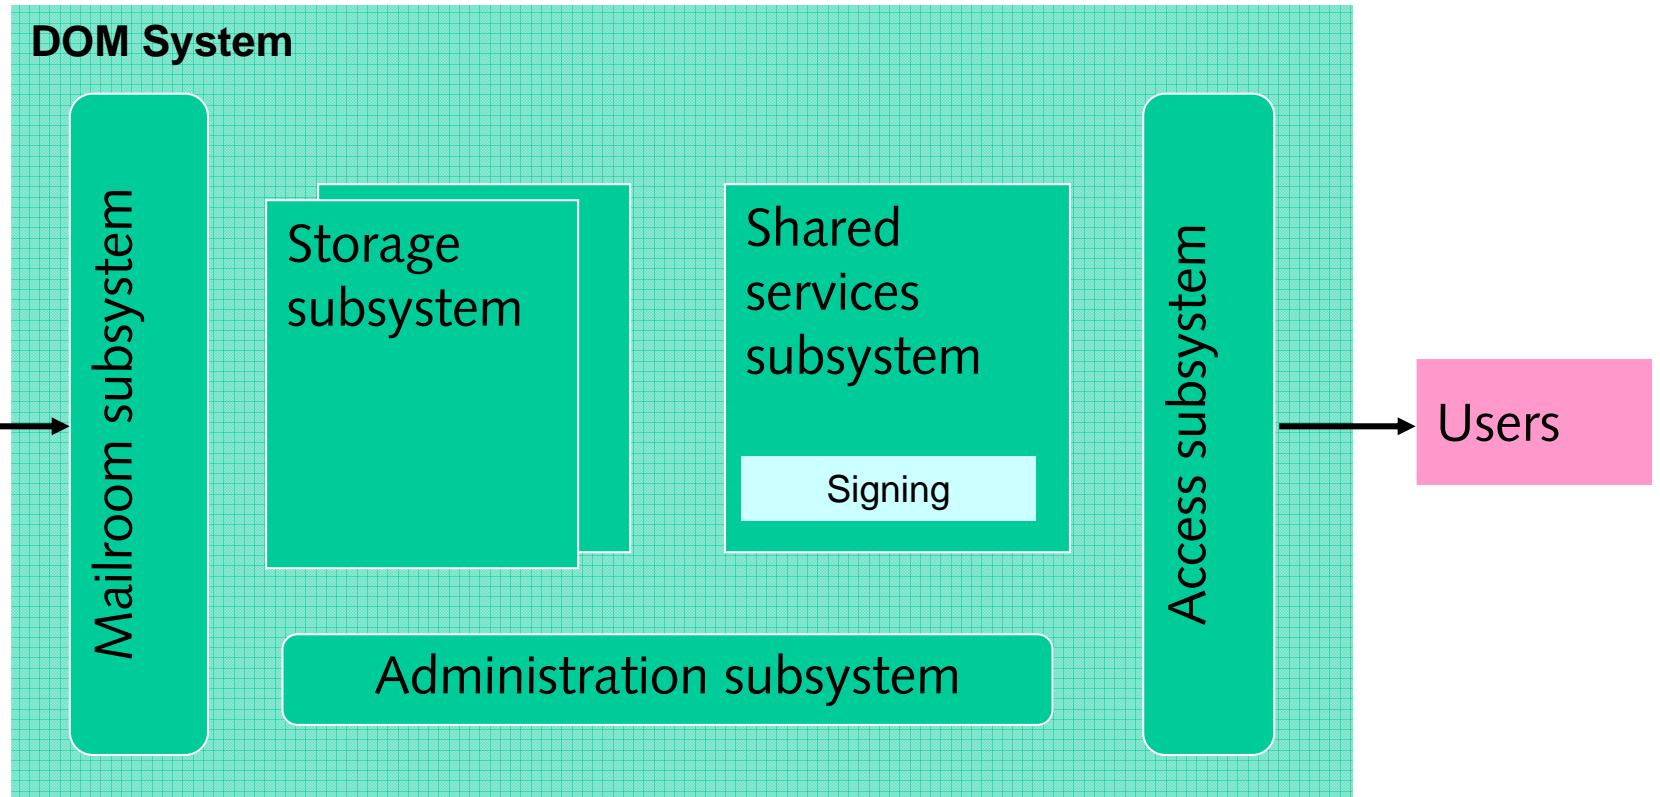
- Significant number of objects will be stored in perpetuity
- Objects can be considered to be invariant and some will be large
- Objects will typically be accessed infrequently

- Each object will have a unique persistent invariant identifier (DOMID)
- All systems external to the Store interact only using a DOMID

- The design of the system must be inherently scaleable in terms of capacity, the number of objects, and the ability to deliver objects
- Need inherent resilience so that object loss is extremely unlikely
- Short interruptions to, or degradation in, service can be tolerated, but extended loss of complete service cannot be tolerated

- Integrate off-the-shelf components
- Be cost conscious

Principles - Overall architecture



Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Integrity and authenticity - conventional approach: The BL stamps content as it is received



Integrity and authenticity – conventional approach: Example uses of BL date stamping



Rhone-Poulenc. Source:

'Reports of Patent, Design, and Trade Mark Cases', 1996 no. 4 p125

A European patent in French covered a paste for the preparation of pharmaceuticals. When translated into English, the phrase "35-80%" was mis-copied as "35-50%". Although this clerical error was spotted and corrected, the incorrect English version had been published and therefore the more limited patent was held to be valid. Rhone-Poulenc appealed. Part of the evidence was the dated BL copy of the correction: a lawyer argued that "on this basis, the date of publication was therefore the date on which the translation was made available to the public in the British Library." But there were all sorts of arguments about dates and legal details: the company was refused permission to correct the official translation, although the Patent Office was allowed to issue an erratum slip

Viziball Ltd. Source:

'Reports of Patent, Design, and Trade Mark Cases', 1988 no. 11 p213

A dispute about whether 1 of 2 co-patentees owned more of the patent than the other was taken to appeal. The patent covered squash balls with a reflective lining. An earlier US patent covered the same ground. The appeal judge said, inter alia, "It is clear from the ... US patent no. 4042236, made available to the public on the shelves of the Science Reference [now British] Library on 6 September 1977, that the use of [this] material on game balls ... was known before the priority date of the present application." Appeal dismissed.

Integrity and authenticity – conventional approach: Physical examination

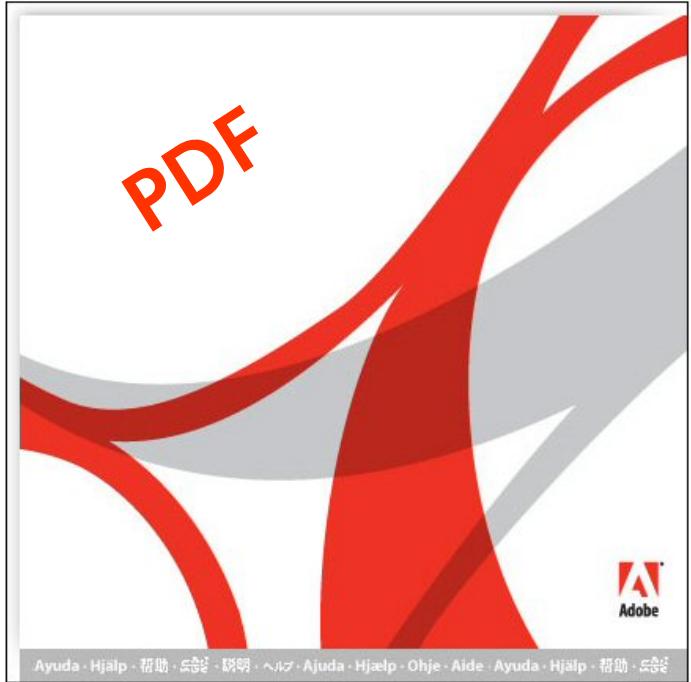


One can examine the chemical composition of the physical item and ensure that:

- The paper is contemporary
- The ink is contemporary
- The binding is contemporary

One can examine the item to detect signs of tampering

Integrity and authenticity - digital objects



Can one can examine the chemical composition of the digital item and ensure that:

- The paper is contemporary ?
- The ink is contemporary ?
- The binding is contemporary ?

NO

How can one can examine the item to detect signs of tampering?

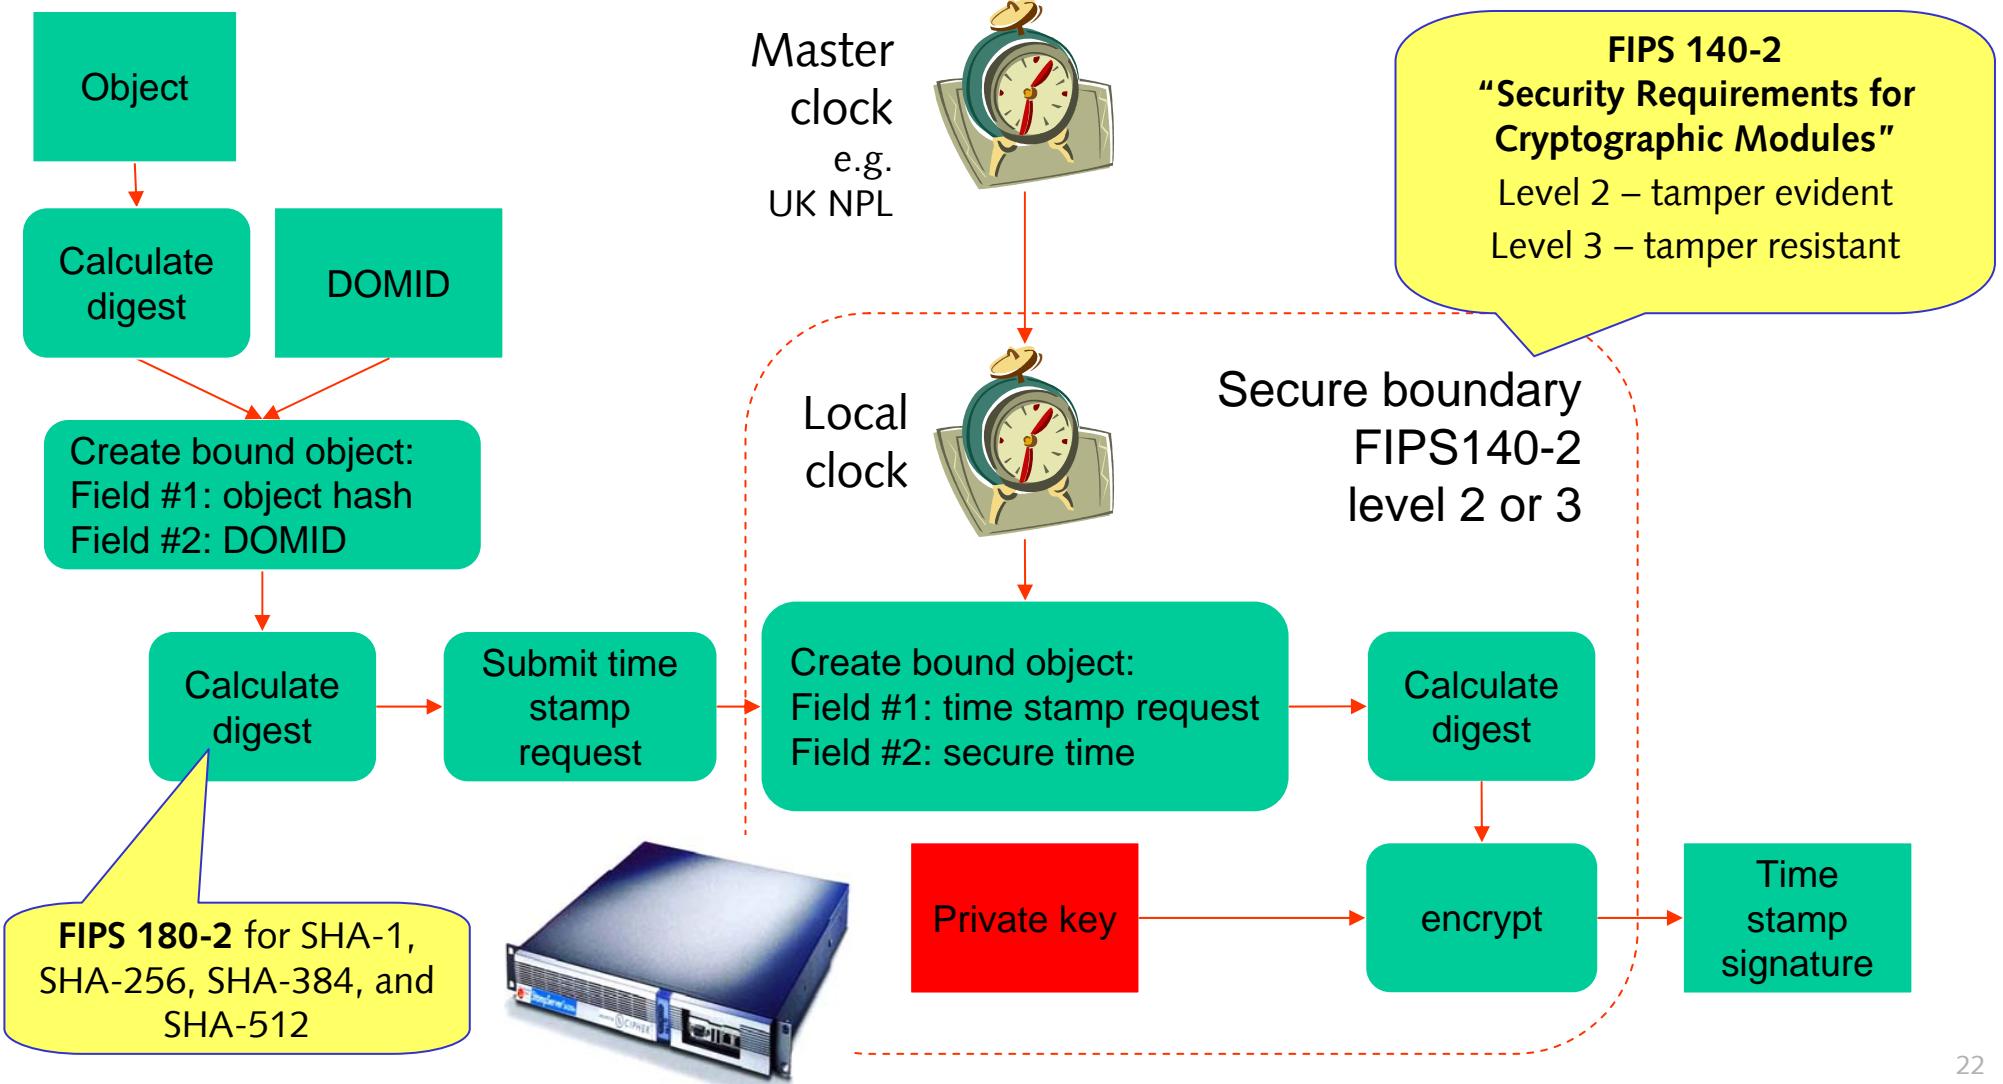
How can one provide assurance that the item is authentic? It may now:

- Be on the 10th generation of hardware
- Passed through the hands of ~20 administrators

Integrity and authenticity - digital objects

- Authenticity:
"Basis for assurance that a re-presented object is identical to the original"
 - Based on the use of cryptographic digital signing techniques
 - Each object is signed and time-stamped when it is ingested
 - The signature is verified when required
 - The signing mechanism is "tightly" controlled
- Integrity: *"Basis for detecting corruption in the store"*
 - System monitors continuously the object store to detect object corruption
 - Based on using a Secure Hash Algorithm (SHA-1 extended to SHA-512)
 - It would then initiate object recovery
- Integrity and Authenticity can be determined locally within the architecture

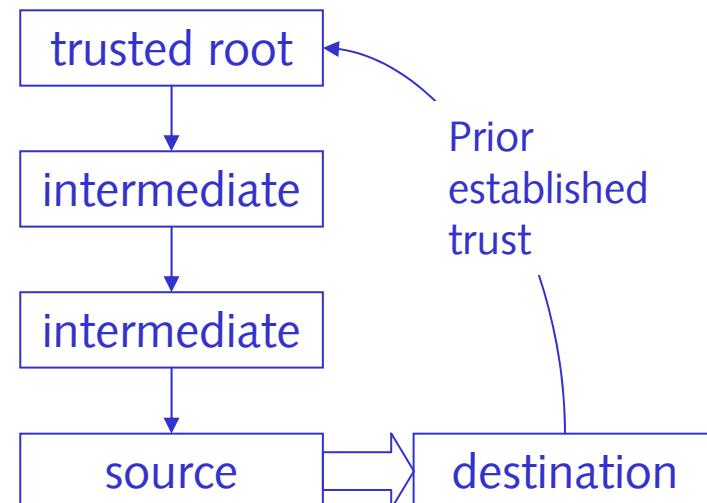
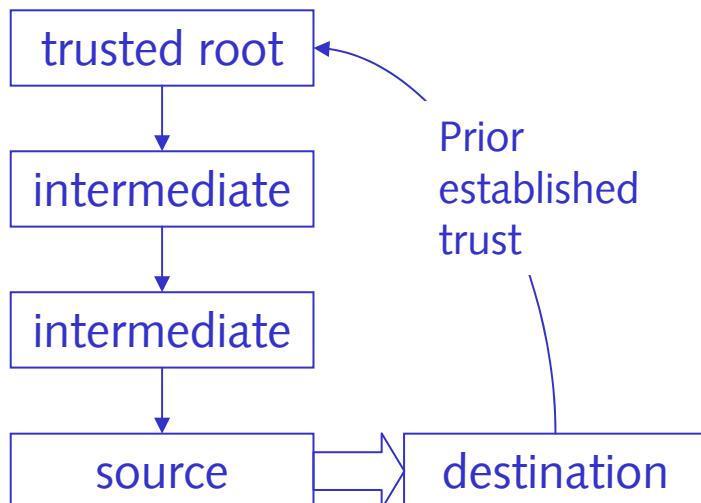
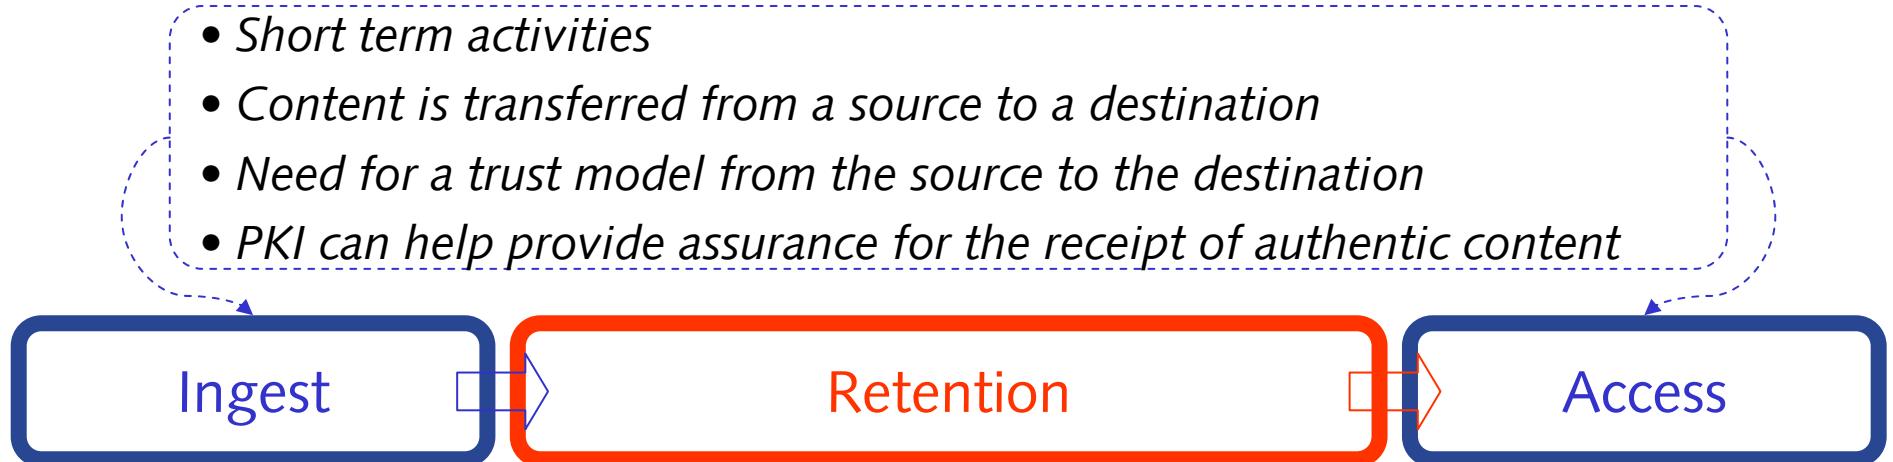
Integrity and authenticity - cryptographic time stamping



Integrity and authenticity - digital objects

Three distinct phases for providing assurance

- Short term activities
- Content is transferred from a source to a destination
- Need for a trust model from the source to the destination
- PKI can help provide assurance for the receipt of authentic content



Integrity and authenticity - digital objects

Three distinct phases for providing assurance

- *Short term activities*
- *Content is transferred from source*
- *Need for a trust model from source to destination*
- *PKI can help provide assurance for the receipt of authentic content*



- *Very long term activity*
- *Content is held within preservation store*
- *No need for a trust model that links two or more parties*
- *The BL is a trusted long term repository for multiple content types*
- *It has no need for assurance provided by short terms organisations such as the current root Certificate Authorities*
- *The Preservation Store will thus incorporate self signed BL root Certification Authority and a time source derived from the UK NPL*

Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- ■ Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Infrastructure: Disaster tolerance rationale for a multi-site design

- One can obtain commercial disaster recovery (DR) solutions for common equipment configurations
- However one cannot obtain such solutions for systems comprising multi-100 Tb systems
- So we must build in the need for DR into the design of the system
- A single site solution, subject to a common-mode disaster, would suffer considerable loss of availability after a disaster, and so is not acceptable
- This implies that we need a multi-site solution
- Conventionally based on a master-standby where 50% of kit delivers service
- Our design is based on the use of multiple autonomous independent peer sites that cross-synchronise so 100% of the kit delivers normal service
- Service continuity: full service, albeit slower, is deliverable by only one site

Infrastructure - context of the storage market: resilience and performance

- The dominant segment of the market focuses on delivering high performance within a highly resilient single site
- However:
 - Many of our objects will be rarely accessed
 - So we do not want to pay for “maximised” performance we do not need
 - We have resilience by using multiple sites, hence we have a reduced need for resilience within a site
 - so we do not want to pay for “maximised” resilience we do not need
- These observations helped us in designing a cost-effective large scale resilient solution

Infrastructure - context of the storage market: procurement and rolling programmes

- A major cost is in physical storage
- The market for storage systems is changing rapidly, and this implies that supplier “lock-in” is not sensible
- We thus need flexibility to change supplier over time
- Cost of storage is reducing by 30-40% per year
- So we procure on rolling basis just ahead of demand
- We also will replace storage on a rolling basis on expiry of warranty

- The rolling procurement & replacement programmes imply the need to be able to support a heterogeneous hardware product solution
- The design of the logical architecture thus needs to support storage sourced from multiple storage vendors

Infrastructure – conclusions in context of storage market:

- We do not want to pay for “maximised” performance that we do not need
- We do not want to pay for “maximised” resilience within a site that we do not need
- We will procure as we need storage, and we do not need to be tied to a single vendor

- These all imply that we can seek to obtain commodity storage hardware solutions from the marketplace, so in that sense:
 - We manage the market
 - It does not manage us

Infrastructure - generic drivers in architectural design

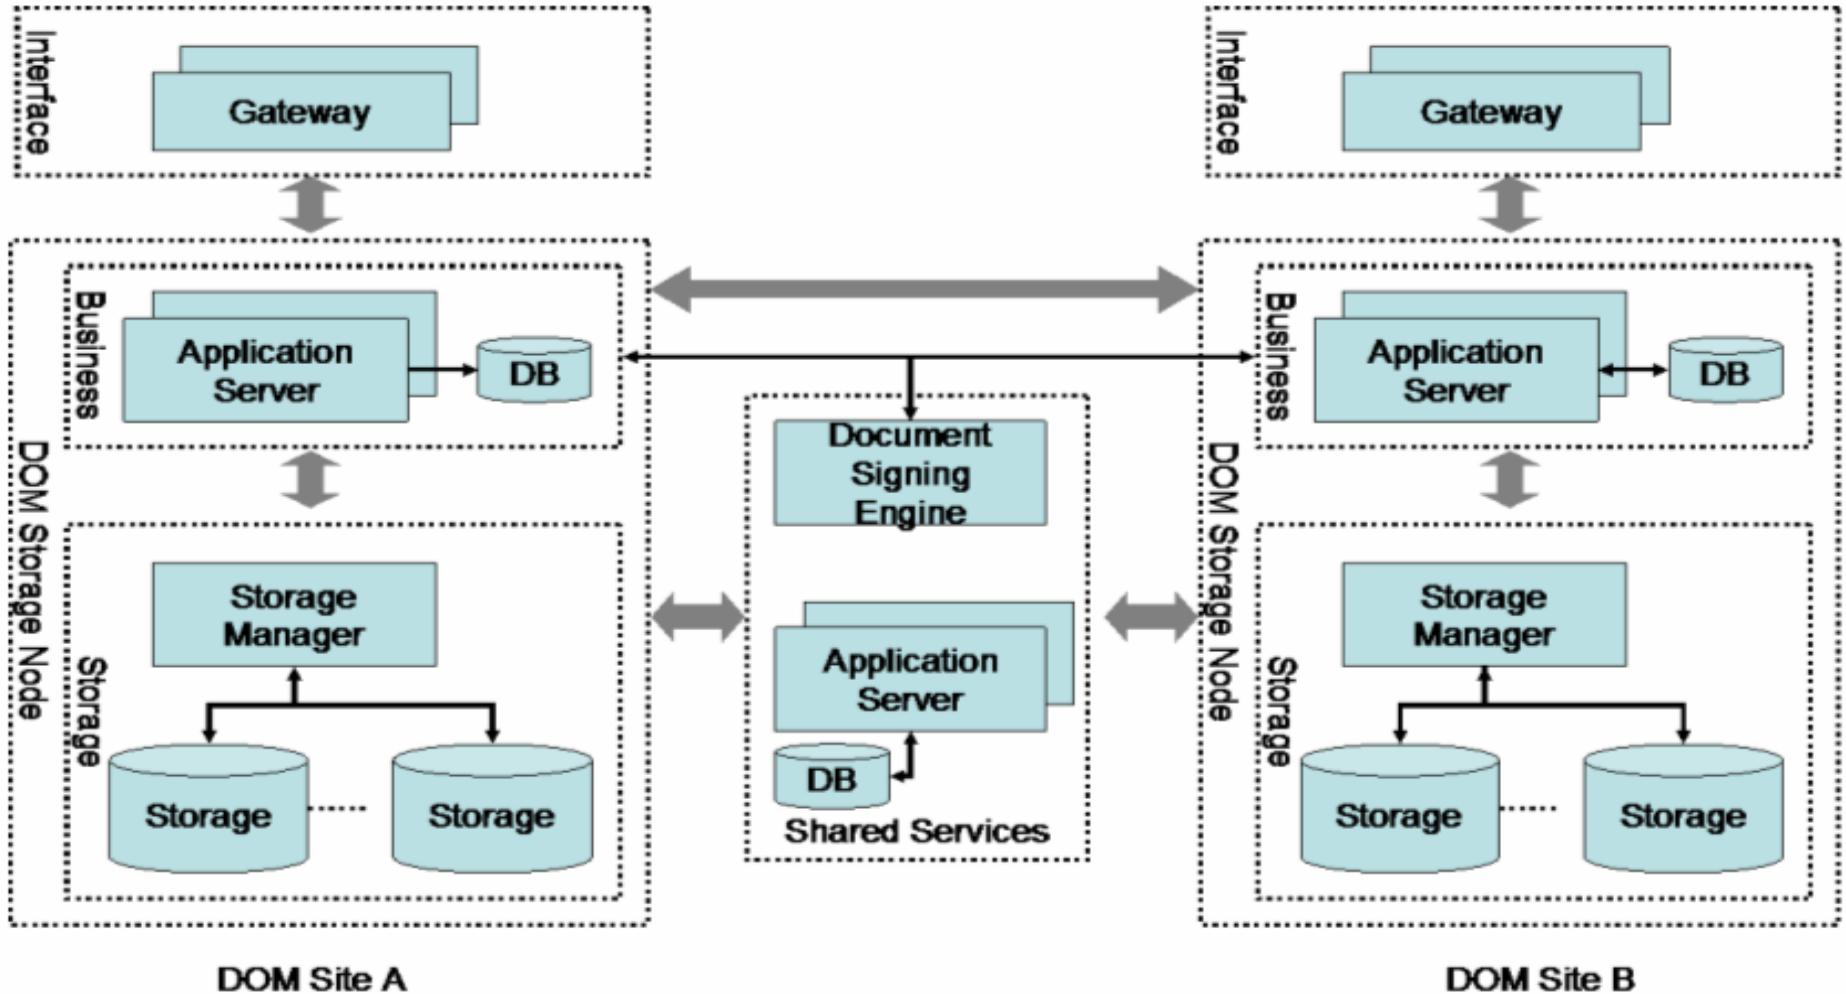
- Manage Total Cost of Ownership covering a complete life cycle:
 - Initial purchase
 - Operations support
 - Data Centre Costs
 - Application support and enhancement
 - Replacement cost (hardware and application)
- Disaster recovery
 - Minimise impact on service through common mode environmental incident / disaster
 - Minimise risk of failure
 - Maximise continuity of service
 - Minimise time to recover
- Adaptability of architecture for anticipated requirements
- Performance (though commodity performance seems adequate)

Agenda

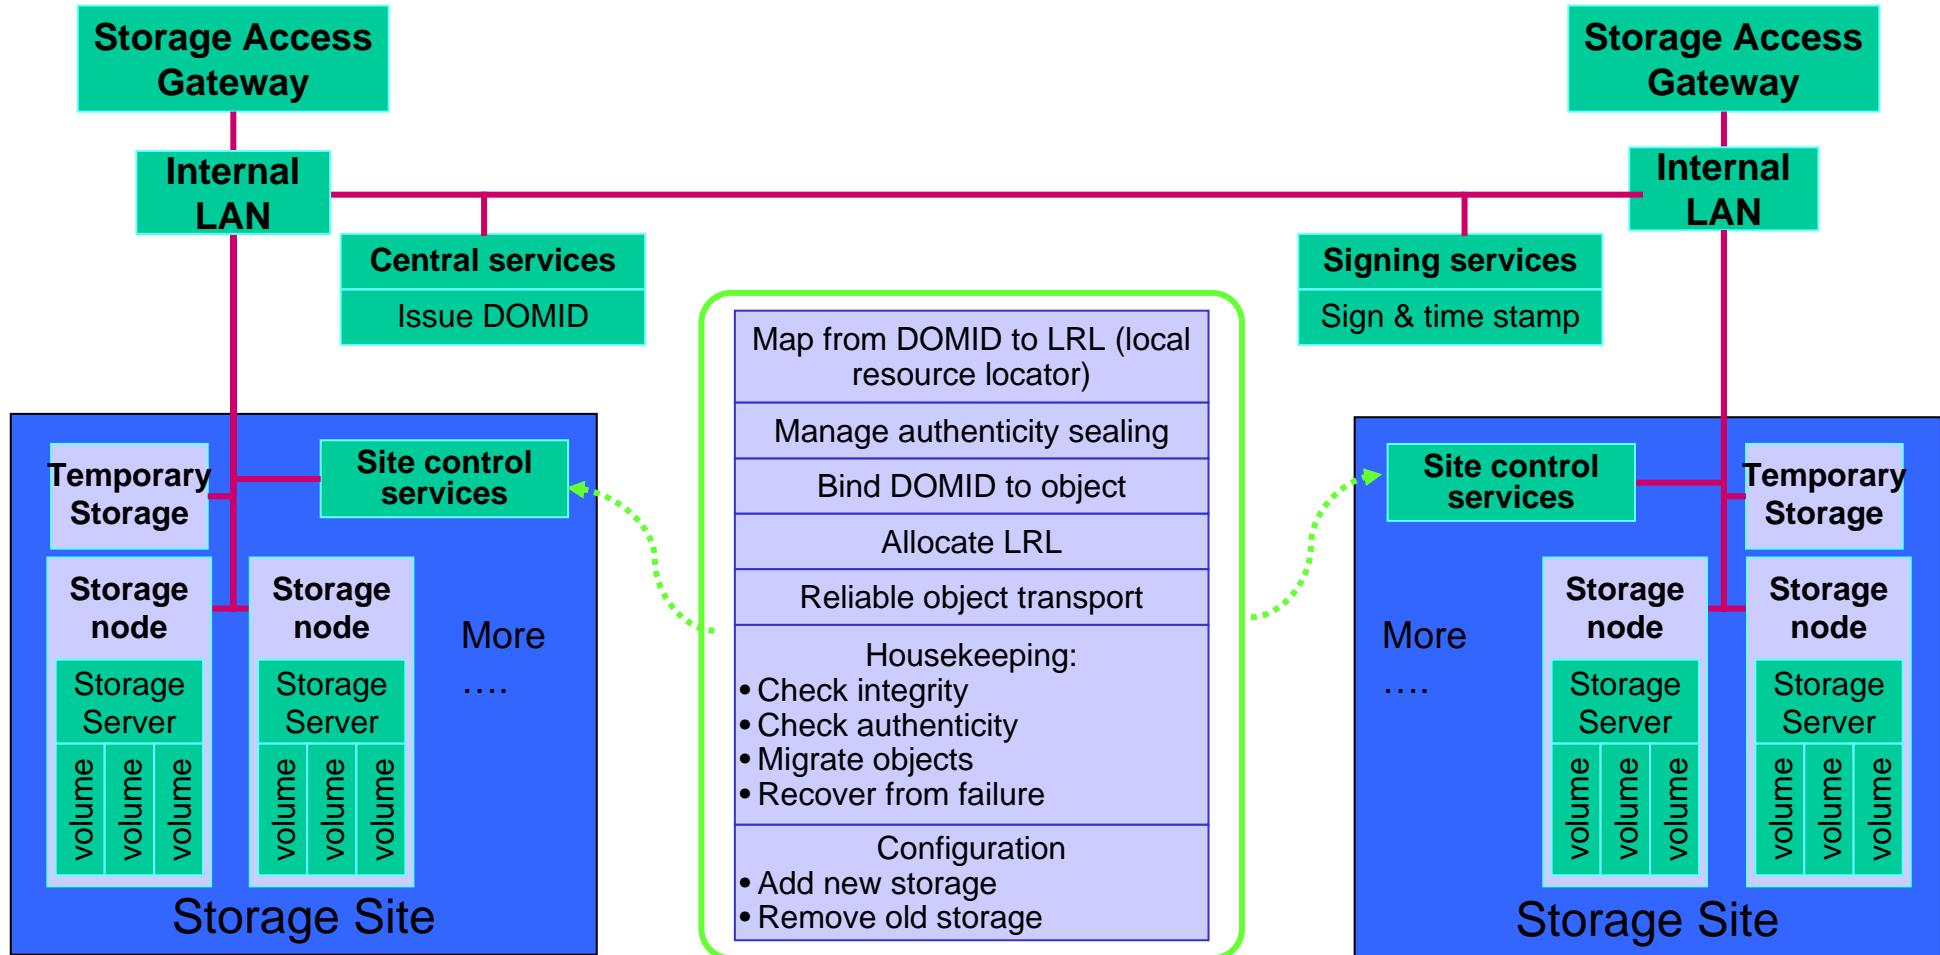
- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments



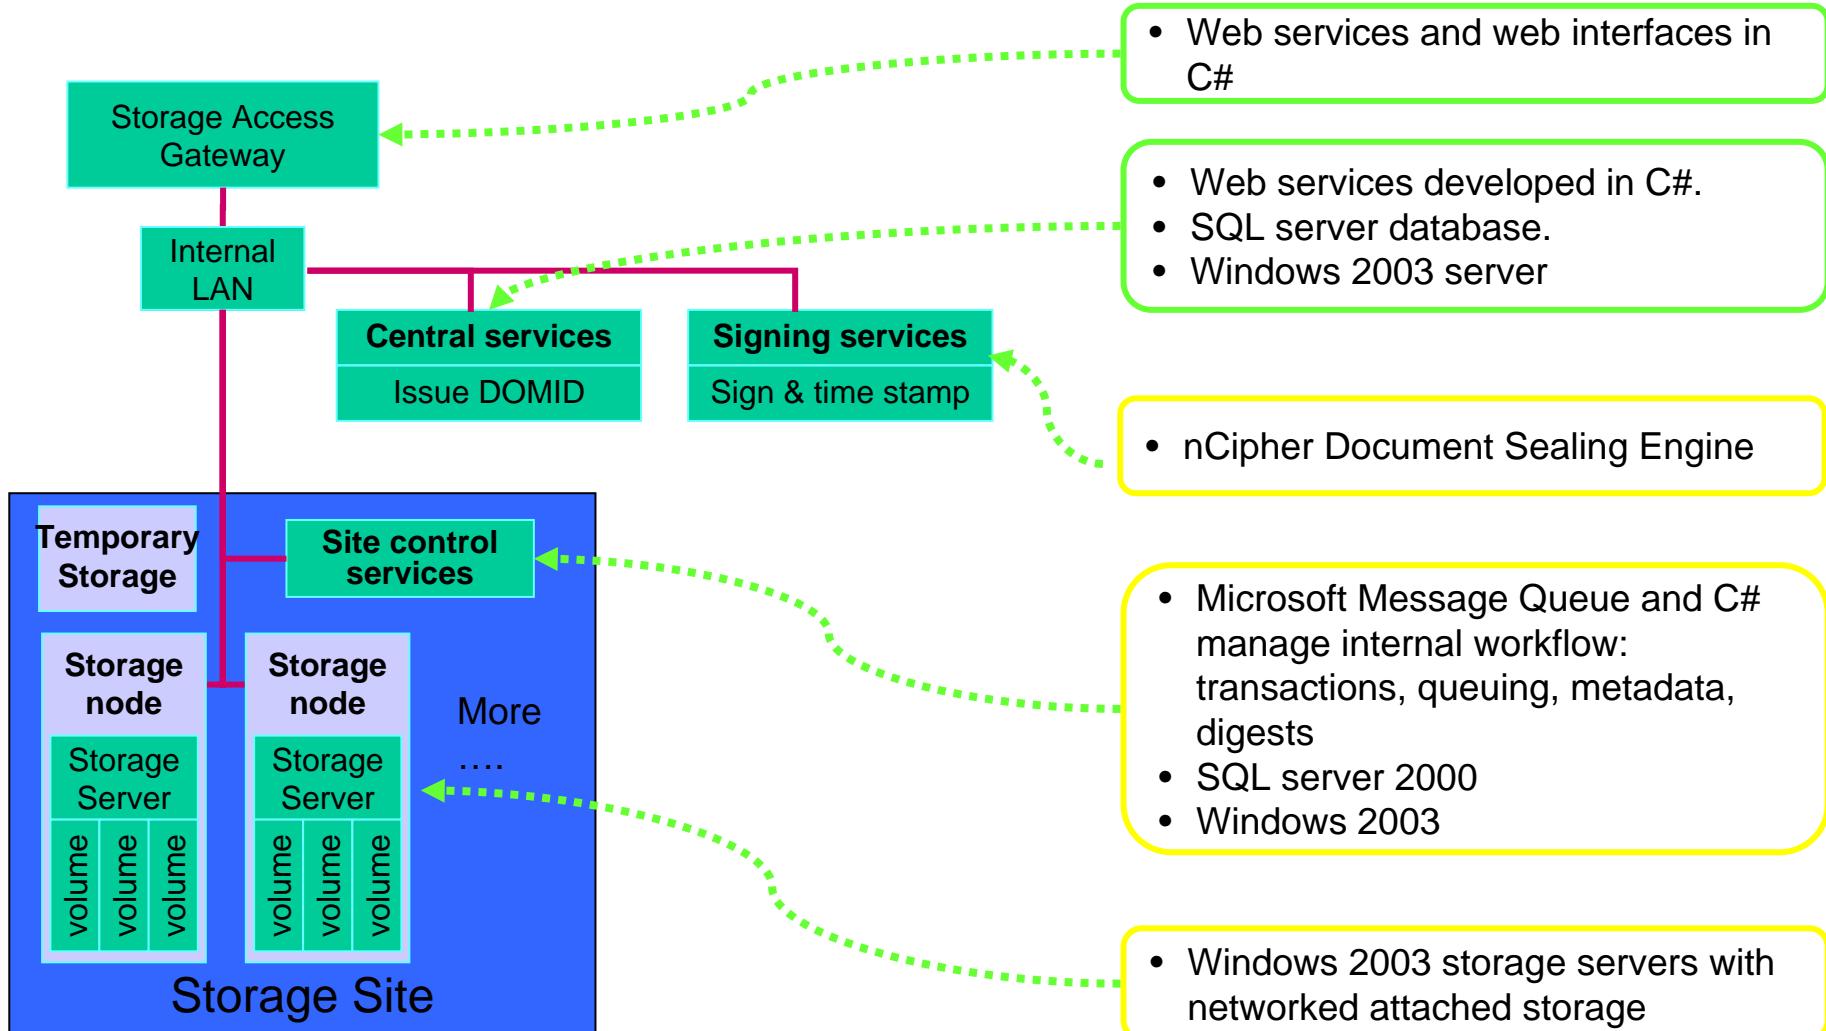
Design of storage service: multi-site architecture



Design of storage service: component design



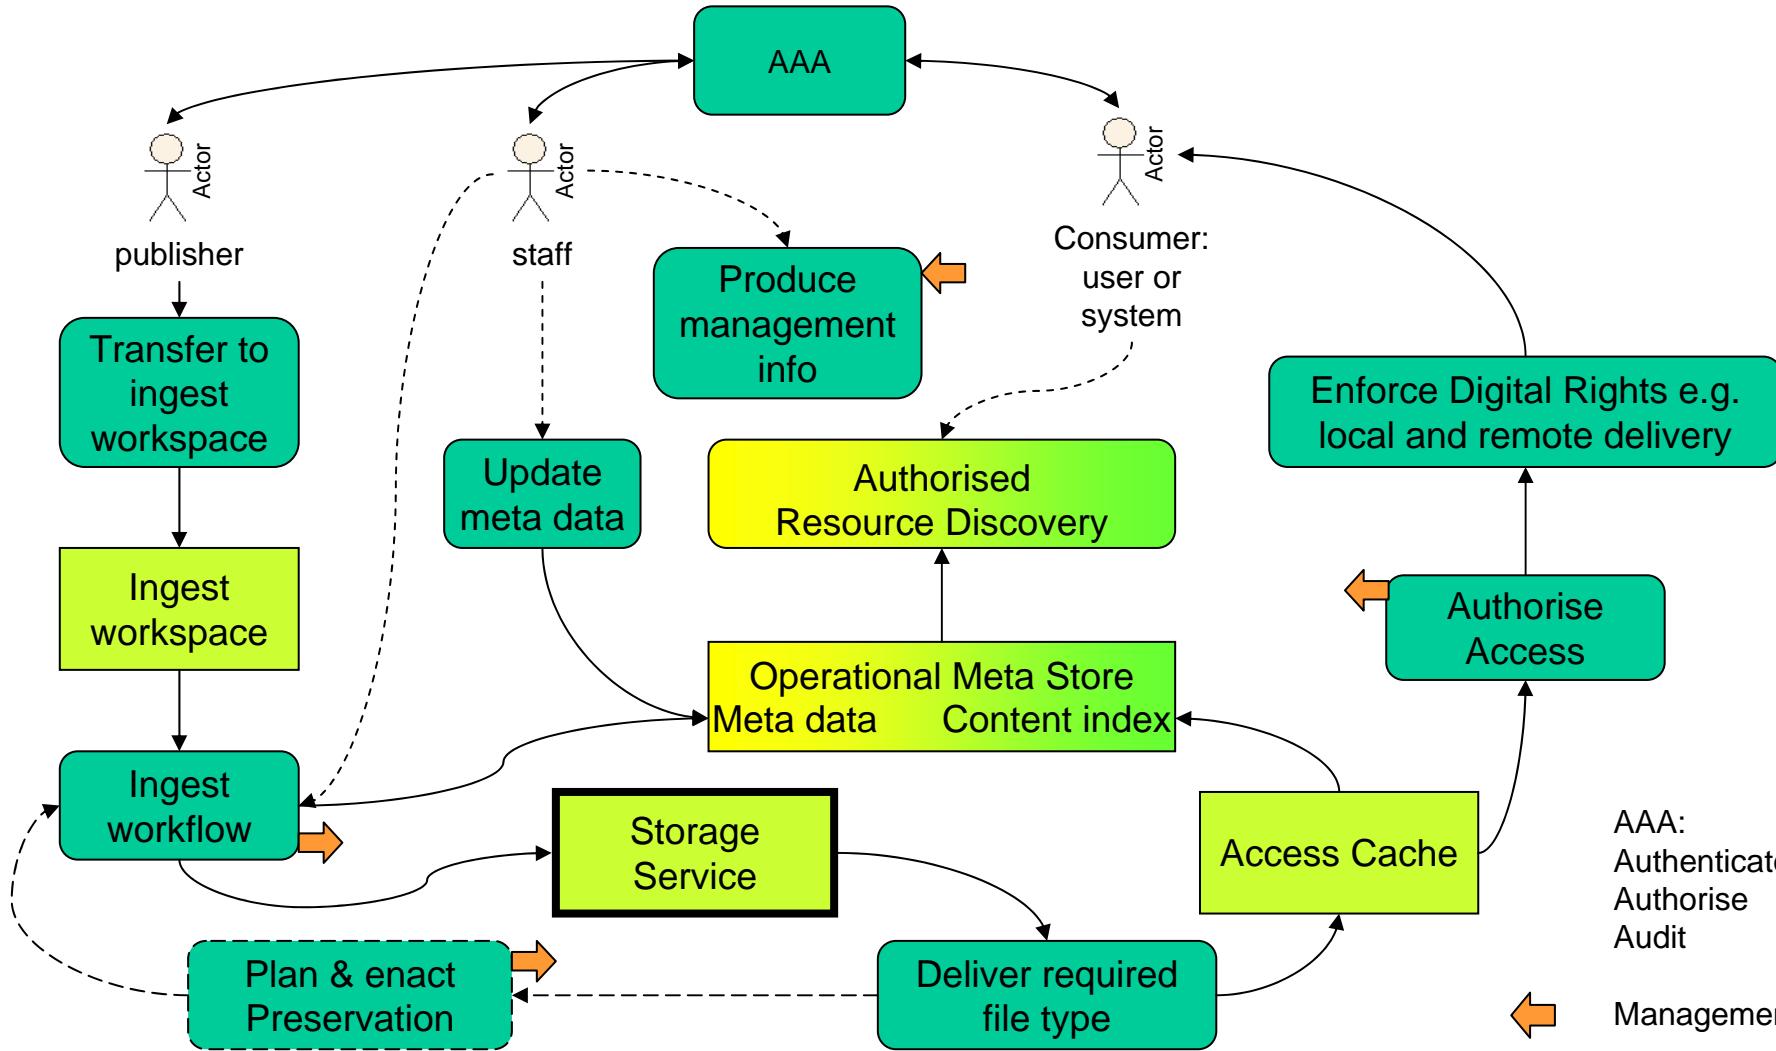
Design of storage service: implementation



Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- ■ Overall service architecture: ingest, storage and access
- Relation to other initiatives
- Concluding comments

Overall service architecture



Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- ■ Relation to other initiatives
- Concluding comments

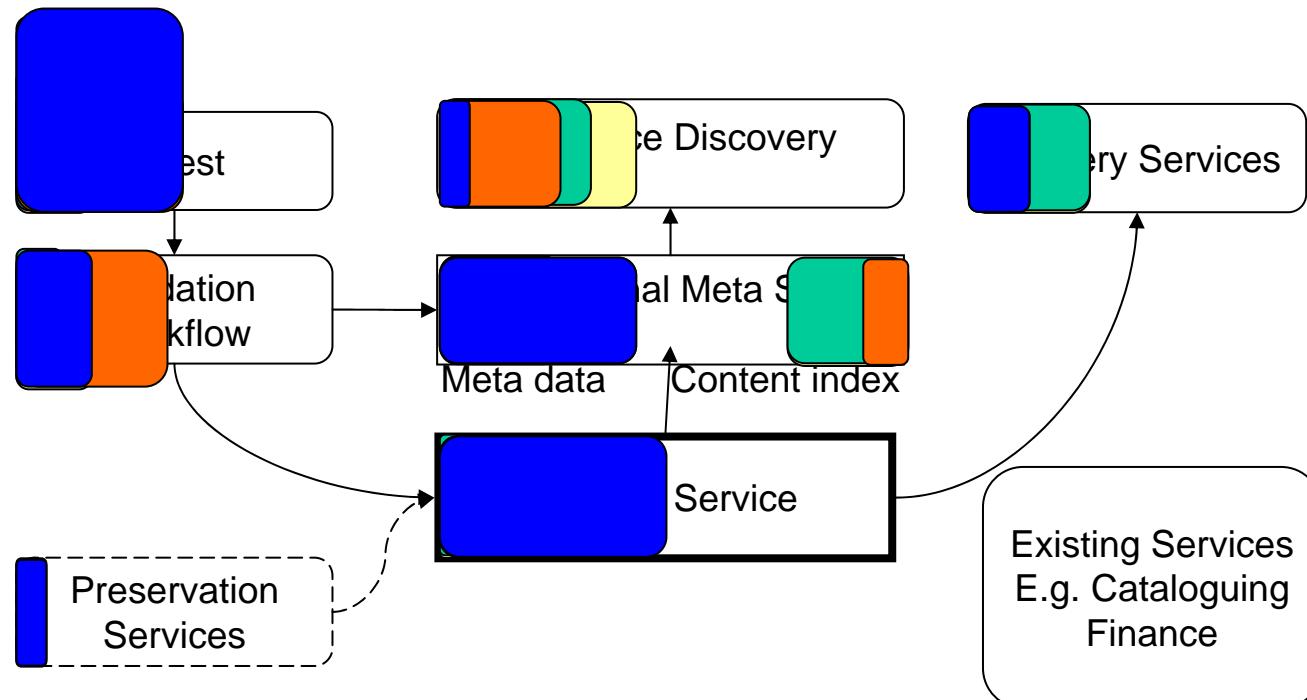
Relation to other initiatives: key DOM requirements

- Authenticity
- Integrity
- Low-cost scaleable storage
- Heterogeneous storage
- Disaster tolerance (multi-site)
- Integrate with BL Infrastructure
- Flexible extensible metadata
- Multiple metadata records per item
- Extensible ingest workflow
- Extensible storage workflow
- Scalable: number & size of objects
- Versioning for objects
- Versioning for metadata
- Bulk ingest
- Automated quality control
- Manual quality control
- Scaleable ingest performance
- Any of content type incl. compound objs
- Format validation and virus checking
- Support for digital preservation services

Relation to other initiatives: tabulated

Key DOM Requirements	DSpace	ePrints	Fedora	SRB
■ Authenticity	N	N	N	N
■ Integrity	Y	Y	Y	N
■ Low-cost scaleable storage	I	I	I	Y
■ Heterogeneous storage	I	I	I	Y
■ Disaster tolerance (multi-site)	I	I	I	Y
■ Integrate with BL Infrastructure	L	L	I/L	Y
■ Flexible extensible metadata	Y	Y	Y	Y
■ Multiple metadata records per item	N	N	Y	?
■ Extensible ingest workflow	L	L	L	L
■ Extensible storage workflow	N	N	L	L
■ Scalable: number & size of objects	L	L	L	Y
■ Versioning for objects	Y	N	Y	L
■ Versioning for metadata	N	N	N	N
■ Bulk ingest	L	L	Y	Y
■ Automated quality control	N	N	N	N
■ Manual quality control	N	N	N	N
■ Scalable ingest performance	L	L	L	Y
■ Any type of content incl. compound objects	L	N	Y	Y
■ Format validation and virus checking	L	L	Y	N
■ Support for digital preservation services	L	L	L	L

Relation to other initiatives: broad comparison



DSpace

ePrints

Fedora

SRB

Agenda

- Brief Introduction to the British Library
- Introduction to DOM programme: mission, vision, scope and drivers
- Principles: programme and system design
- Integrity and authenticity: conventional and digital
- Infrastructure: disaster tolerance and the storage market
- Design of storage service: multi-site
- Overall service architecture: ingest, storage and access
- Relation to other initiatives
- ■ Concluding comments

Concluding comments: further information

See the British Library Web pages under
'About us / Policies & Programmes'

<http://www.bl.uk/about/policies/dom/homepage.html>



Technology Watch Report

Related Publications:

"The large-scale archival storage of digital objects." Feb 2005

<http://www.dpconline.org/graphics/reports/>

The large-scale archival storage of digital objects

Jim Linden, Sean Martin, Richard Masters, and Roderic Parker
The British Library

DPC Technology Watch Series Report 04-03
February 2005

"Design for the Long Term: Authenticity and Object Representation"

<http://www.bl.uk/about/policies/dom/pdf/archiving2005l.pdf>

"Why Traditional Storage Systems Don't Help Us Save Stuff Forever"

http://www.stanford.edu/~candea/hotdep/papers/baker_forever.pdf

Concluding comments: any finally

- The British Library has a substantial long term vision for electronic resources
- We believe that we have the correct processes & design principles to build the capability we need
- The current budget is £1.5m per year
- We have made good initial progress although not without setbacks
- We have designed a highly scalable, fault tolerant, redundant, self-validating architecture
- We are seeking partners to collaborate in the more complex areas of the architecture, such as Digital Preservation & Digital Rights Management