IEEE Archive Storage Life Cycle Workshop

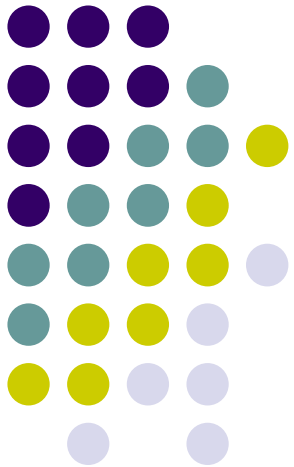

IEEE Computer Society

Mass Storage Systems and Technology

San Diego

September 24, 2007

IEEE COMPUTER SOCIETY

Workshop objectives

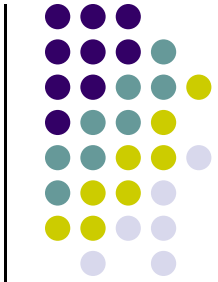

- Bring together experienced professionals representing data resource services providers and consumers in government, academia and industry to discuss, document and publish ideas and suggestions to improve data migration and operational effectiveness for large institutional data repositories built upon the file system metaphor (e.g., in excess of 40 petabytes). Also consider implications for metadata and life cycle policy management

- Discuss the notion that, for the file system metaphor, ownership and stewardship of data and the digital data media assets on which the data resides requires (a) an interface or (b) disclosure that provides all of the information (protocols, formats, etc.) that enable vendor neutral data and storage migration (e.g., move media instead of copying data to facilitate a move from one vendor's file system and storage services to another).

Agenda

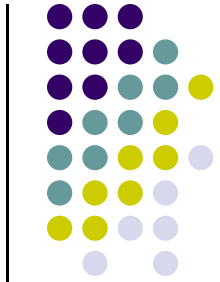

8:30 AM	Introduction and Welcome
	Discuss workshop objectives and challenge - Bob Coyne
9:30 AM	Data Migration and Policy Management - Reagan Moore
10:30 AM	Break
10:45 AM	SNIA on Data Migration - Michael Peterson
Noon	Lunch
1:00 PM	Group Discussion - Robert Chadduck
3:00 PM	Break
3:15 PM	Individual Responses – All
5:00 PM	Adjourn
	Develop and publish report (post workshop task)

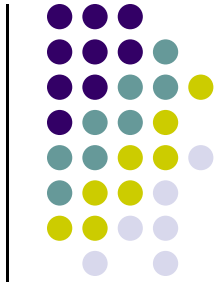

Workshop Terms, Problem, Challenge & Mission Environment

Differences between an Archive, Repository, and Records Management

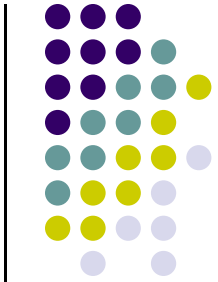

- Institutional Repository – A system for collecting, preserving, and disseminating mission content
- Archive – A collection of data that is maintained as a long-term record of a business, application, or information state. Archives are typically kept for auditing, regulatory, analysis or reference purposes rather than for application or data recovery. - SNIA
- Records Management – The systematic control of records throughout their life cycle. – ARMA

Provided courtesy of Millman and Thielen [1]

Workshop Data and Storage Services Spectrum

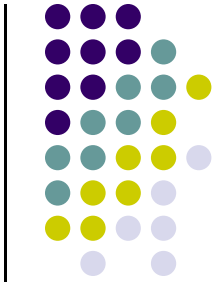

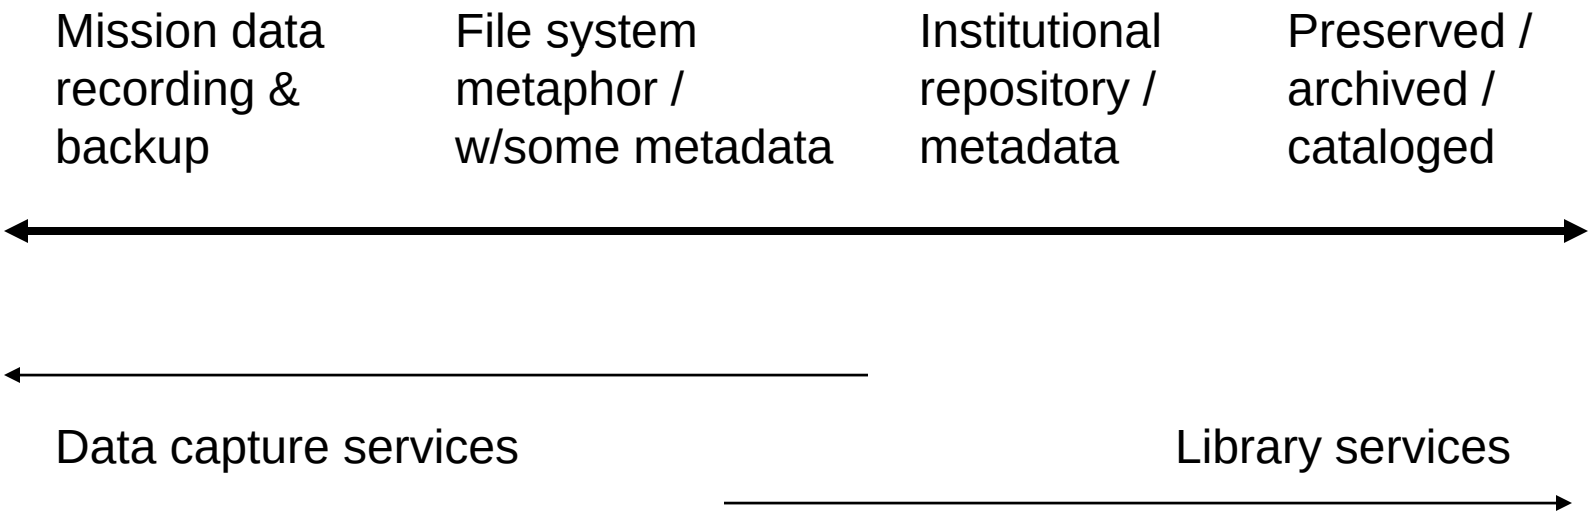

Provided courtesy of Millman and Thielen [1]

Workshop Discussion Focus 1

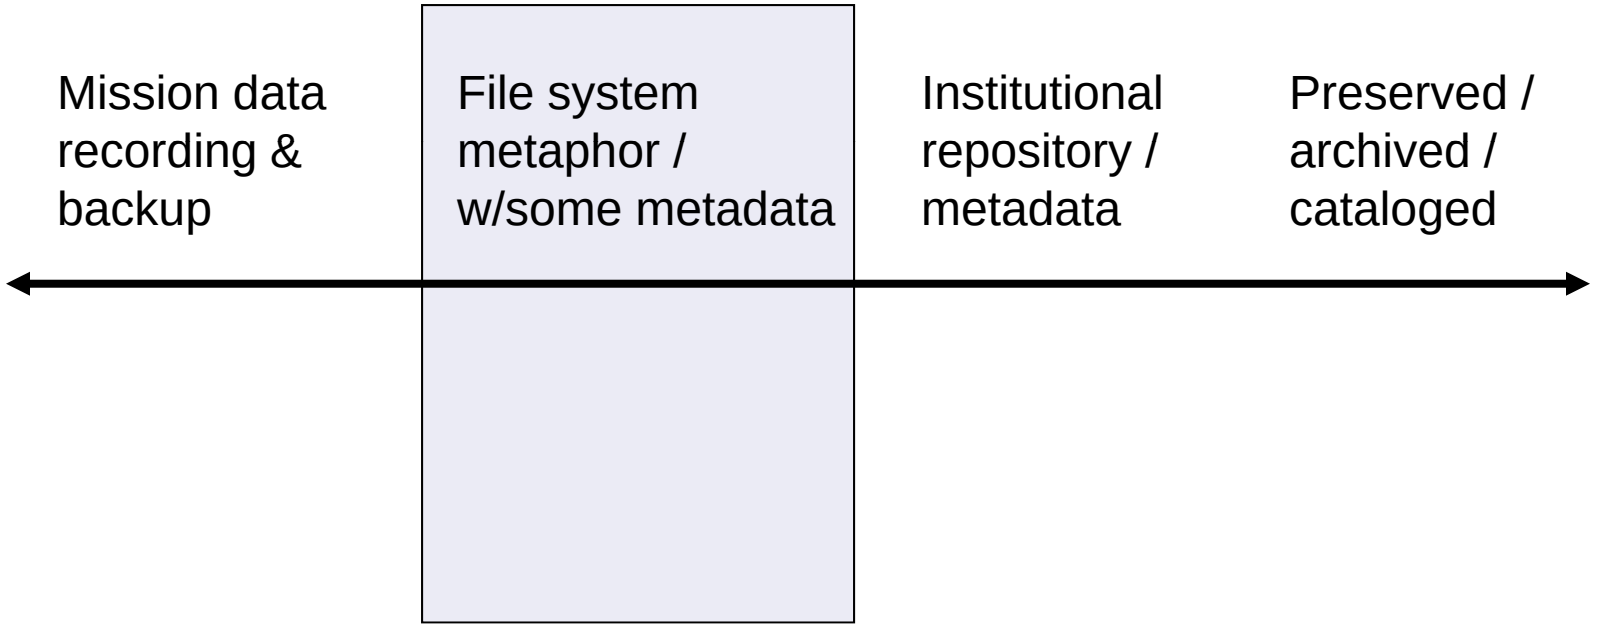

Bitfile Migration

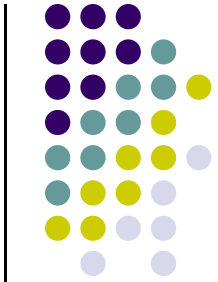

- Bitfiles [7] created using software dependent file and metadata formats and hardware dependent platforms can become inaccessible over time. The challenge here is to undertake proactive steps that ensure that bitfiles are accessible over time, independent of software and hardware platforms. [6]

- Logical and physical migration are only two dimensions of the preservation problem. The list of challenges is long and includes many technology and operations factors. [4]

Note: Bitfile services are content independent.

Bitfile Migration

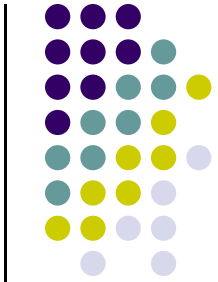

- NARA offers these rules for physical migration: First, use the most current storage technology, then if on disk, MIGRATE every 3 years, if on tape, MIGRATE every 5 years [4]

  - NASA Planetary Data Systems: …. no technology so far has proven itself to be affordable, reliable, and of high enough density to suit current needs. Because of this we have been planning not only multiple copies but also distant storage locations, ideally to prevent nearly any form of event or catastrophe from damaging more than one copy of the archive. We are also convinced that this alone is not sufficient. Regular verification and restoration of these archives is necessary to counter the inevitable effects of long-term storage. [11]

Bitfile Migration

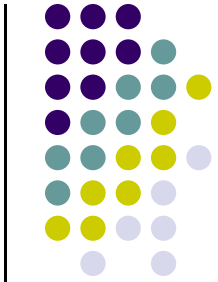

- How painful and costly is it to migrate a 8, 16 or 100 petabytes per year?
  - NCAR: A substantial portion of our user data has actually been migrated twice in the last six years …. data migration has already become a continuous task, often with multiple separate migrations running in parallel. More and more of our system resources are being used for data migration, which adversely affects users who are trying to read or write their MSS files …. we estimate that we will be getting close to five Petabytes of data …. We will need to sustain an aggregate migration data rate of one to two Petabytes/year in order to complete the entire migration in a timely manner. This is going to put a very large load on our already overloaded production system [9]

Bitfile & Storage System Migration

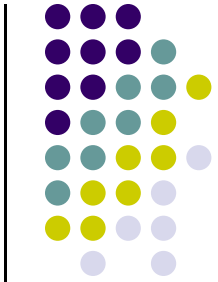

- Workshop challenge:
  - Maintain a policy of migrating data on tape to new removable media and/or new technology every 5+ years.
  - Move 40 PB of bitfile system data from vendor A's bitfile and storage services to vendor B's in less one year.
  - Move 40 PB of a second copy of bitfile system data from vendor A's bitfile and storage services to vendor B's in less one year.
  - Define implications on management policies during the migration (authenticity, integrity, disposition, retention).

Bitfile Migration Quantified

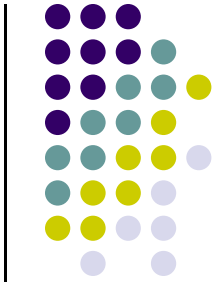

- Movement of a petabyte of data per year requires sustained data transfer of 32 MBytes/sec
  - Modern tape drives sustain 120 Mbytes/sec data transfer rates
  - Implies at least 20 tape drives are needed to continuously read and write 40 PB per year plus other necessary operating overhead, capacity and media.

Workshop Challenge: Mission Services Environment

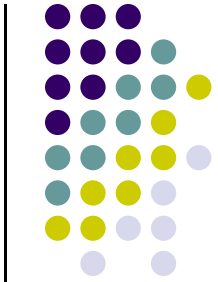

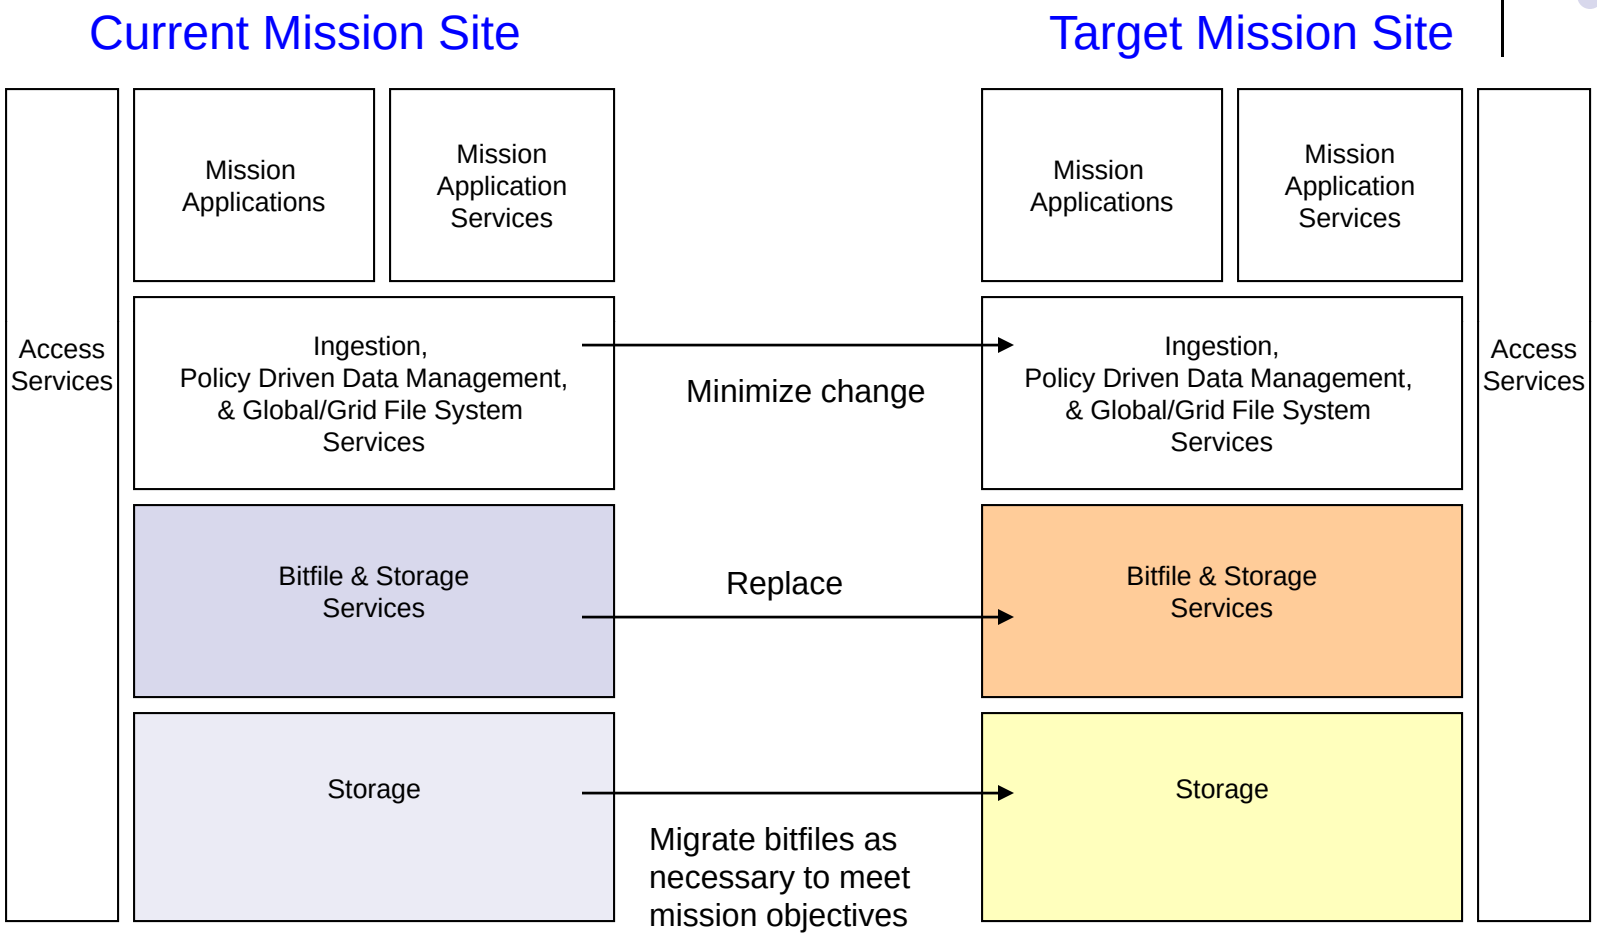

Workshop Problem & Challenge

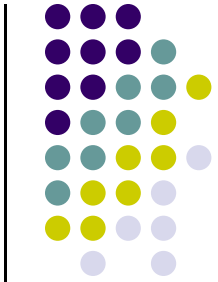

- Problem: Replace bitfiles and storage systems services managing 40PB of bitfiles and 40PB of second copy bitfiles from one vendor's file system and storage services to another in less than 1 year.
  - Preserve bitfiles, name space structure and file & storage metadata of potential mission value
  - Enable portability of mission data and provide file & storage services required to support migration of policy driven data management, content ingestion, global metadata and access services to new technology and locations with the least impact practical
  - Sustain operational effectiveness
- Challenge: Move, do not copy, removable media containing 40PB primary bitfiles and 40PB second copy bitfiles from one vendor's file system and storage services to another in less than 1 year.

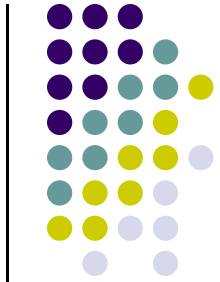

Improving Data Migration Effectiveness

Workshop Discussion Focus 2

Improve operations effectiveness
Integrate IT & IM practices

Represent & manage media content and the context in which media is used

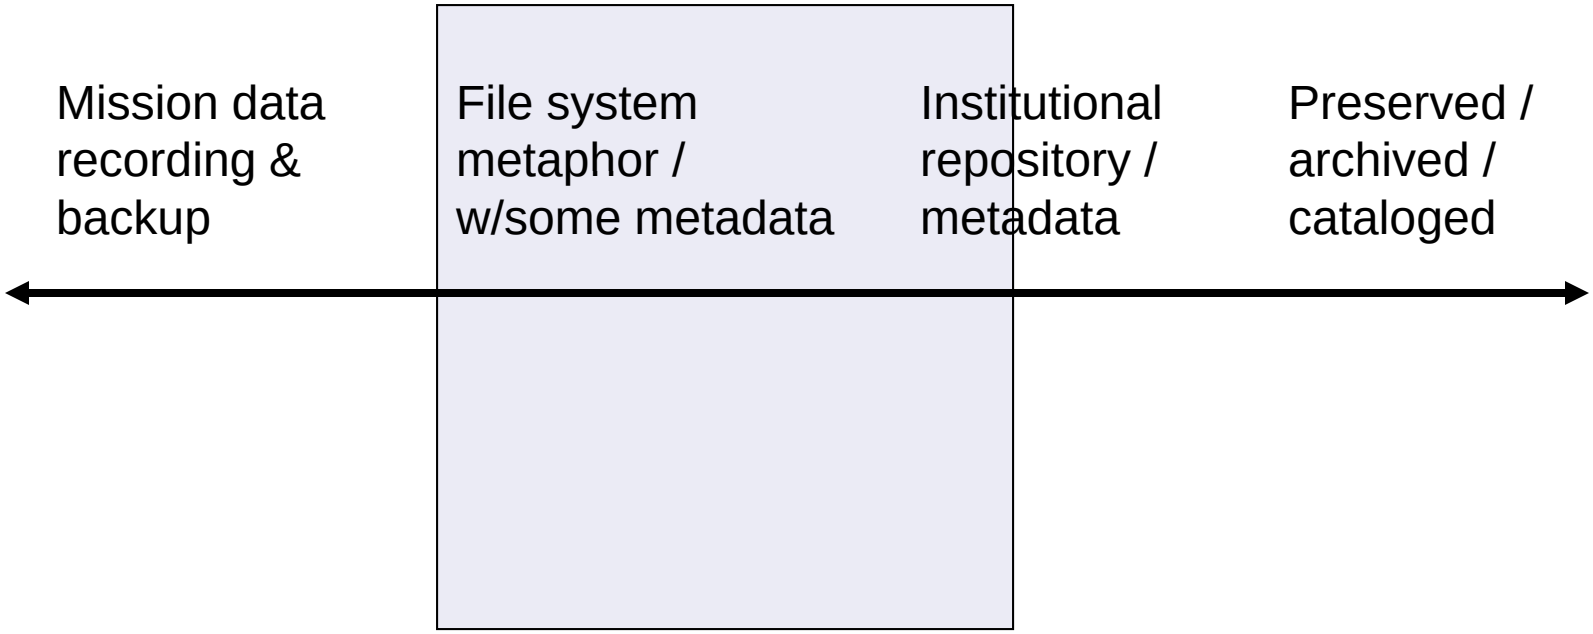

Representation Information for Data Management Systems

adapted from integrated Rule-Oriented Data System [3]

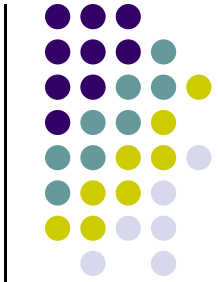

- It is not sufficient to characterize only the media structure for media migration between vendors
- Also need to characterize the management policies and processes of the original system - comprises **representation information for the system**
  - Management processes map to the operations performed upon the media
  - Management policies map to rules that control operations performed upon the media
- The context in which the media is used is as important as the content in the media

Simple Site Repository Model

adapted from OAIS functional model [5]

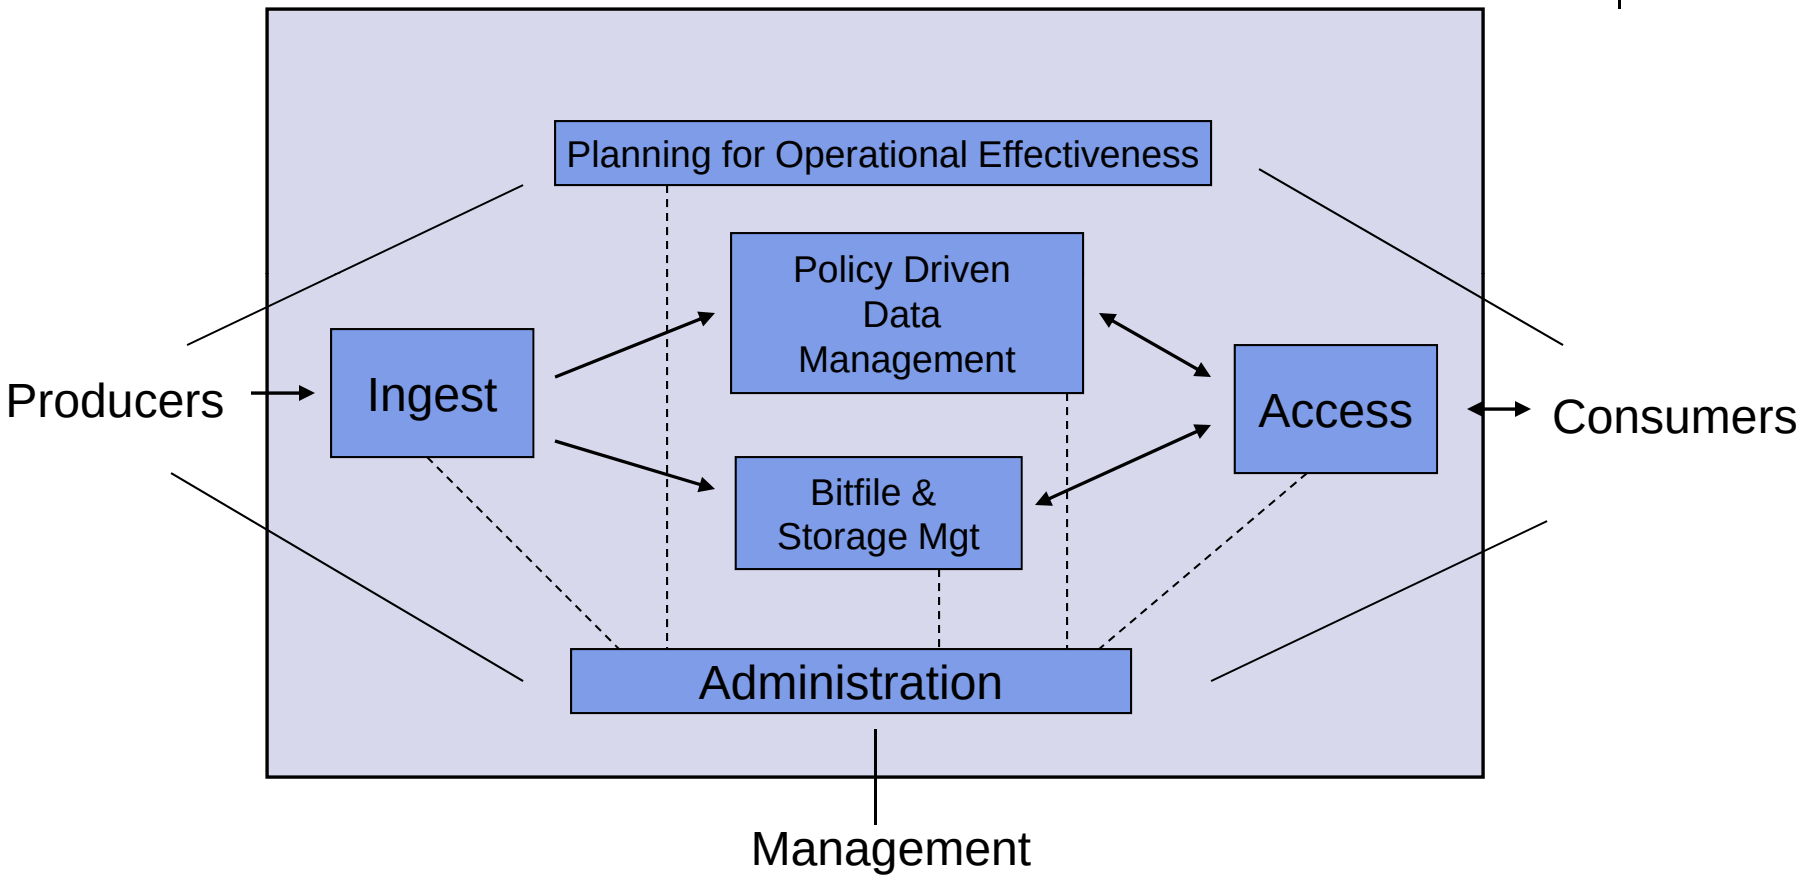

Attendees' Suggestions to Improve Operational Effectiveness

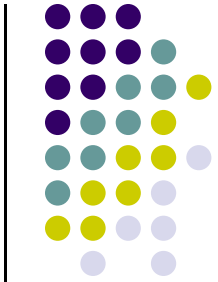

- On-going refinement and automation of file life cycle policies [2]
- Deploy policy driven and rule based data management [3]
- Undertake proactive steps that ensure that bitfiles are accessible over time, independent of software and hardware platforms. [6]
- Control/schedule data lifecycle management activities such as data expiration, migration, replication to minimize performance impact on mission applications. [8]
- Build-in the migration of all of the mission or business digital data repository's technology, hardware and software. [12]
- Define best practices and industry standard(s) for encryption and key management for very large heterogeneous repositories [16]

Attendees' Suggestions to Improve Operational Effectiveness

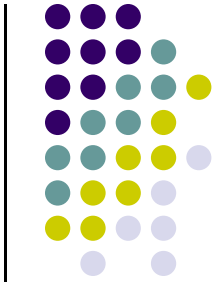

- Produce a "best practices for long-term digital information retention" reference model that covers the information-storage domain – the technology domain unaddressed in all existing 'archive' standards and best practices such as ISO 14721:2002 - "Open Archival Information Systems", OAIS or the Sedona Conference [10]

- Integrate ILM-based practices into the long-term digital information retention process to sustainably automate IT infrastructure in support of business and information requirements [10]

- Define reference models and possible technical standards that solve and provide for scalable physical and logical migration – the two 'big challenges' of preservation [10]

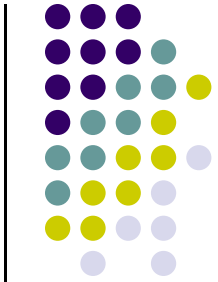

Leveraging archival and records management initiatives for large institutional data repositories

To what extent?
Bulk data migration?
Aggregate & migrate a bitfile namespace?

File System Metaphor

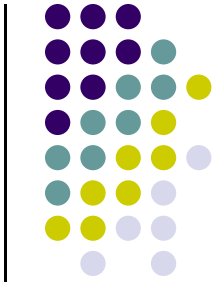

Comments from XAM advocate Chuck Hollis, EMC [13]

- There's a very strong case to have certain kinds of metadata stored with the file or object itself in a permanent and inseparable kind of way. If there's a security concern with a file or object, you'd probably like to embed that sort of information within the object itself, rather than trust an external repository that may or not be around.

- "The Filesystem Problem"
  - Simply put, the technology is not keeping up with requirements. Today, we work with people who are already responsible for billions of stored objects, with more coming.
  - Creating lots of little file systems (even TB-sized ones), and managing them all just won't work. Anyone want to predict how long a PB-sized fsck (file system check) would take? Or how many filesystems you'd need to store a petabyte? Or how you'd manage all of that at an application or infrastructure level?
  - Having applications store pathnames to files just isn't workable at scale either.
  - What you'd like to have is what I call a "claim check API".

XAM, ILM and Vendor Plug & Play

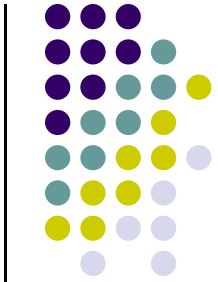

Comments on XAM from Jim Armstrong, LasCon Storage [14]

- SNIA specifications will mean that ILM metadata will be made vendor independent in terms of location and field availability, so ILM policies can be applied as standard over any compliant vendor's hardware or software. This will be especially important to avoid vendor locking for long term data retention.

ILM, according to SNIA-DMF [15]

- ILM is described as ‘information-based management’. Or more expanded, a ‘standards-based, business driven management practice’. It is about establishing requirements and operating to standard practices and service levels. By applying retention, placement, protection, business continuity, archive, and compliance policies to information, based on its value, as it is created and as it ages the practices can be automated and operations costs reduced. Many new tools are available now to implement automated ILM polices and many others will come to market next year.

SNIA Next Generation Storage Repository

Origin OSD Use Cases, Seagate, October 2006

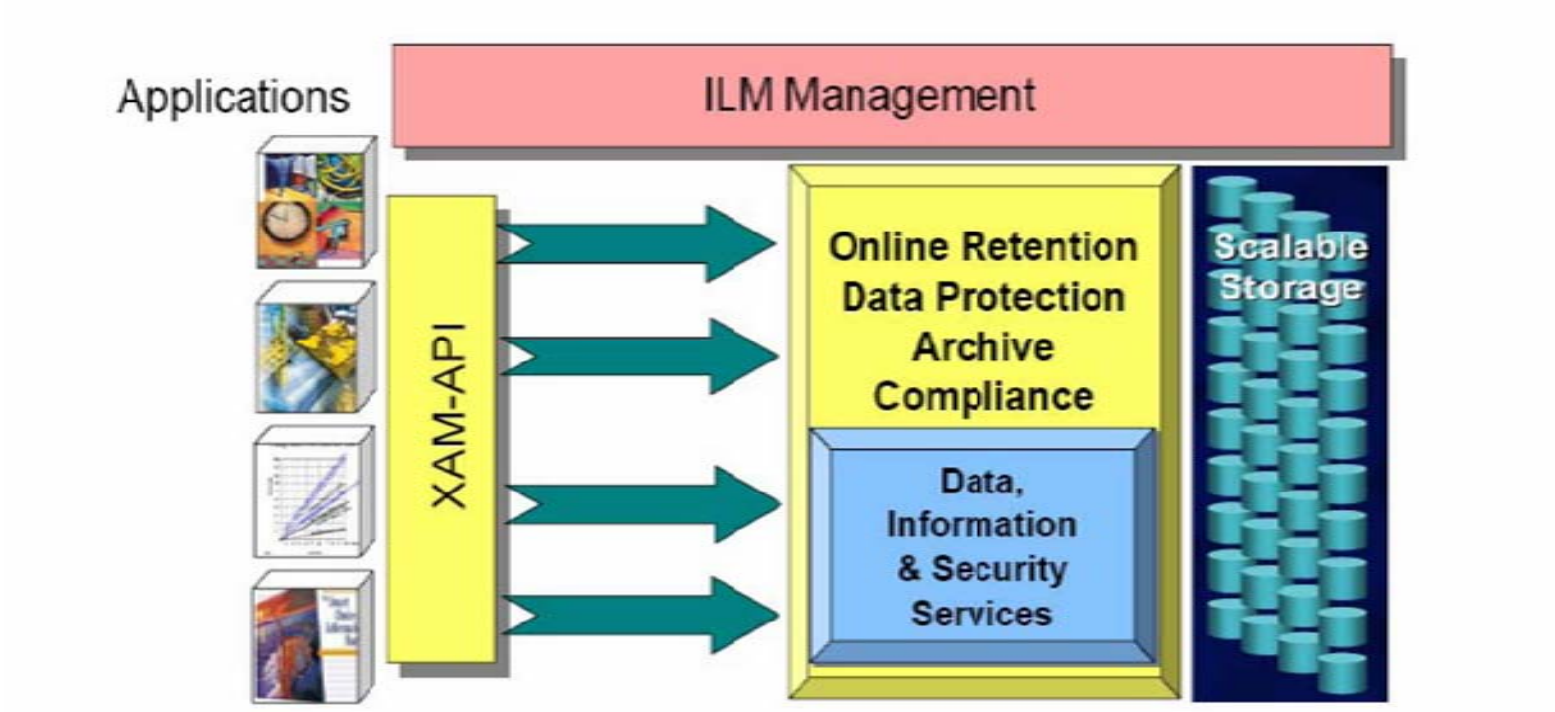

XAM – How it Works

Origin Christina Casten, t*he Power of XAM,* Chair XAM Committee, SNIA, July, 2006

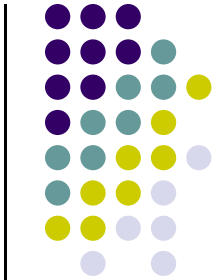

- The XAM Interface specification defines a standard interface (access) method between “Consumers” (application and management software) and “Providers” (storage systems) to manage information. XAM annotates objects with metadata providing for the management of information at a semantic level. This coupling allows external ILM-based policy services to make intelligent decisions about the management of objects without referring back to the application and without impacting the application.
- As an interface, XAM abstracts the access method from the storage. This supports the mobility of information independent from the storage to allow longevity, distribution, and management of information. The XAM Interface is intended to achieve interoperability, storage transparency, and automation for Information Lifecycle Management-based practices, long-term records retention, and information assurance (security).

Data Preservation Widgets & OAIS Functional Model

Adapted from Dalit Naor, Preservation DataStores:
An architecture for Preservation Aware Storage, IBM Research, Haifa, 2007

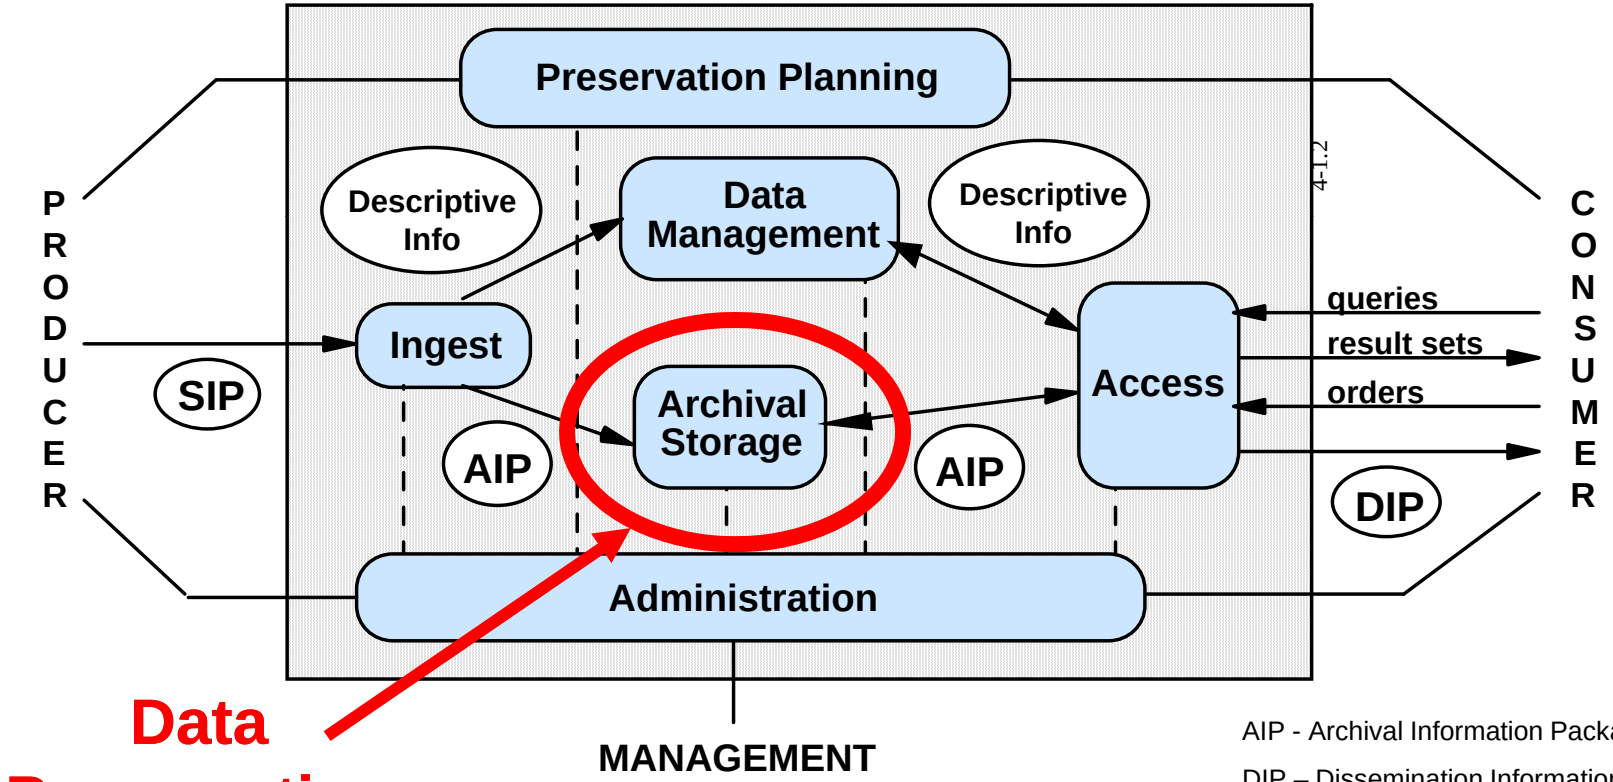

Data Preservation Widget

AIP - Archival Information Package

DIP – Dissemination Information Package

SIP - Submission Information Package

Data Preservation Widget

Adapted from Dalit Naor, Preservation DataStores:
An architecture for Preservation Aware Storage, IBM Research, Haifa, 2007

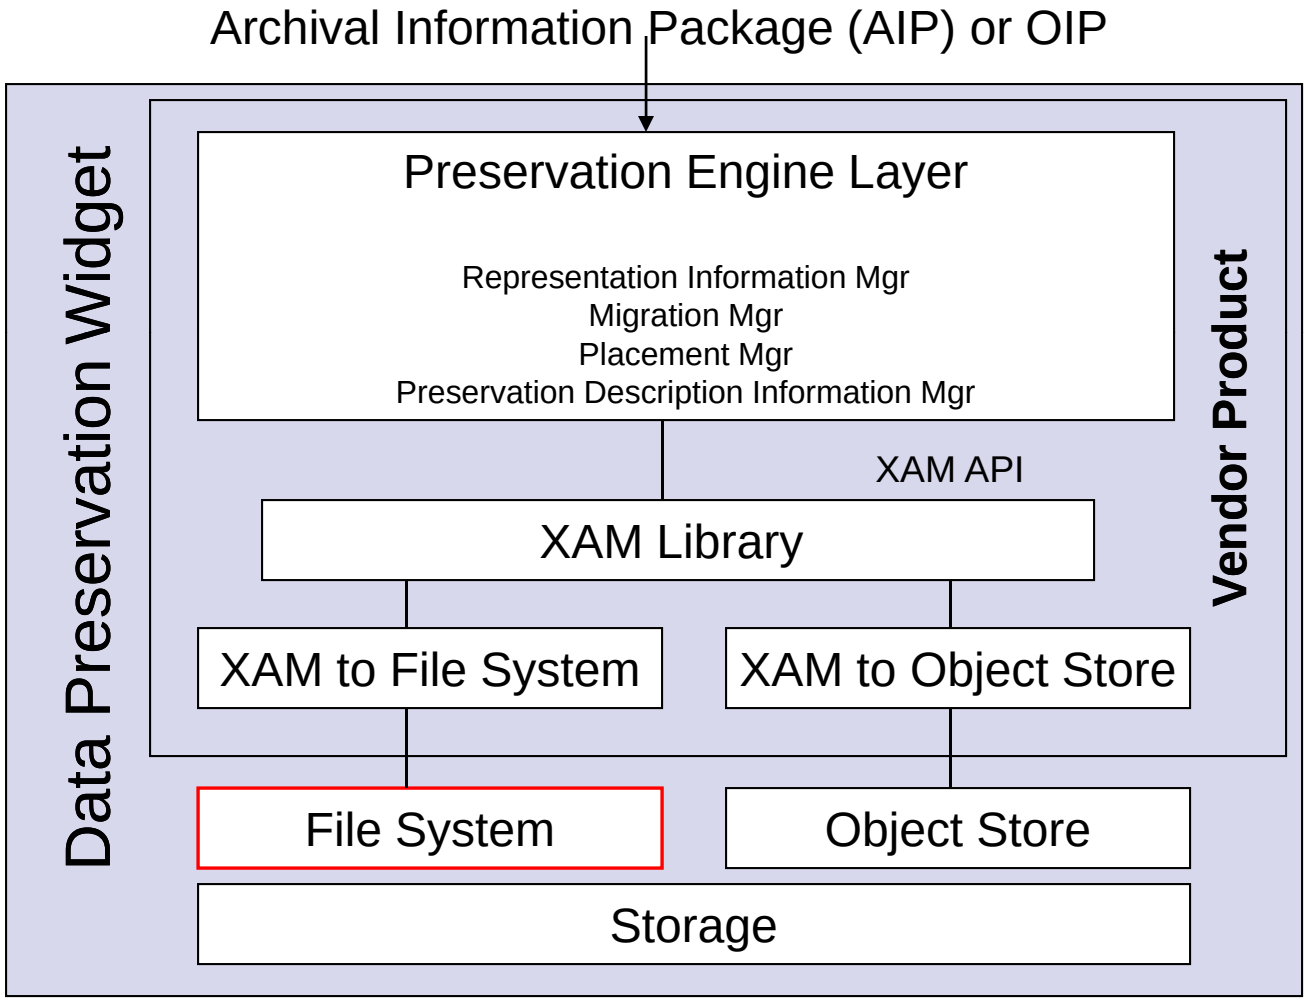

Simple Site Repository Model

adapted from OAIS functional model [5]

File systems developers continue to add ILM functionality and will continue to integrate IT with IM

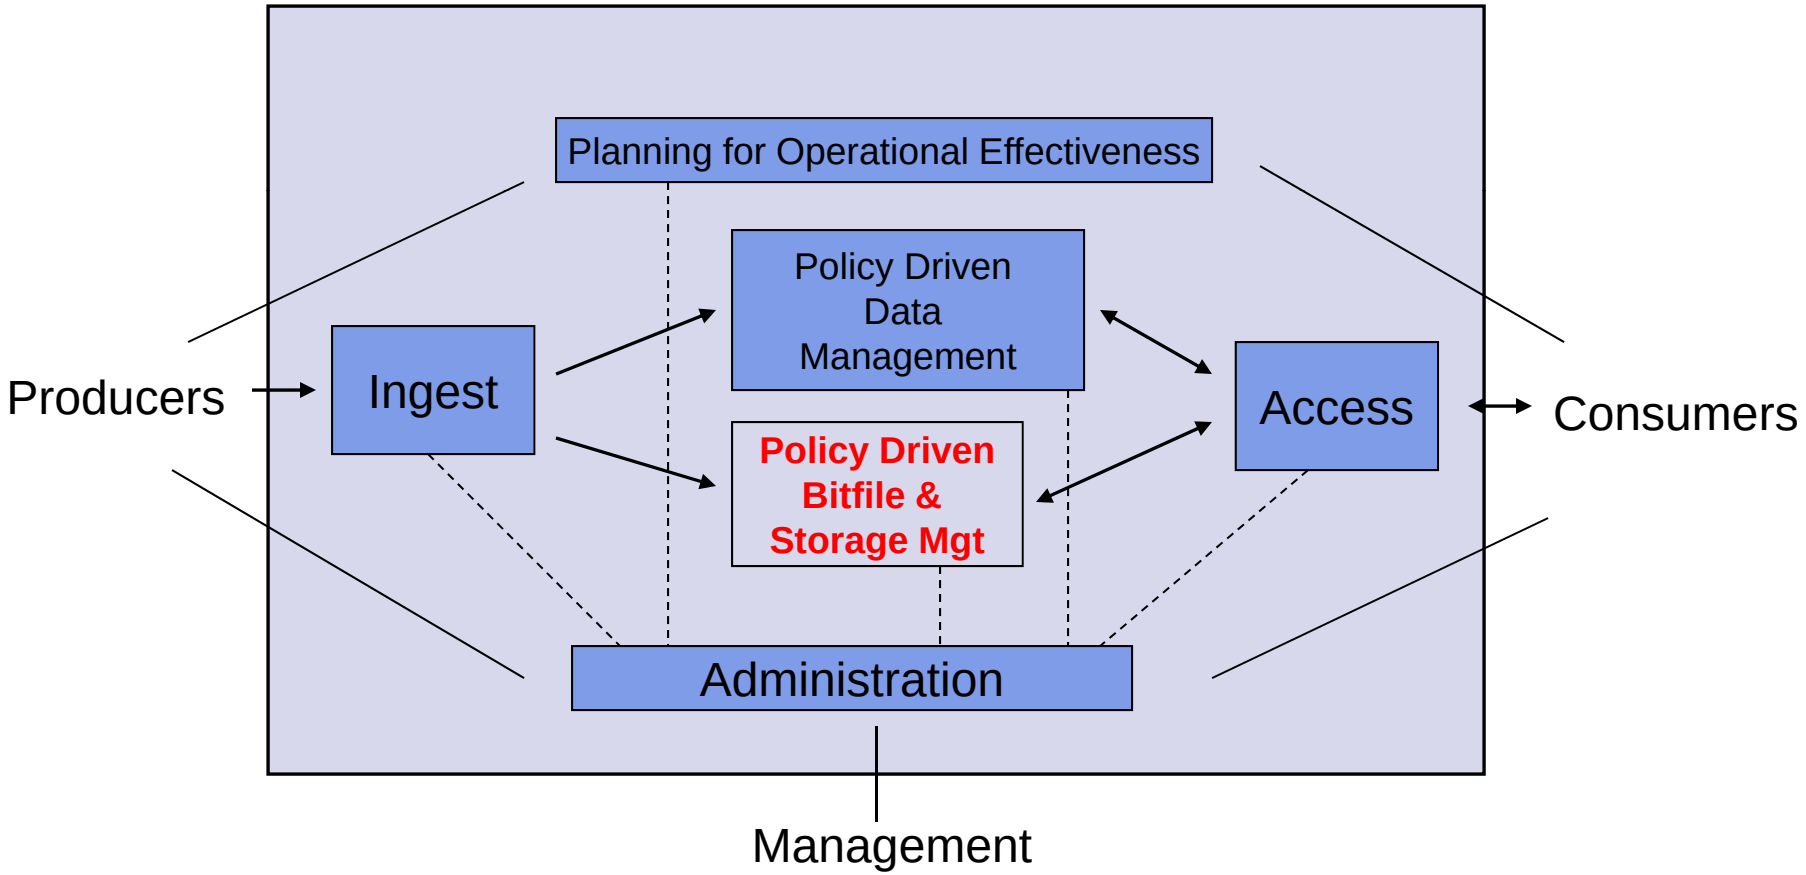

Acknowledgements

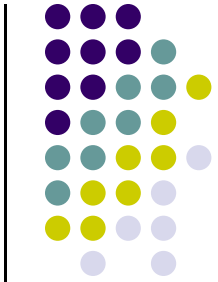

[1] David Millman, Columbia University, and Ron Thielen, University of Chicago, *Archival, Digital Preservation, and Records Management*, Spring 2006 Common Solution Group, May 12, 2006

[2] Robert H. McDonald and Chris Jordan, *Archival Storage Lifecycle Management in High Performance Computing Data Centers,*
San Diego Supercomputer Center, July 2007

[3] Reagan W. Moore , *Large-scale Data Management*, San Diego Supercomputer Center, July 2007

[4] Michael Peterson, *The Coming Archive Crisis,* SNIA - Data Management Forum, November 16, 2006

[5] Reference Model for an Open Archival Information, RECOMMENDATION FOR SPACE DATA SYSTEM STANDARDS, System (OAIS) CCSDS 650.0-B-1, BLUE BOOK, January 2002

[6] Nadi Mirvahabi, Archival Preservation and Mass Storage at the U.S. National Archives, NARA, May 2007

[7] IEEE Storage System Standards Working Group, *Reference Model for Open Storage Systems Interconnection—Mass Storage System Reference Model Version 5*. New York: IEEE, September 1994

[8] Lan Huang, *Non-Intrusive Information Lifecycle Management*, IBM Almaden Research Center, July 2007

[9] John Merrill, *Data Migration at NCAR*, NCAR MSS Group, June 2007

[10] Gary Zasman, *Comments on IEEE Archive Storage Life Cycle Workshop topics*, SNIA 100 Year Archive Task Force, July 2007

[11] Bill Harris, *Past Failures in Archive Maintenance*, NASA Planetary Data Systems, July 2007

[12] Dave Cavena, Chris Wood, Jeff Bonwick, Guy Steele, Michael Selway, *Archiving Movies In a Digital World* ©, Version 2.1, Sun Micorsystems, June 2007

[13] Jim Armstonrg, LasCon, http://www.lascon.co.uk/d004000.htm

[14] Chuck Hollis, EMC, Why XAM Is Very Very Cool, http://chucksblog.typepad.com/chucks_blog/2007/06/why_xam_is_very.html, June 2007

[15] ILM and Tiered Storage, http://www.drunkendata.com/wp-content/TSEDIT.pdf

[16] Dennis Kuehn and Marc Peters, D*ata Stewardship in Very Large Data Environments,* Boeing, September 2007

Acknowledgements

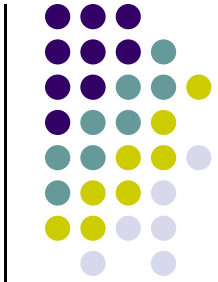

Workshop Organizing and Program Committee

Bob Coyne –Vice Chair for Metadata Meetings,
IEEE MSSTC Executive Committee, coyne@us.ibm.com

Reagan Moore – Member,
IEEE MSSTC Executive Committee, moore@sdsc.edu

Robert Chadduck, NARA, robert.chadduck@nara.gov

Harriet Coverston, Sun, harriet.Coverston@Sun.COM

Harry Hulen, IBM, hulen@us.ibm.com

Workshop Sponsor

IEEE Mass Storage Systems & Technology Technical Committee

Merritt Jones – Chair, Executive Committee

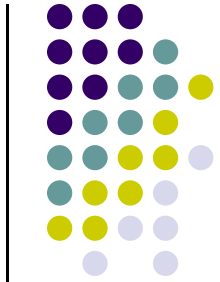

Backup Charts

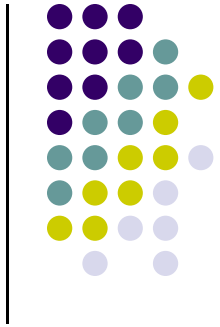

Industry, Government and Initiatives Scan

Provided courtesy of Millman and Thielen [1]

State of the IT Industry

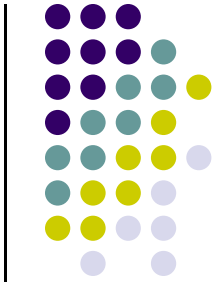

- Used to be all about compliance
- Increasing awareness that there are other reasons for archival
- Scan of IT Industry Organizations
- Scan of IT Vendors
- Scan of Government Initiatives
- Scan of Higher Education Initiatives

IT Industry Organizations

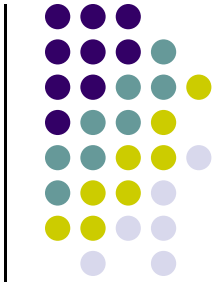

- SNIA (Storage Network Industry Association) Data Management Forum (DMF)
  - LTACSI (Long Term Archive and Compliance Storage Initiative)
    - 100 Year Archive Task Force
    - SDDF (Self Describing Data Format) Task Force
- ARMA - Association for Records Managers and Administrators (aka RIM Professionals) – Working with the SNIA
- AIIM – Association for Information and Image Management – Believes that ISO adoption of PDF/A is the way to address preservation

Survey of Government Authorities and Initiatives

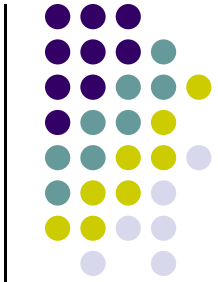

- LOC “Library of Congress”
- NARA “National Archives and Records Administration”
- NDIIPP “National Digital Information Infrastructure and Preservation Program”

Survey of Higher Education and Library Initiatives

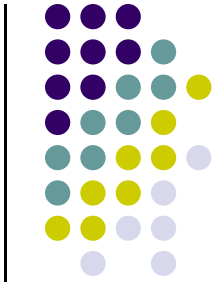

- DSpace *(an institutional repository, not an archive)*
- FEDORA *(ditto)*
- Stanford LOCKSS *(Lots of Copies Keep Stuff Safe)*
- DAITSS *(Dark Archive in the Sunshine State)*
- NEDLIB *(Networked European Deposit Library)*
- JORUM *(repository service, U.K.)*
- Columbia *(DSpace pilots; FEDORA in Socioeconomic Data Center Long-Term Archive)*
- CDAD *(Chicago Digital Archive Depository)*
- RLG Digital Repository Certification
- UCSD / SRB (Storage Resource Broker)
- JHOVE *(Harvard--object validation service)*

Standards
(formal, ad-hoc, and otherwise)

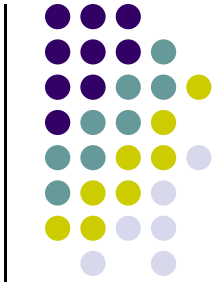

- OAIS “Open Archival Information System”
- PREMIS “Preservation Metadata Standard”
- METS “Metadata Encoding and Transmission Standard”
- EAD “Encoded Archival Description”
- MADS “Metadata Authority Description Schema”
- MODS "Metadata Object Description Schema"
- DOD 5015.2 “Design Criteria Standard for Electronic Records Management Software Applications”
- ISO 15489 (Records Management)
- and on … and on … and …

Scan of IT Vendors

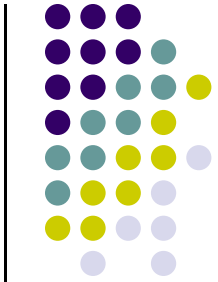

- Niche
  - Archivas, Permabit, Yosemite
- 800 lb Gorillas
  - HP, IBM, EMC, Sun (StorageTek)
- “Archival” Vendors
  - Commvault, Zantaz, ZipLip, iLumin, …

Five Steps to Archival

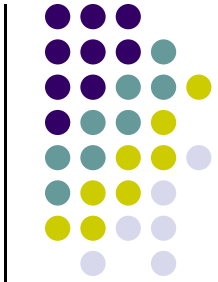

- Backup and Data Recording - a backup is not an archive, but backup and capture processes, support personnel and infrastructure may (or may not) support parts of the archival infrastructure
- Simple Bitstream Preservation - keep from losing the information; adds fixity checking, digital media asset management to backup
- Records Management - adds policy based classification and information life-cycle management
- Intellectual Content Preservation - keep the format current; migrate (or emulate) formats & structures
- Archival - adds bibliographic and administrative metadata

Provided courtesy of Millman and Thielen [1]

Types of Data Migration

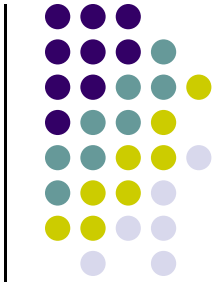

- **Refreshment -** effect is to replace a media instance with a copy that is sufficiently exact that all mission hardware and software continues to run as before. [5]
- **Replication**: no change to the repository information management data (e.g., packaging and content information). Refreshment is also a Replication, but may require changes to the storage mapping infrastructure. [5]
- **Repackaging**: some change in the bits of the repository information management data. [5]
- **Transformation**: some reversible or non-reversible change in the repository information management data while attempting to preserve the full information content. [5]