Education

Storage Tiering for File Systems and NAS

Ashvin Kamaraju, Symantec Corporation

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
- The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

Storage Tiering for File Systems and NAS

- This session will appeal to Data Center Managers and those that are seeking a fundamental understanding of storage tiering and the various approaches to implementing storage tiering and life cycle management of data stored in file systems and NAS filers.

  This session will survey the various storage tiering technologies available today, in file systems and NAS systems, and the policies that automate tiering. It will cover HSM, file systems that can inherently handle multiple classes of disk storage in a single name space/file system and File Area Network (FAN) switches that can migrate data between NAS systems. The objective of the session is to provide the attendees with the benefits of storage tiering and the policies that automate tiering to manage costs.

The expanding digital universe

Information created, captured and replicated

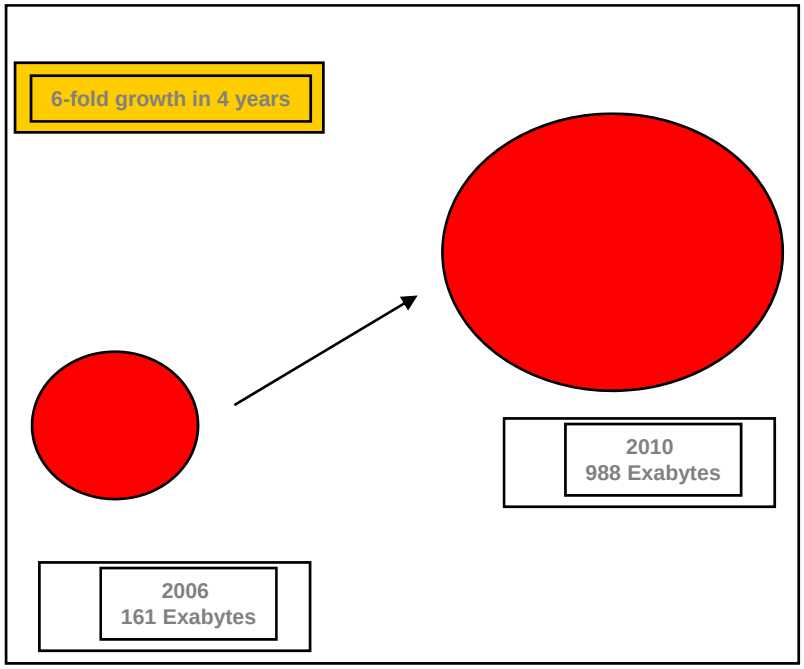

• IDC/EMC Study – The Expanding Digital Universe (March 2007)

Key growth drivers

- Analog to digital conversion
- Image capture devices
- Replicated content in e-mail, backups and web sites.
- Data retention/archives to comply with regulatory requirements.

Enormous pressure on current data centers, driving IT organizations to develop more information centric architectures

Storage Tiering & Information Life Cycle Management

- Storage Tiering:
  - Based on price, performance and availability characteristics of the media.
  - Data value matched to the tier through out the life cycle.
- Life Cycle Management:
  - Storage policies are defined according to business requirements
    - Placement, access, retention and deletion.
  - Automation is key.

Managing Storage – How Storage Tiering can help

Storage Management Challenges

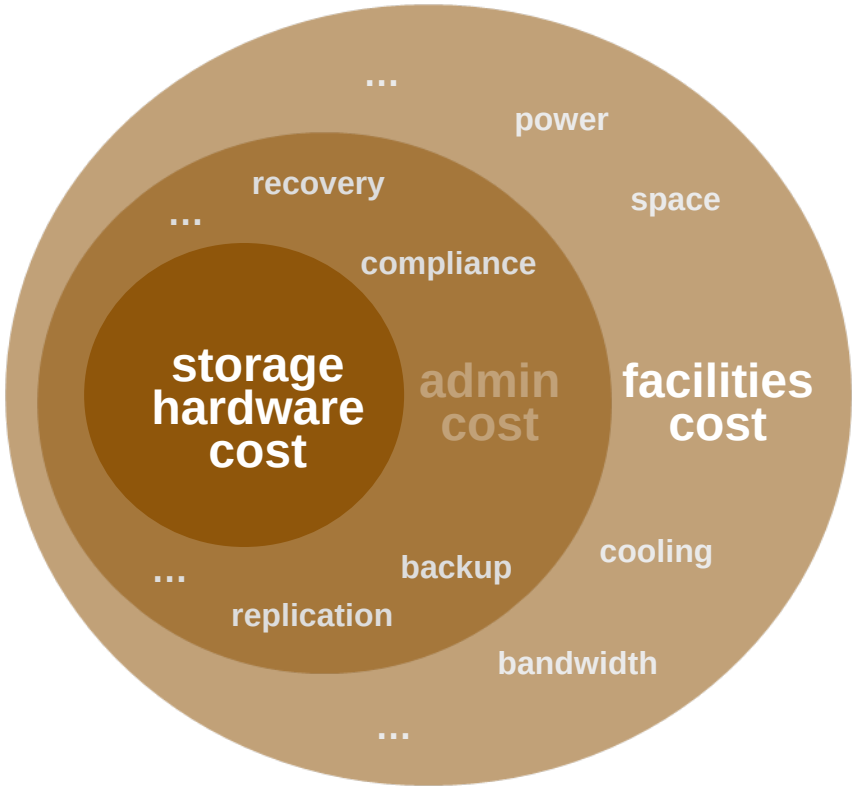

Benefits of storage tiering

- Match storage properties with the data's business value
  - Performance
  - Exploit properties of media e.g. tape, optical, MAID, SSD.
- Lower storage costs
  - Better utilization of costlier storage tiers through reclamation.
  - Minimize backups.

Who can benefit most from tiering?

- Telco
- Financial
- Hosted Service Providers
- Electronic Design and Automation
- Oil and Gas
- Healthcare
- Entertainment

Procedure for Storage Tiering

- Define policies that meet the business requirements and cost objectives.
- Classify storage based on cost, performance & retention characteristics.
- Select the technology/product that can implement and automate the policies and migrate data between tiers according to the policies.

Approaches to Storage Tiering

- Static (Adhoc) data migration
- Hierarchical Storage Management (HSM)
- Multi-Device file systems (Dynamic Storage Tiering)
- Network based storage tiering
- File Area Network switches based storage tiering

Static (Ad-hoc) Tiering

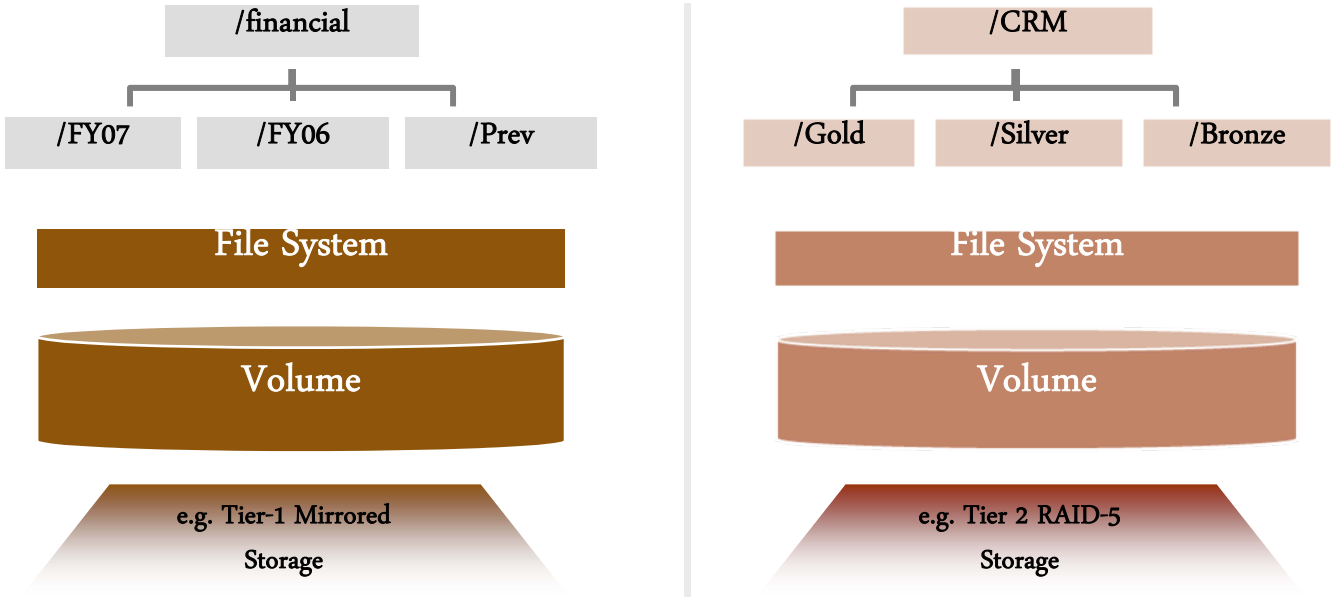

- Simple but inflexible
  - Ignores data value transformation within each application context
  - High administrative costs with procedural complexity
  - Too coarse grained

hsm

Hierarchical Storage Management (HSM)

- Software is closely integrated with the file system.
- Migrates files managed by the file system:
  - Leaves a stub in file system meta data to indicate location of migrated file (file still appears in the file system name space).
  - Restores files from secondary storage upon access by applications.
- Policy for migration is typically based on most recent access:
  - Move files that have not been accessed in *x* days to location *y*.

Hierarchical Storage Management (HSM)

- Inflexible Migration Policies
  - Uses only activity patterns to migrate files to 2nd tier
- Access Time Delays
  - Applications must wait for files to be reverse-migrated to primary storage
- Space Management Challenges
  - Cannot predict retrieval patterns or space needed for large files

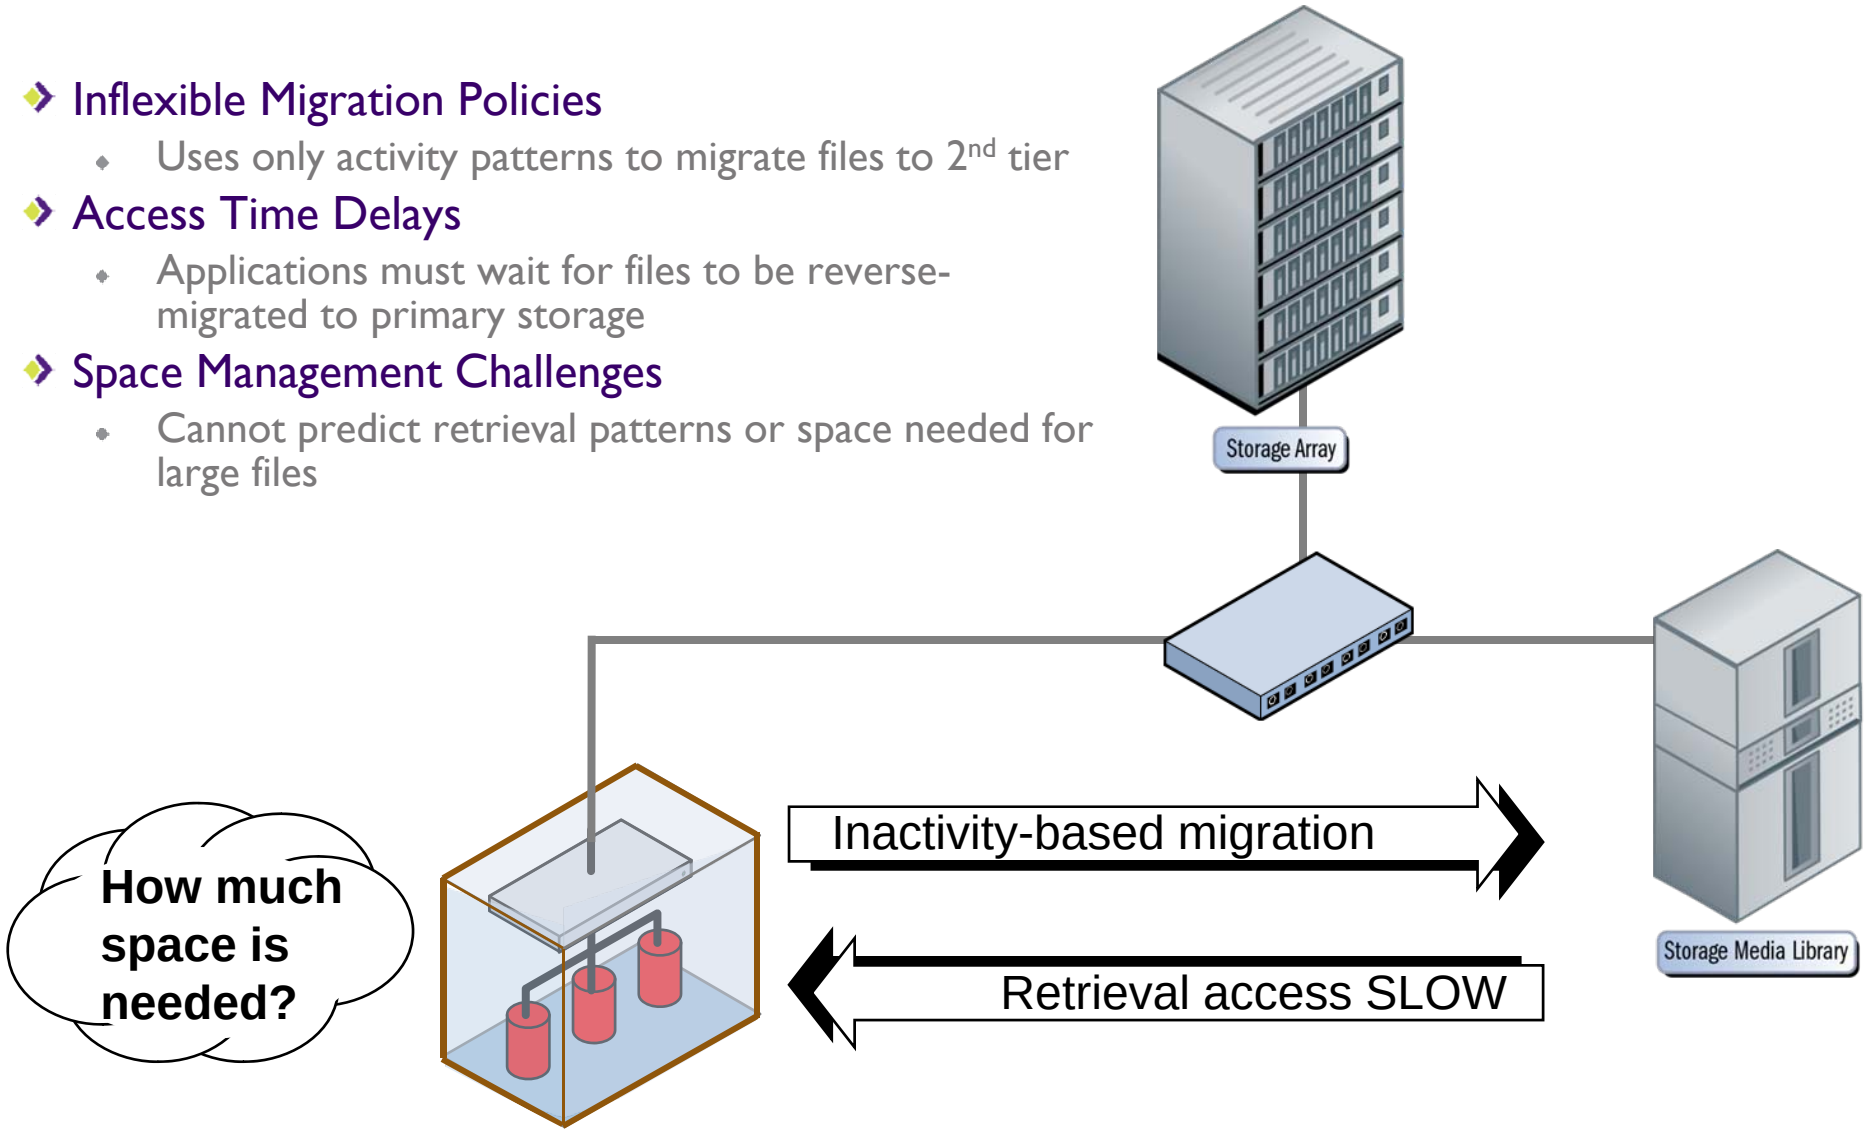

soln

Multi-Device File Systems (Dynamic Storage Tiering)

- Storage tier aware file system and volume manager
- Transparent file migration between tiers
  - No stubs for migrated files.
  - No retrieval penalty.
- Customizable behavior
  - Fine grained control
  - Policies: Per file system or per application
- Simplified Space Management
  - Behavior is customizable – can mimic traditional file system (overflow to other tiers until FS is full) or fail writes to full tiers

Multi-Device File Systems

Education SNIA

Key enabling technologies:

- **Multi-Device File System:**
  - One FS namespace over multiple tiers, consisting of multiple volumes, carved out of heterogeneous storage hardware
- **Policy Framework:**
  - Define initial placement; when, why, & where to relocate files
- **Separate tiers for meta data and data**
- ***Migration without stubs:***
  - Move files transparently to users, applications, tools, scripts

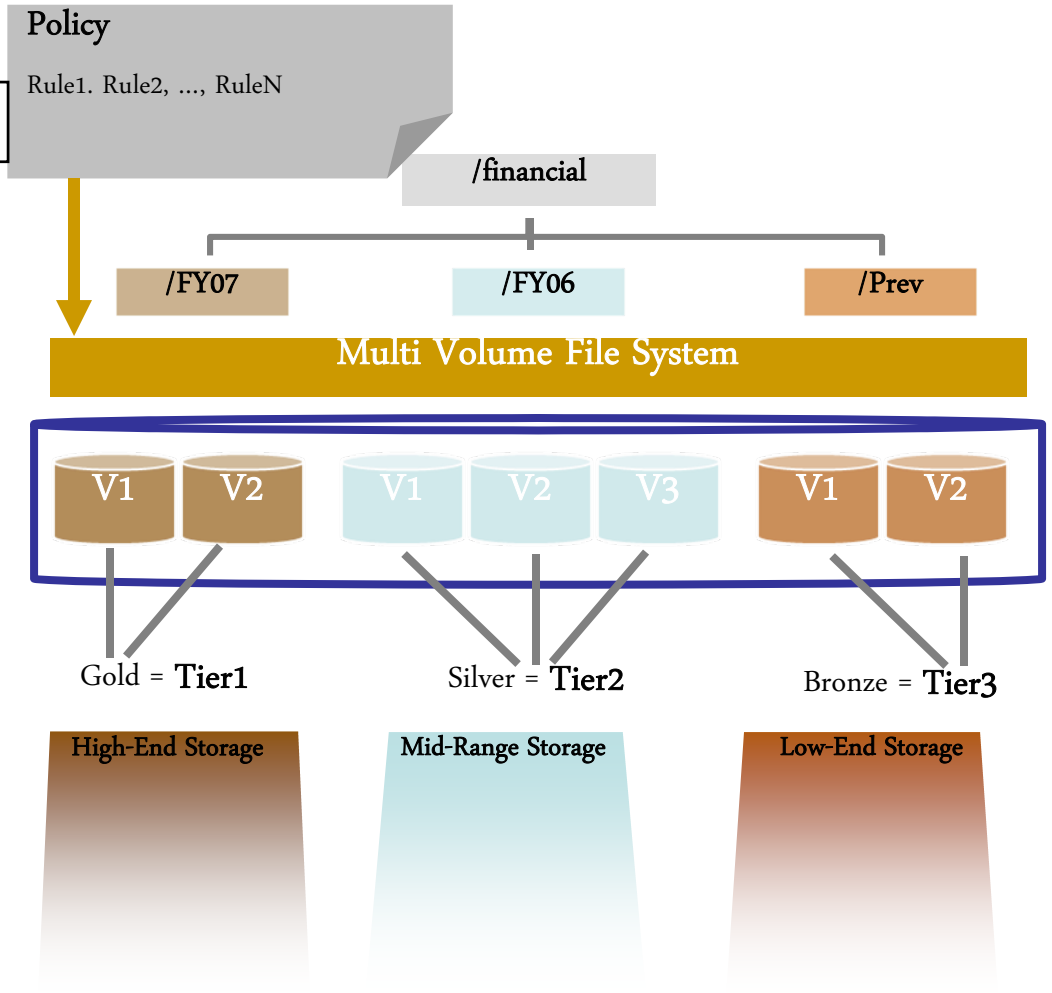

kernel

Multi-Device File Systems – Placement Classes & Policies

- Tier:
  - A volume or group of volumes with similar attributes
- Tag:
  - A mnemonic for a tier. E.g. “Gold”, “Silver” or “Bronze”
- Placement class:
  - One for each tier and referred to by its tag
  - Better abstraction than volumes
- Placement policies:
  - Set of rules that govern initial placement and subsequent relocation of a set of files. Rules specified in terms of placement classes.
  - One active placement policy per file system

Multi-Device File Systems – Placement Classes & Policies

- Placement rules:
  -  Directory (logical Location)
  -  File name (Name)
  -  User or Group (Owner)
  -  Modification time or Access Time (Aging)
  -  Access or I/O temperature (Activity)
  -  Space allocated to the file (Size)
  - Current Tier (Physical Location)

Multi-Device File Systems- Scale with Centralized Policy Management

- Define tiers with standard naming conventions
- Create a 'master' policy
- Assign it to multiple file systems across multiple hosts

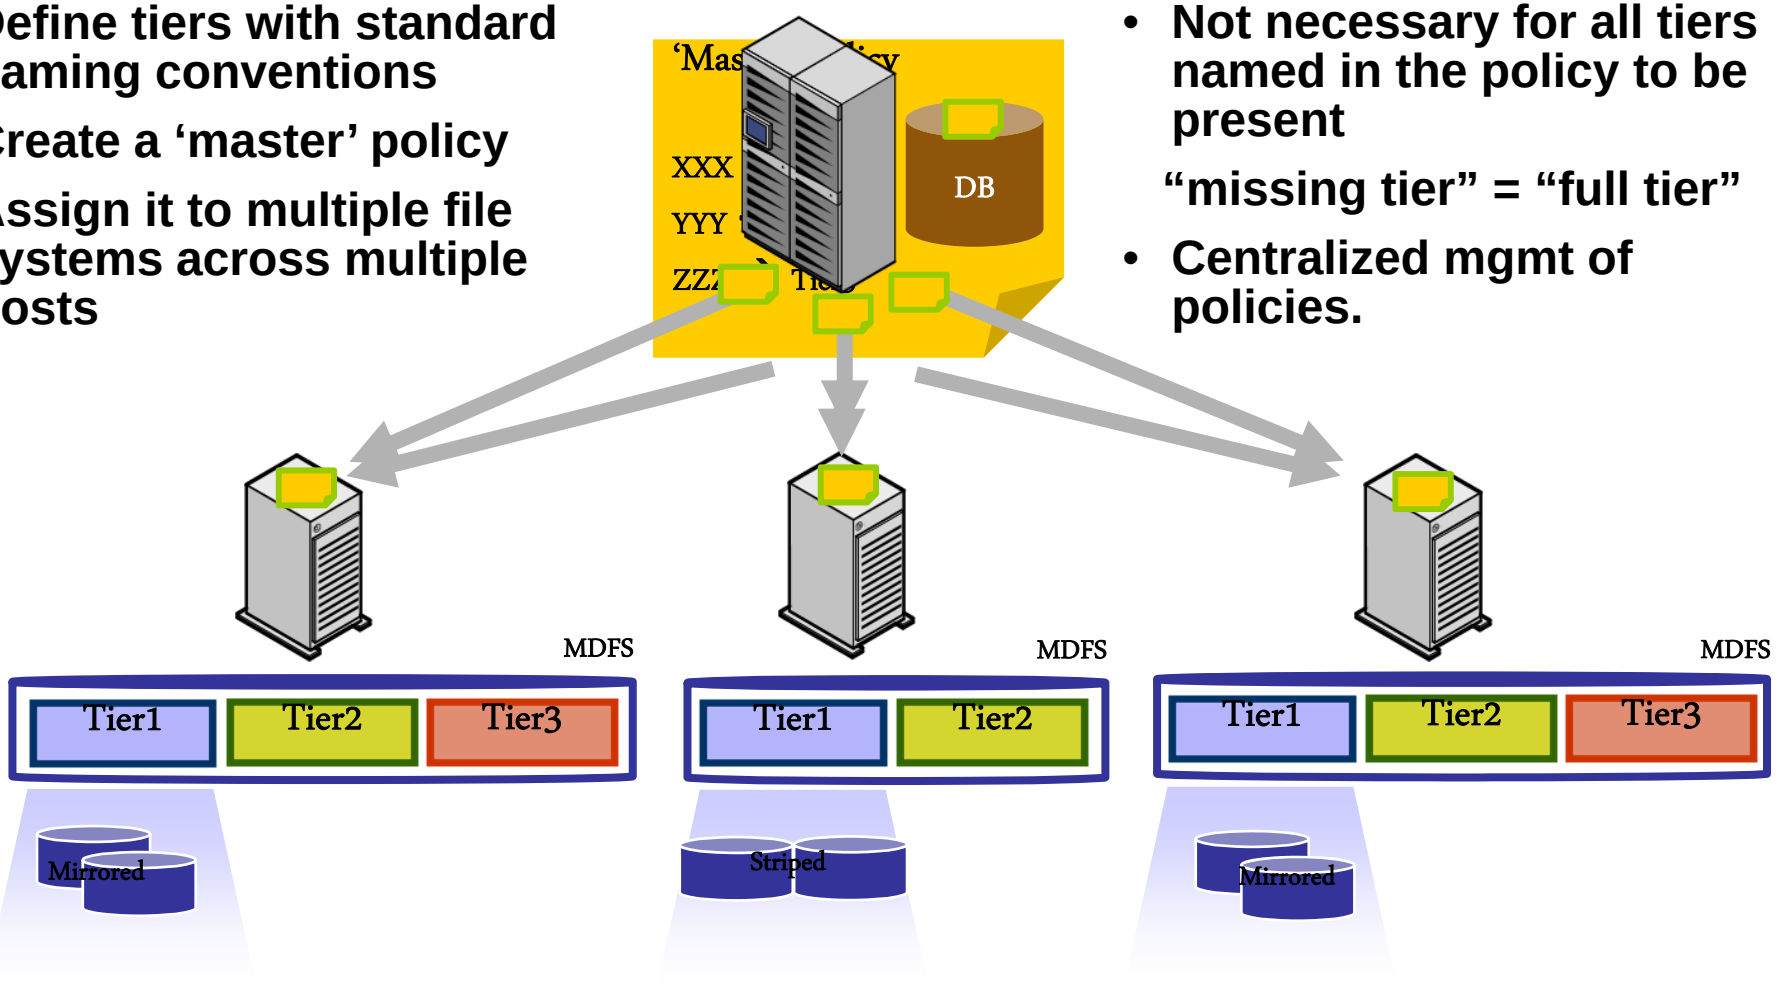

- Not necessary for all tiers named in the policy to be present

  "missing tier" = "full tier"
- Centralized mgmt of policies.

returns

Multi-Device File Systems - Limitations

- Technology is new and few commercially available file system have this capability.
- Tiering is limited to data managed by a single file system, although placement policies can be centrally assigned to many file systems.
- All different classes of storage must be connected to the host configured with the multi-device file system.
- Supporting media other than disk (e.g. tape or optical) will require significant changes to the file system.

Network based approaches – Centralized policy managers

- Centralized server (or a cluster of servers) takes on the role of policy manager and data mover for Network File Servers (NAS)
- Policy engine scans all shared file systems and migrates files that meet the criteria specified in the policies.
- Migrated files are replaced with stubs that point to the actual location of the files.
- Policies are typically access based with additional actions related to compliance.
- In some cases the policy manager registers with the file server to receive notifications for all access to files under the control of the policy manager.

Network based approaches – Centralized policy managers

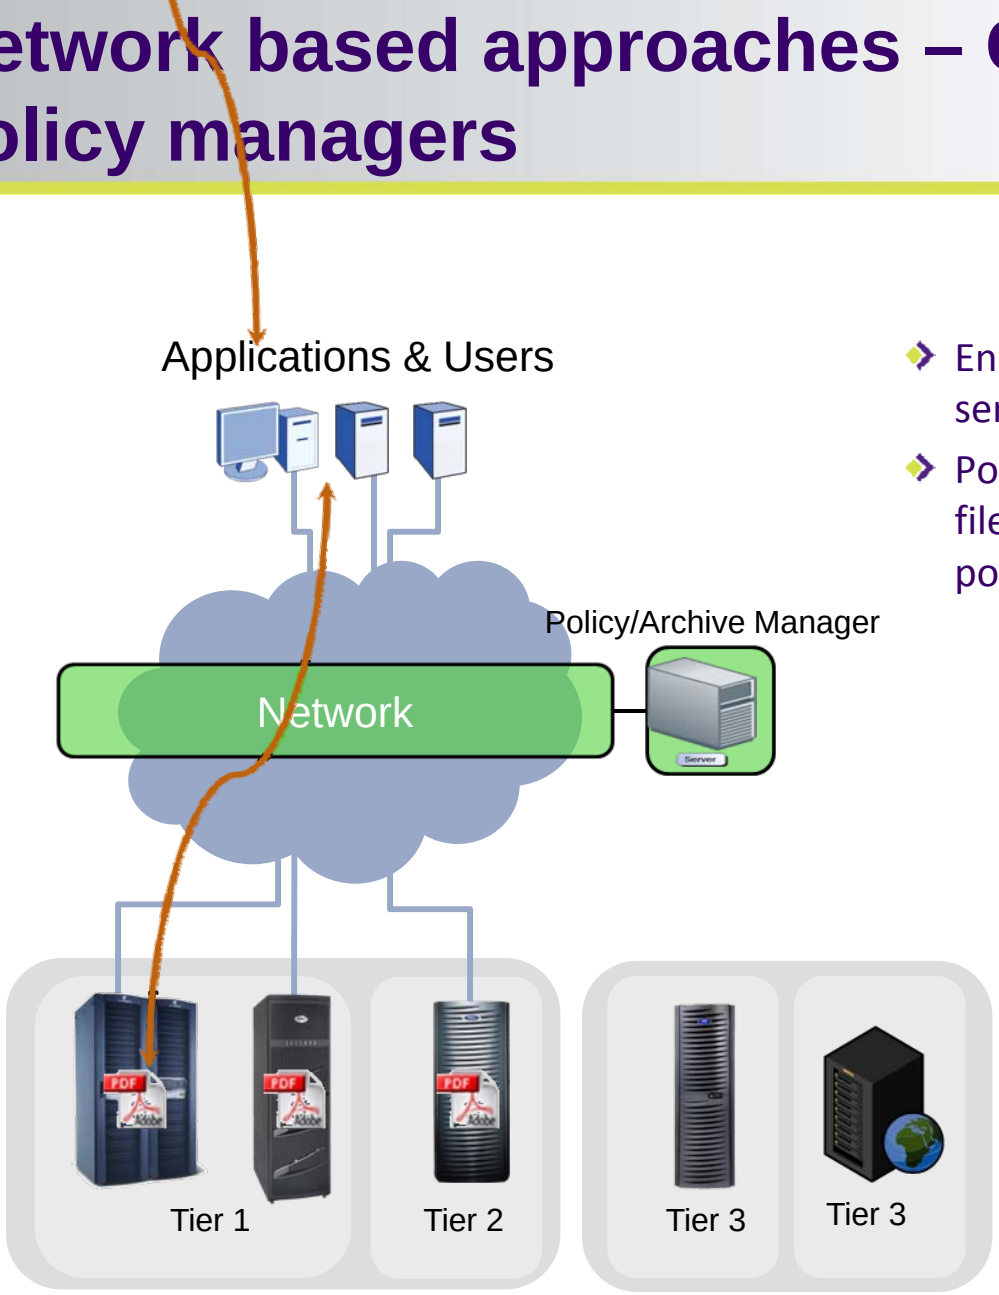

- End Users access files on Tier 1 and Tier 2 file servers through shared file systems.
- Policy manager crawls the network, scanning file systems for files that meet specified policies.

Network based approaches – Centralized policy managers - Archiving

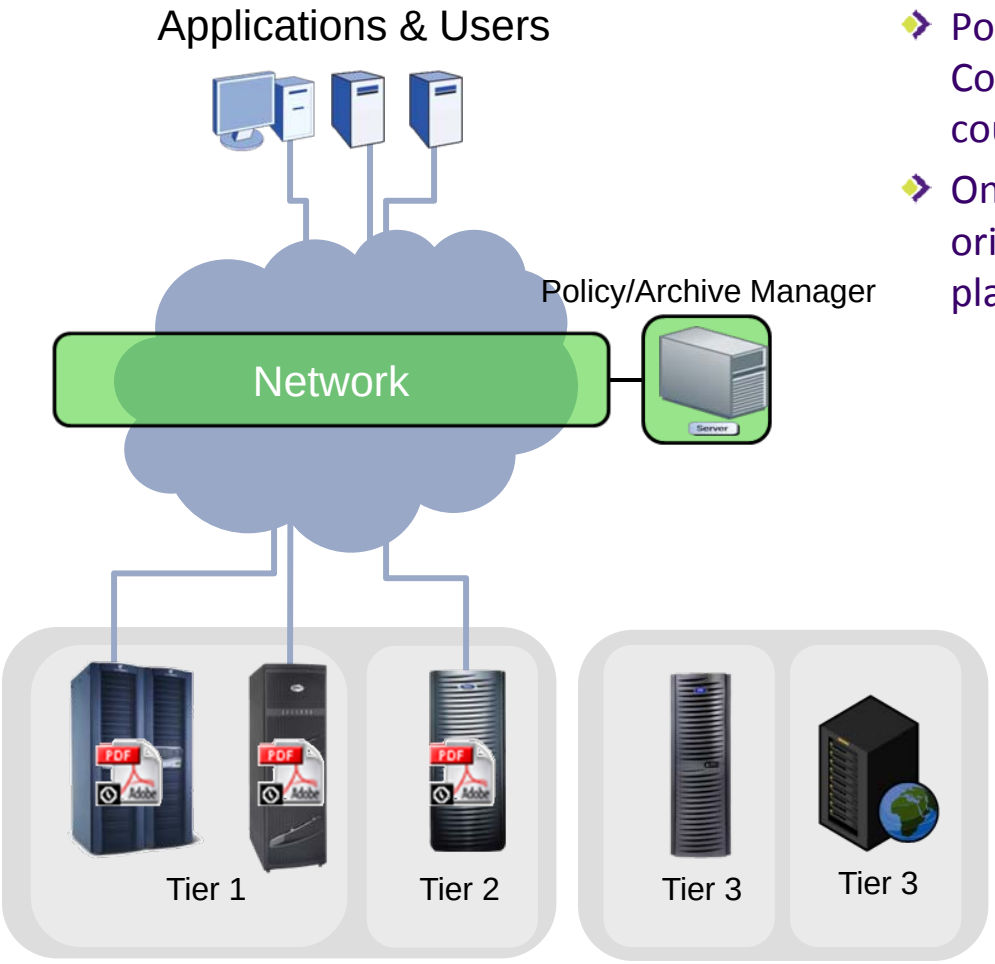

- Policy Manager will archive a copy of the file. Compression, Single-Instancing and Indexing could occur.
- Once the file is successfully archived, the original files are replaced with space saving placeholders.

Network based approaches – Centralized policy managers - Recall

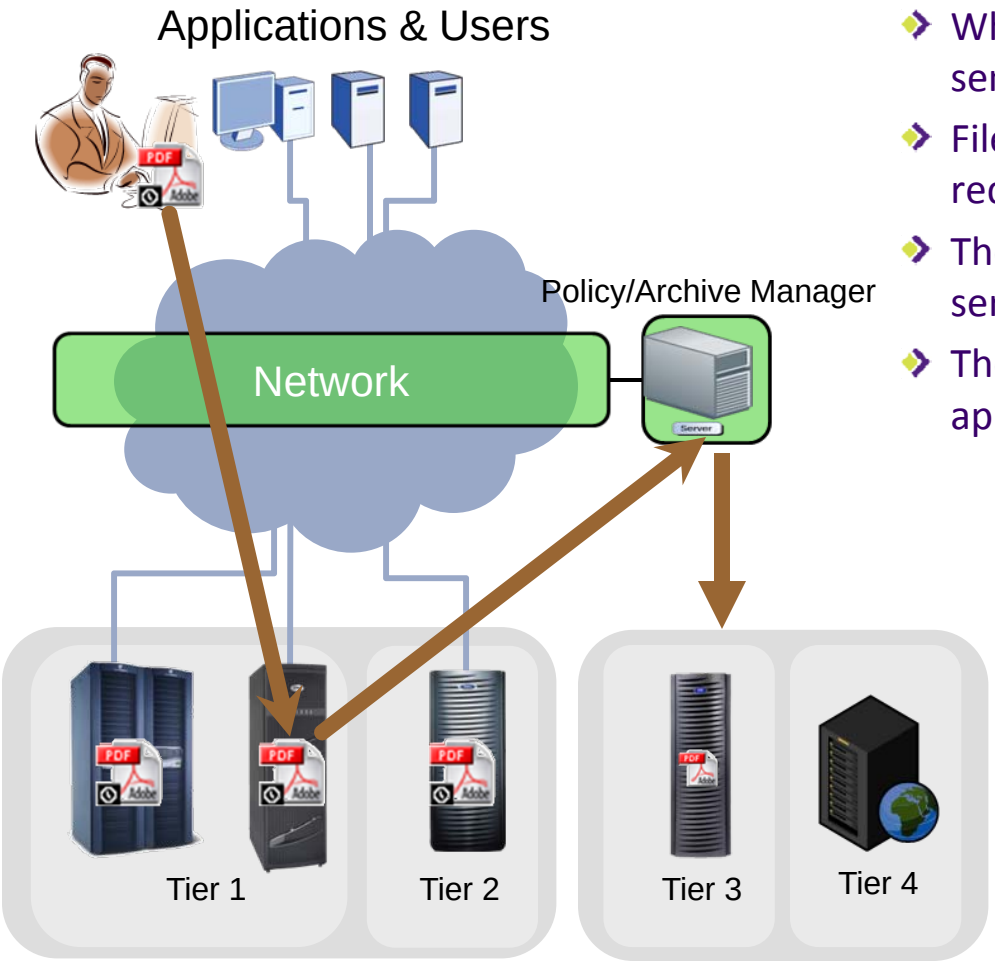

- When a placeholder is accessed, a request is sent to the file server.
- File server notifies policy manager about request to access file.
- The file is copied from the archive to the file server.
- The file data is returned to the calling application/user by the file server.

Network based approaches – Centralized Policy Manager - Limitations

- Similar limitations as HSM:
  - Simple access based policies.
  - File stubs to denote migrated files.
  - Access time issues for recalled files.
  - Space management issues.
- Crawling the network to scan file systems can consume significant network bandwidth.
- Policy manager requires access to all network file servers.
- Policy manager must inter-operate with all types of servers.

File Area Network based approaches

- File Area Network:
  - A collection of storage and network devices and software technologies that organize, route, switch, place and provide access to all the files in an enterprise.
  - Facilitates the implementation of data management services (e.g. storage tiering/ILM, search, replication etc.) in the network.
- File Area Network switches:
  - Virtualize file servers accessed with NFS/CIFS protocols.
  - Provide a large virtual name space constructed from the physical name spaces of file servers.
  - Automate policies for storage tiering/ILM.

File Area Network based approaches – Data Migration

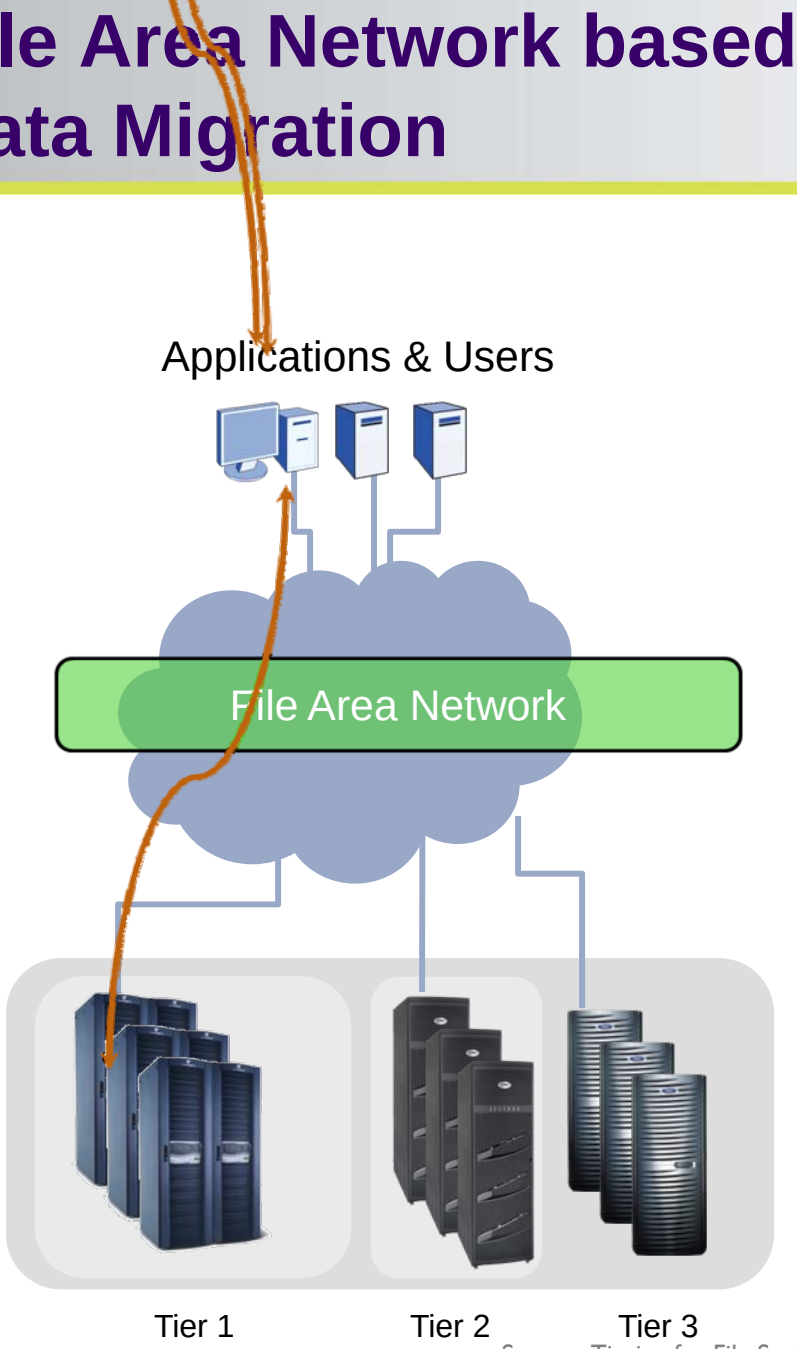

- End users/Applications access files on different tier file servers through virtualized name space.
- Policy manager resides in a “filer virtualization” switch.
- Policies specify file placement for creation and relocation.
- Policies tend to be access time based.
- Policies could allow duplicate copies) of files to be created.
- Data migration between tiers is transparent to applications.
- No stubs. No recalls. No space management issues.

File Area Network based approaches – Limitations

- Relatively new technology; primarily deployed to consolidate Network File Servers
- Requires investment and training in new network infrastructure hardware, software and administration.
- Tiering limited to file servers accessed using NFS/CIFS protocols.
- Cannot inter-operate with purpose built archive servers accessed using protocols other than NFS/CIFS.
- Simplistic policies.

Conclusions

Approach to Tiering	Capability				Where to deploy
	Transparent Migration	Richness & flexibility of policies	Stub migrated files	Space Management issues	
Adhoc	**No**	**Yes**	**No**	**No**	Custom policies that vary from time to time.
HSM	Yes	No	Yes	Yes	Simple access time policies are sufficient. Migrated data accessed infrequently or there is very high tolerance for access times.
Multi-Device File System	Yes	**Yes**	**No**	**No**	Wider range of policies than just ones based on access times. Use of separate tiers for meta data. Low latency requirements when accessing data from any tier.
Network based policy managers	Yes	No	Yes	Yes	Migrate files from Network File Servers. Very good for migrating files to purpose built archive servers. For meeting regulatory & compliance requirements.
File Area Networks	Yes	No	**No**	**No**	Migrate files between file servers (tiers). Simple access & creation time policies with load balancing are sufficient. Low latency requirements when accessing data from any tier

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: trackfilemgmt@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Paul Massiglia	**Data Storage Consultant**
Murthy Mamidi	**Symantec Corporation**
Oliver Robinson	**Symantec Corporation**