October 12-14, 2004

Distributed Hierarchical Storage Management (DHSM)

John Hayden – EMC Engineering

EMC²
where information lives

Agenda

- Customer Challenge
- Distributed HSM
- Influences & Design Goals
- DHSM leveraging the DHSM API

Customer Challenge

- File System growth is rapid – storage requirements compounded
  - The value of information changes over time
  - Storage resources vary in performance, cost, manageability and levels of data protection
  - Not all files should be sitting on primary storage

Customer Challenge

- Managing the movement of files to more cost effective storage such as ATA based storage
  - Manual process can be time consuming
  - Data location transparency should be kept
  - Manage cost of backup and make it feasible

File-Based HSM or "Distributed HSM"

- Gartner
  - Distributed HSM is file-based archiving – not database data, not e-mail data, and not mainframe HSM
- EMC
  - Distributed
  - File migration and management
  - Policy-based
  - Open

Guiding Design Principles of DHSM

- Maintain the NAS head as the customer-facing device
- Leverage existing features and core competencies
- Open architecture that allows us to partner to leverage expertise
  - Encourage 3rd party integration
- Provide a scalable and robust architecture

Existing NAS HSM Implementations

- Not completely transparent
  - Relies on shortcuts and symlinks
  - Clients must access both primary and secondary storage to access migrated data
- No automatic, user-driven mechanism for data to be recalled back to the primary storage in real time
- Single-protocol environments only

Existing DHSM – Architectural Overview

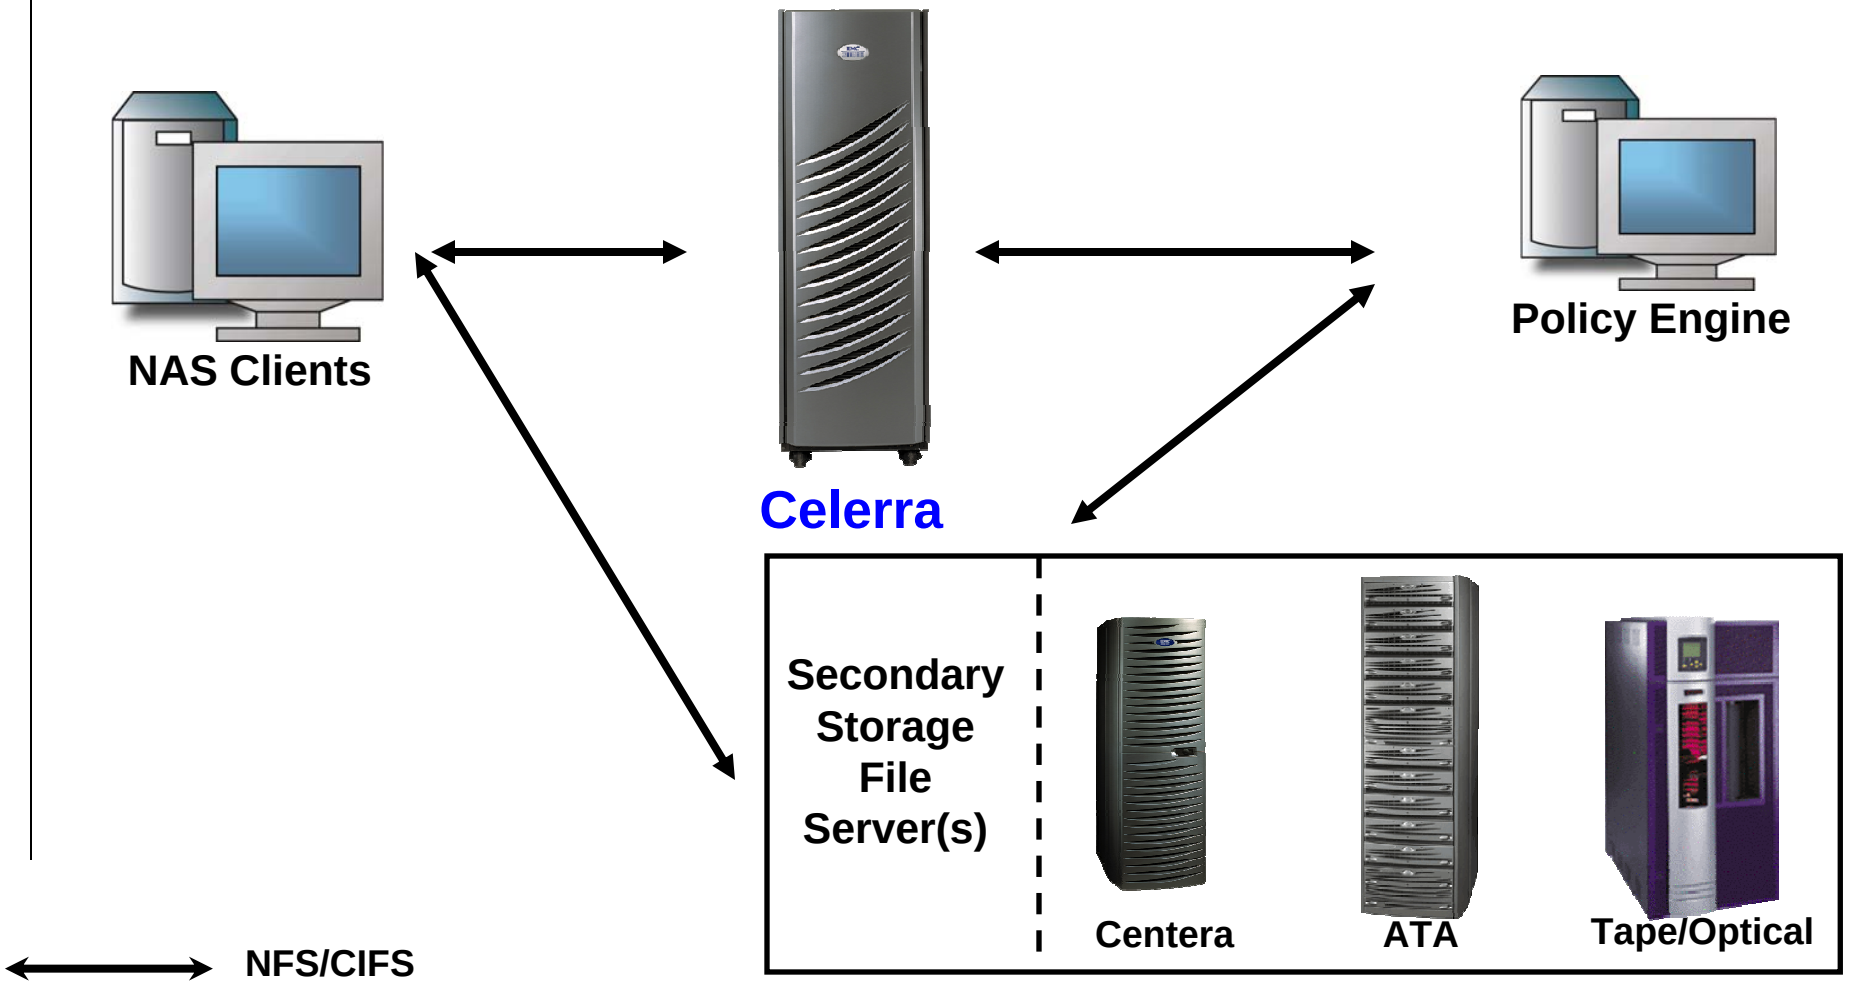

Existing DHSM – Write Path

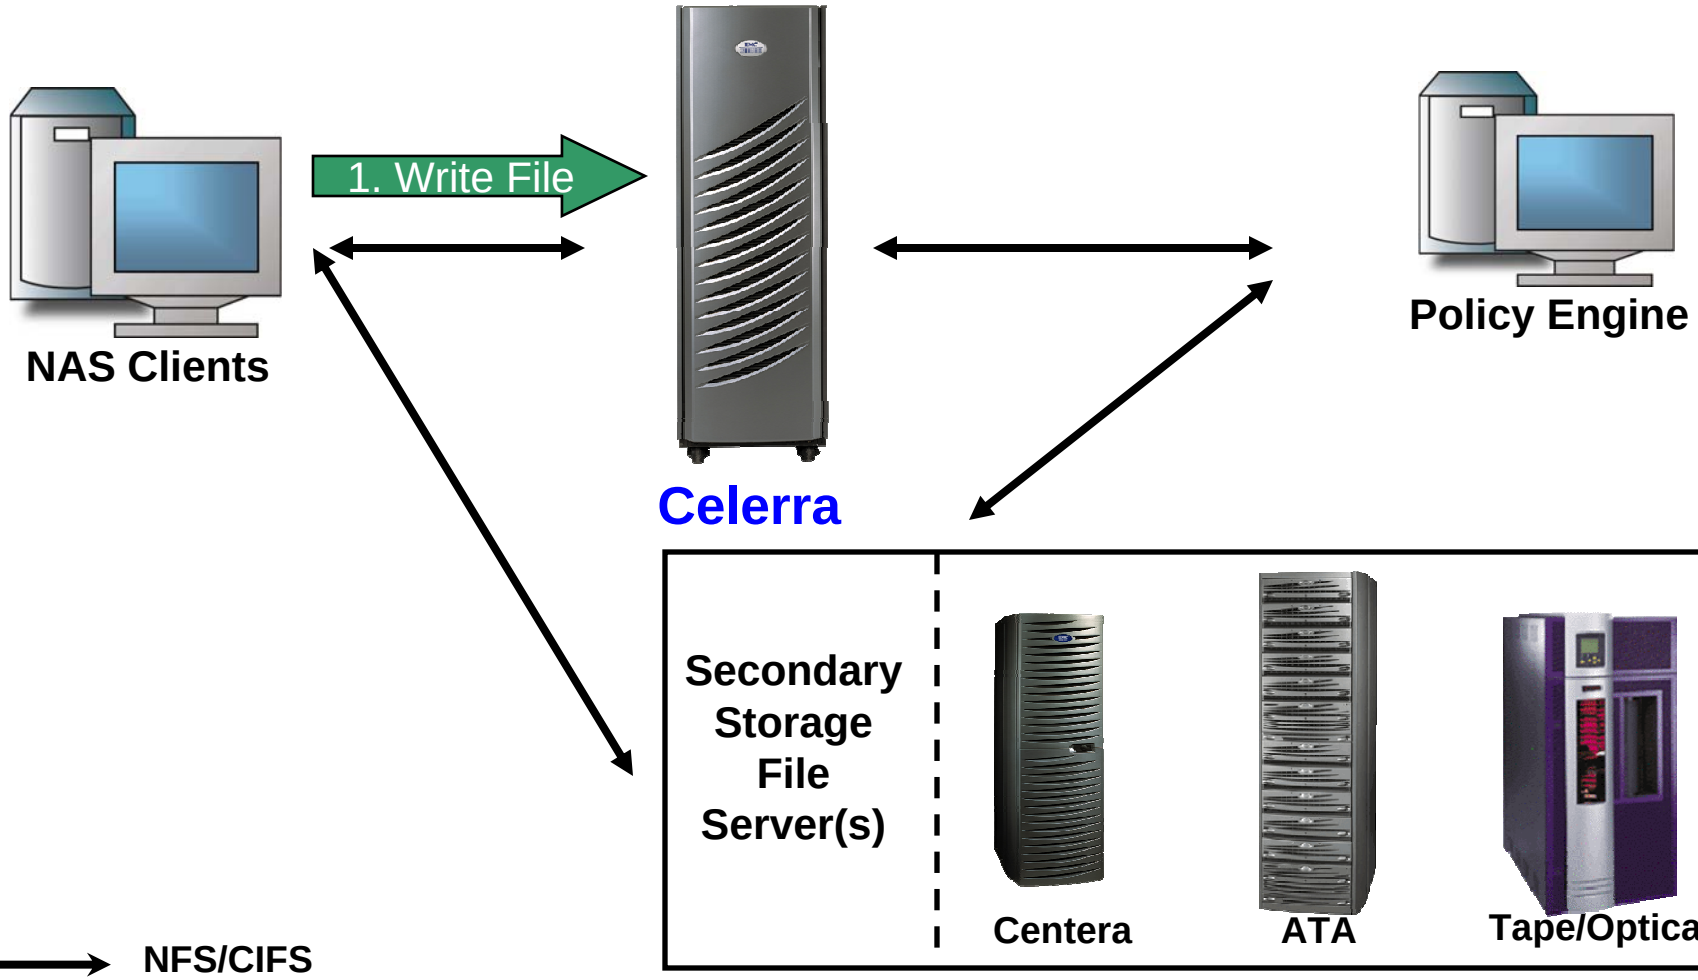

Existing DHSM – Migrate Path

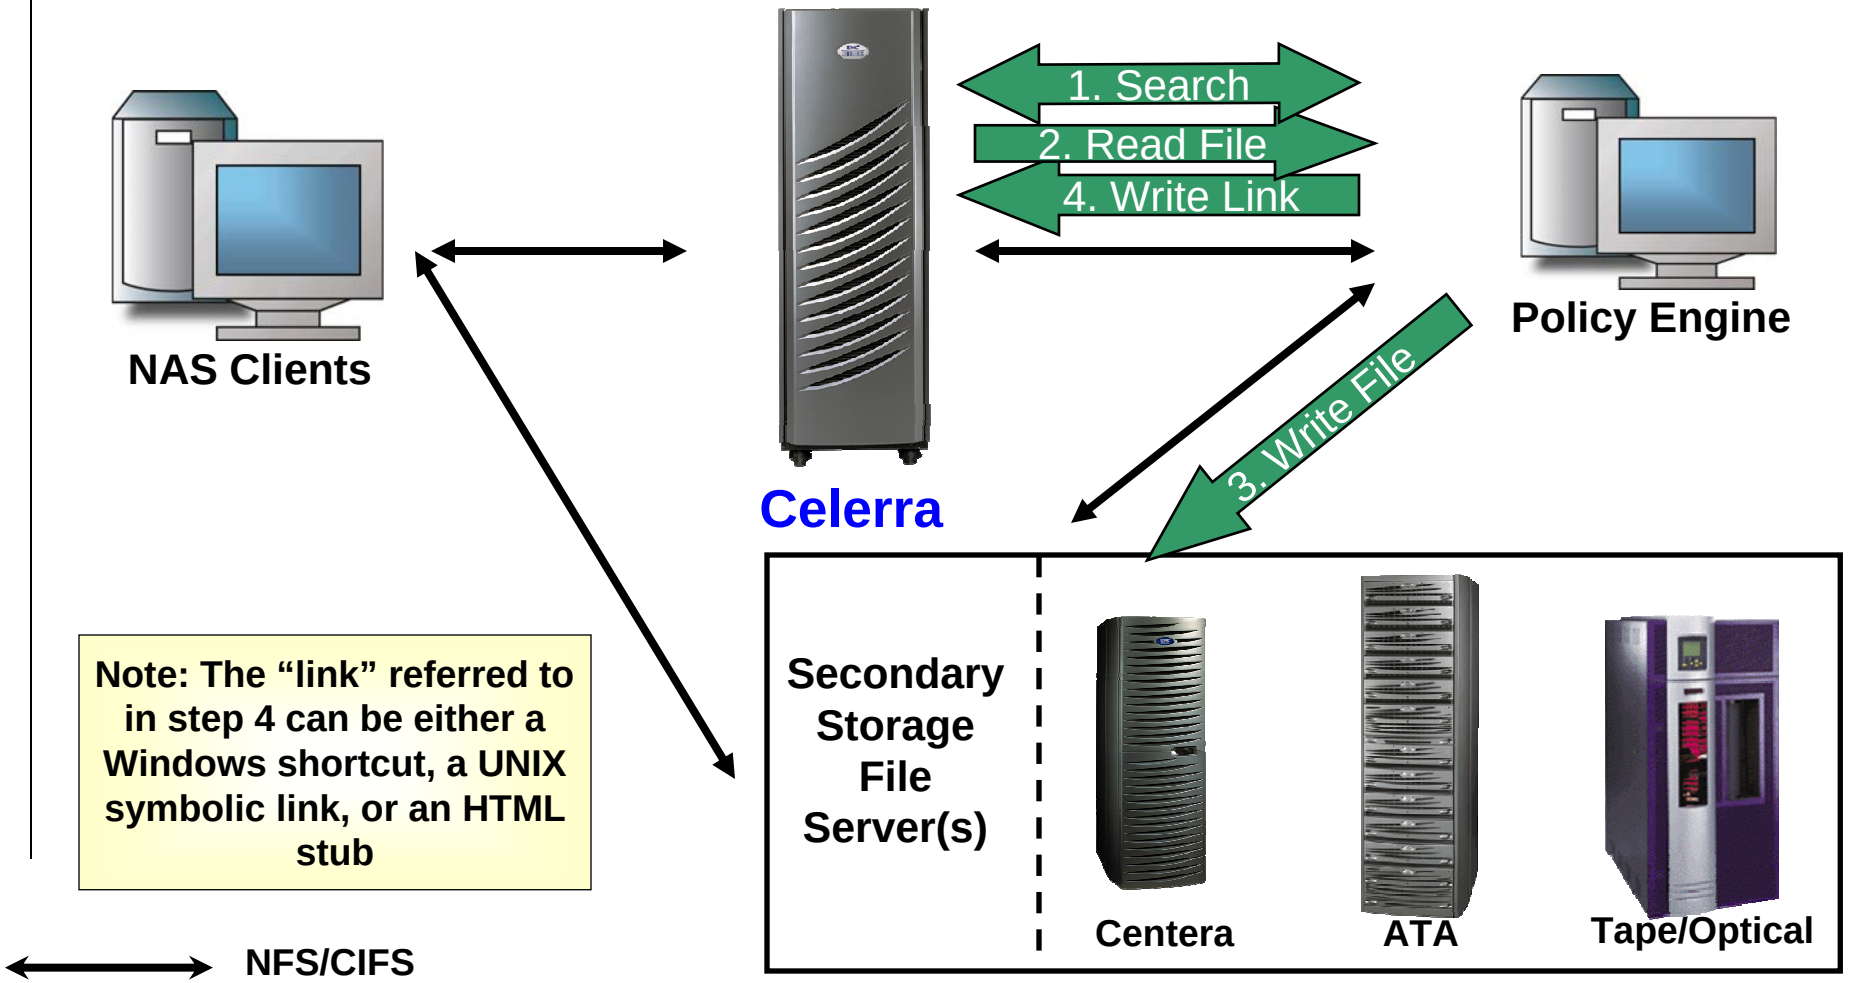

Existing DHSM – Read Path

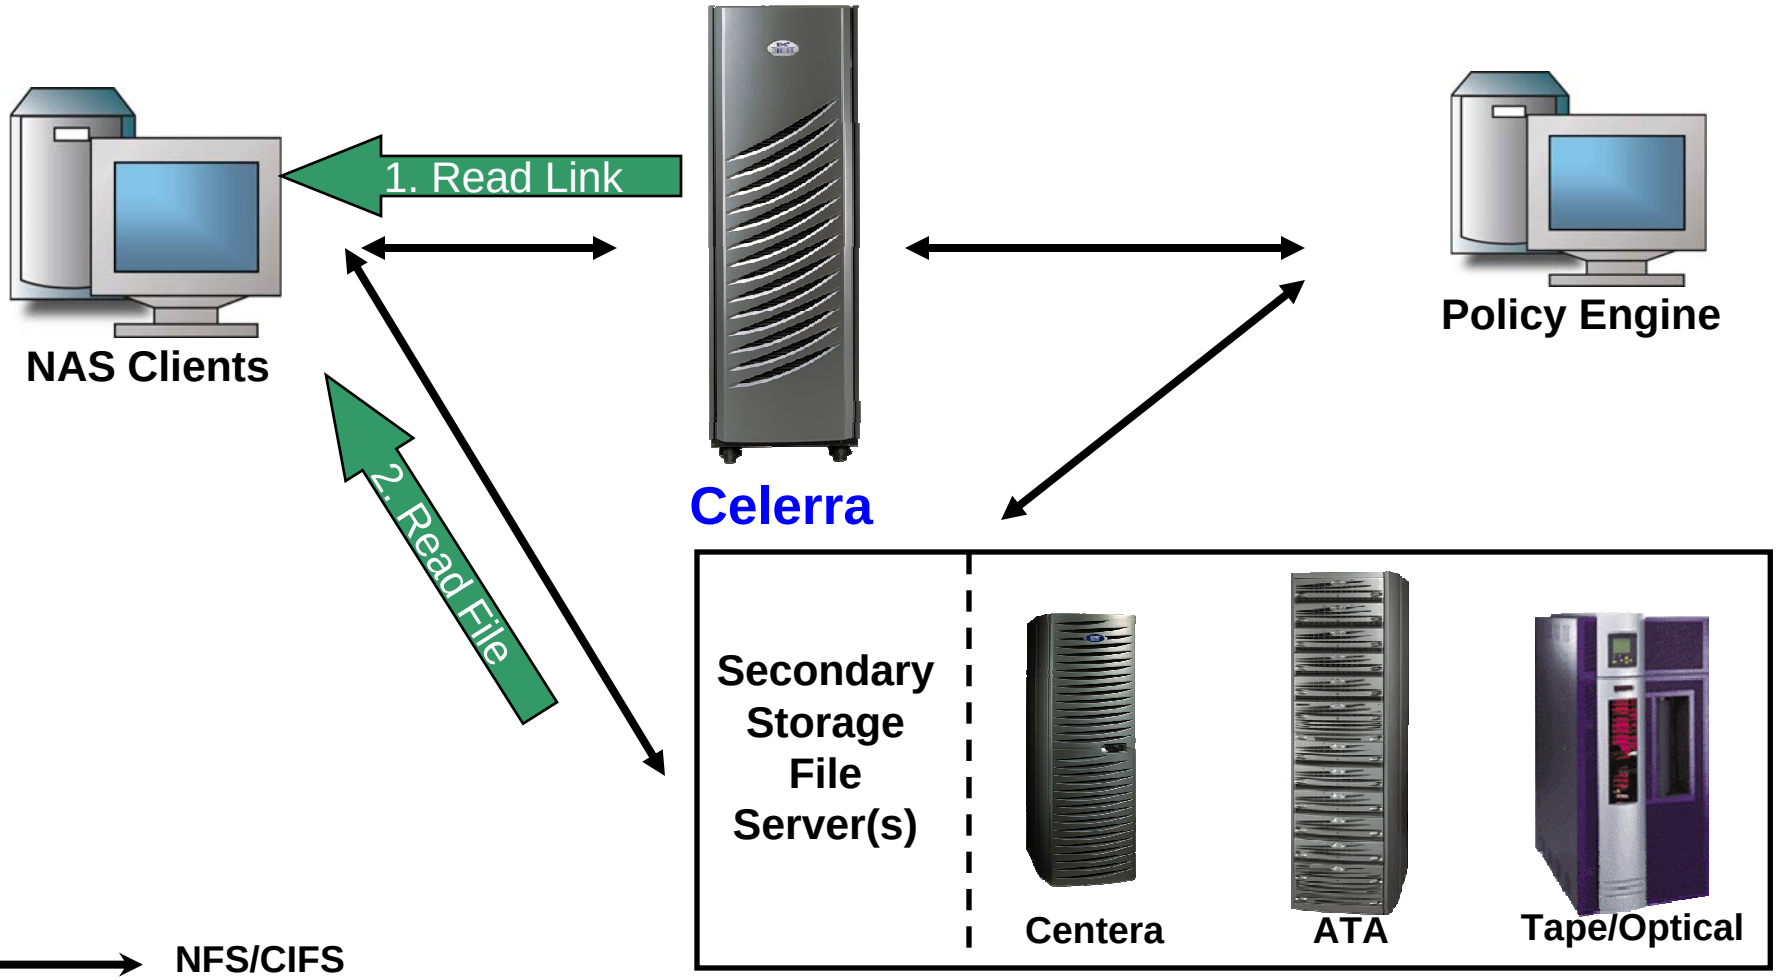

What can we leverage?

- Proven components in the Celerra Data Migration Service (CDMS) functionality
  - NFS and CIFS clients built in
  - Offline inodes
  - Connection database

What was missing?

- An offline inode API to the DART
- Policy and data migration engines that understand the Celerra offline inode API

Celerra DHSM

- Policy engine data migration
  - Periodically copies primary files to secondary store.
  - Overwrites primary files with CDMS style offline files.

Celerra DHSM

- Transparent data recall
  - May or may not migrate back based on configuration
  - Data access to secondary store is handled internally
  - Policy Engine not involved here
- Eases backup/storage costs
  - Decreased frequency of backups on secondary
  - Lower secondary storage costs

DHSM – Architectural Overview

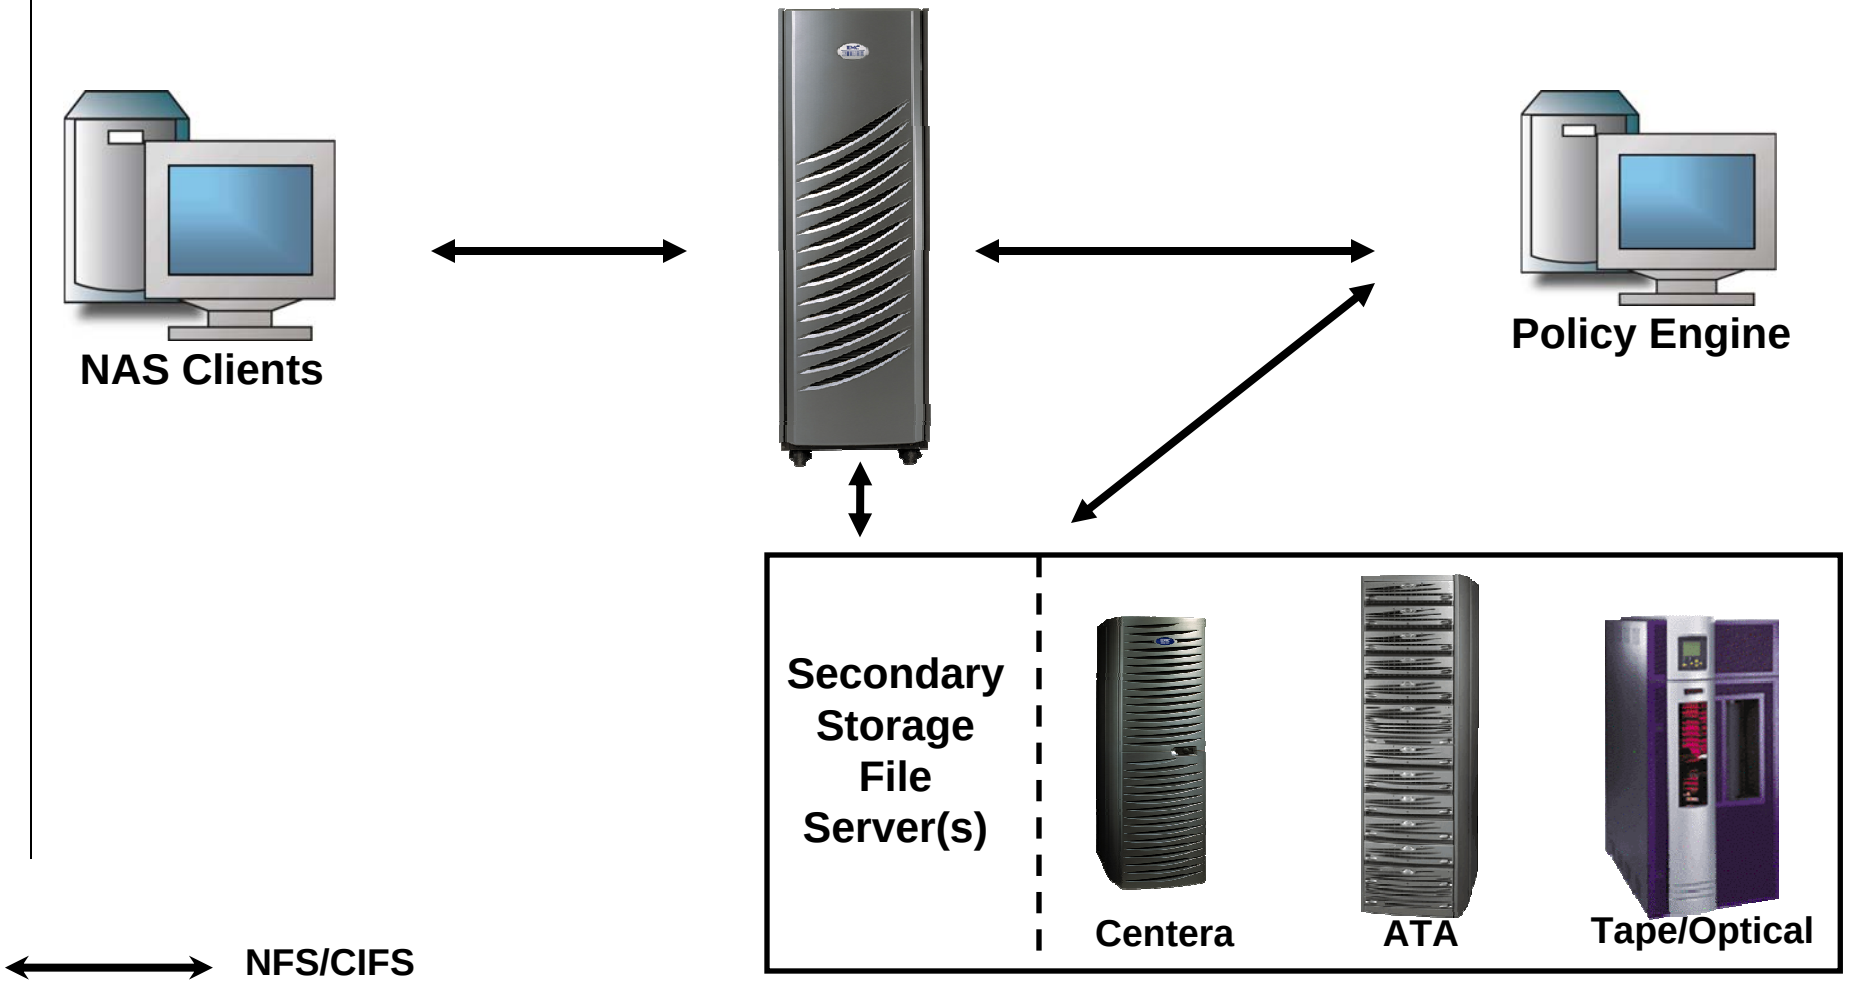

DHSM – Write Path

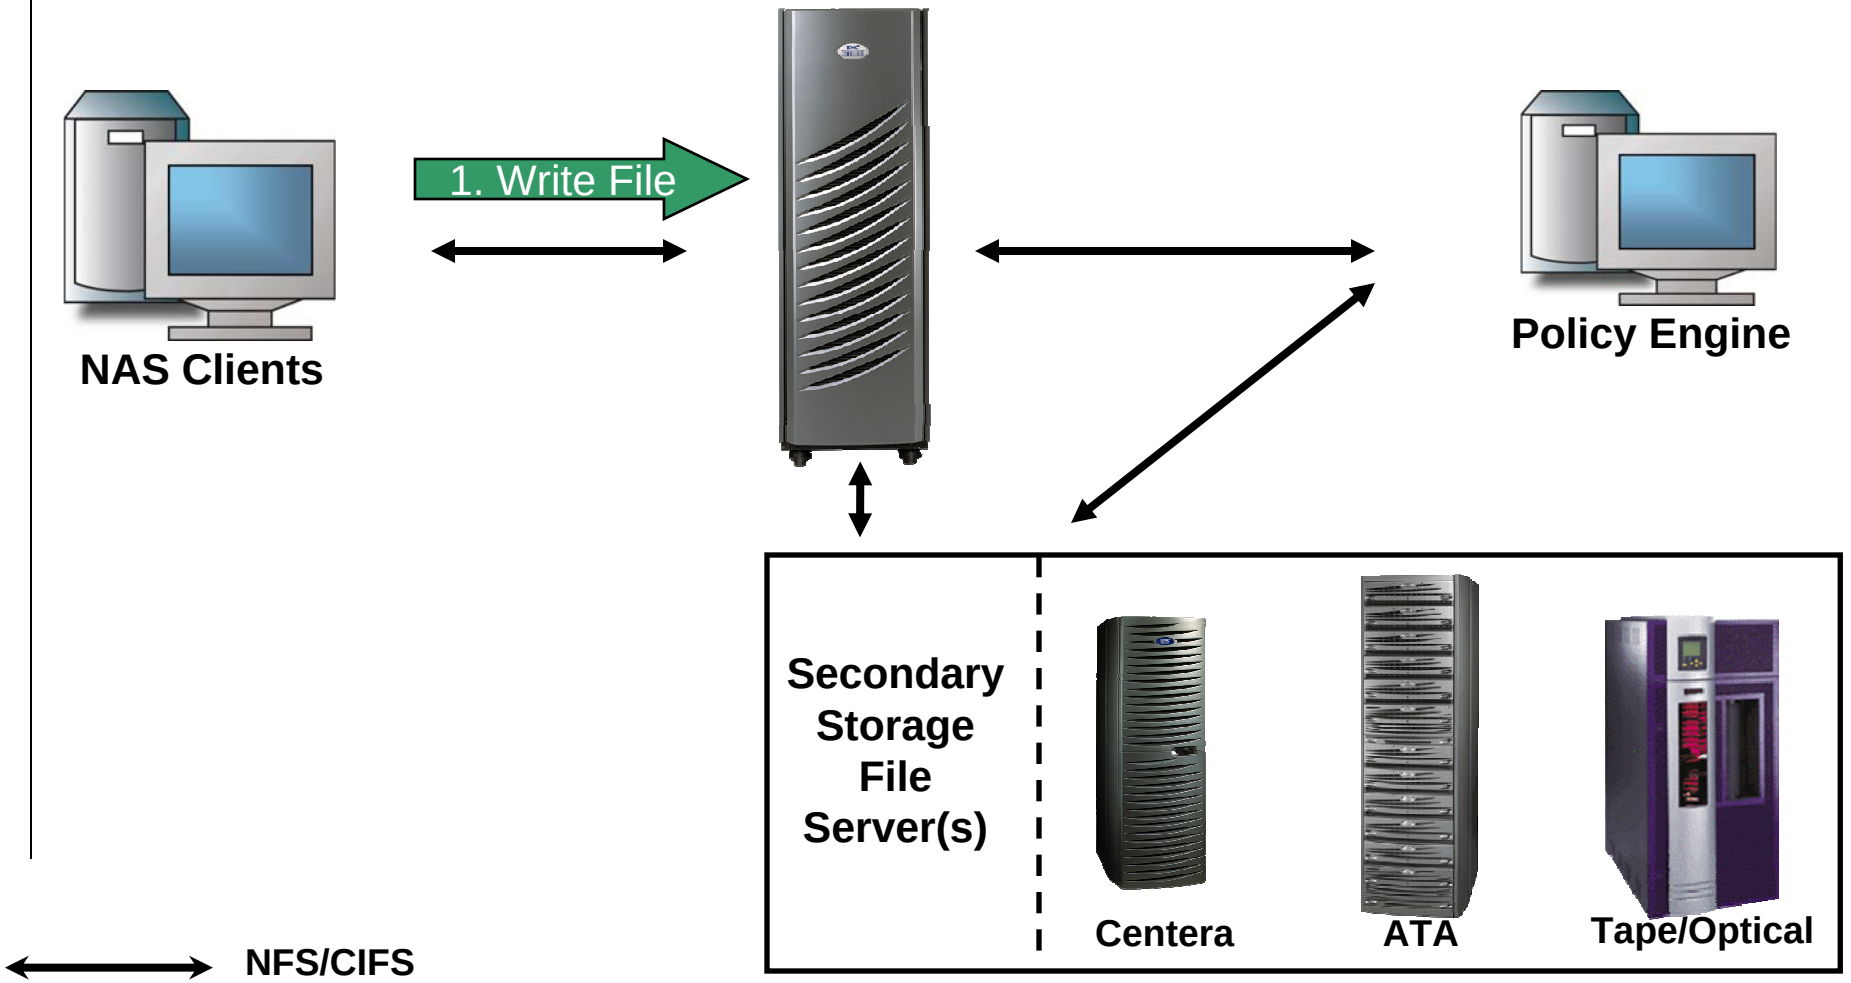

DHSM – Migrate Path

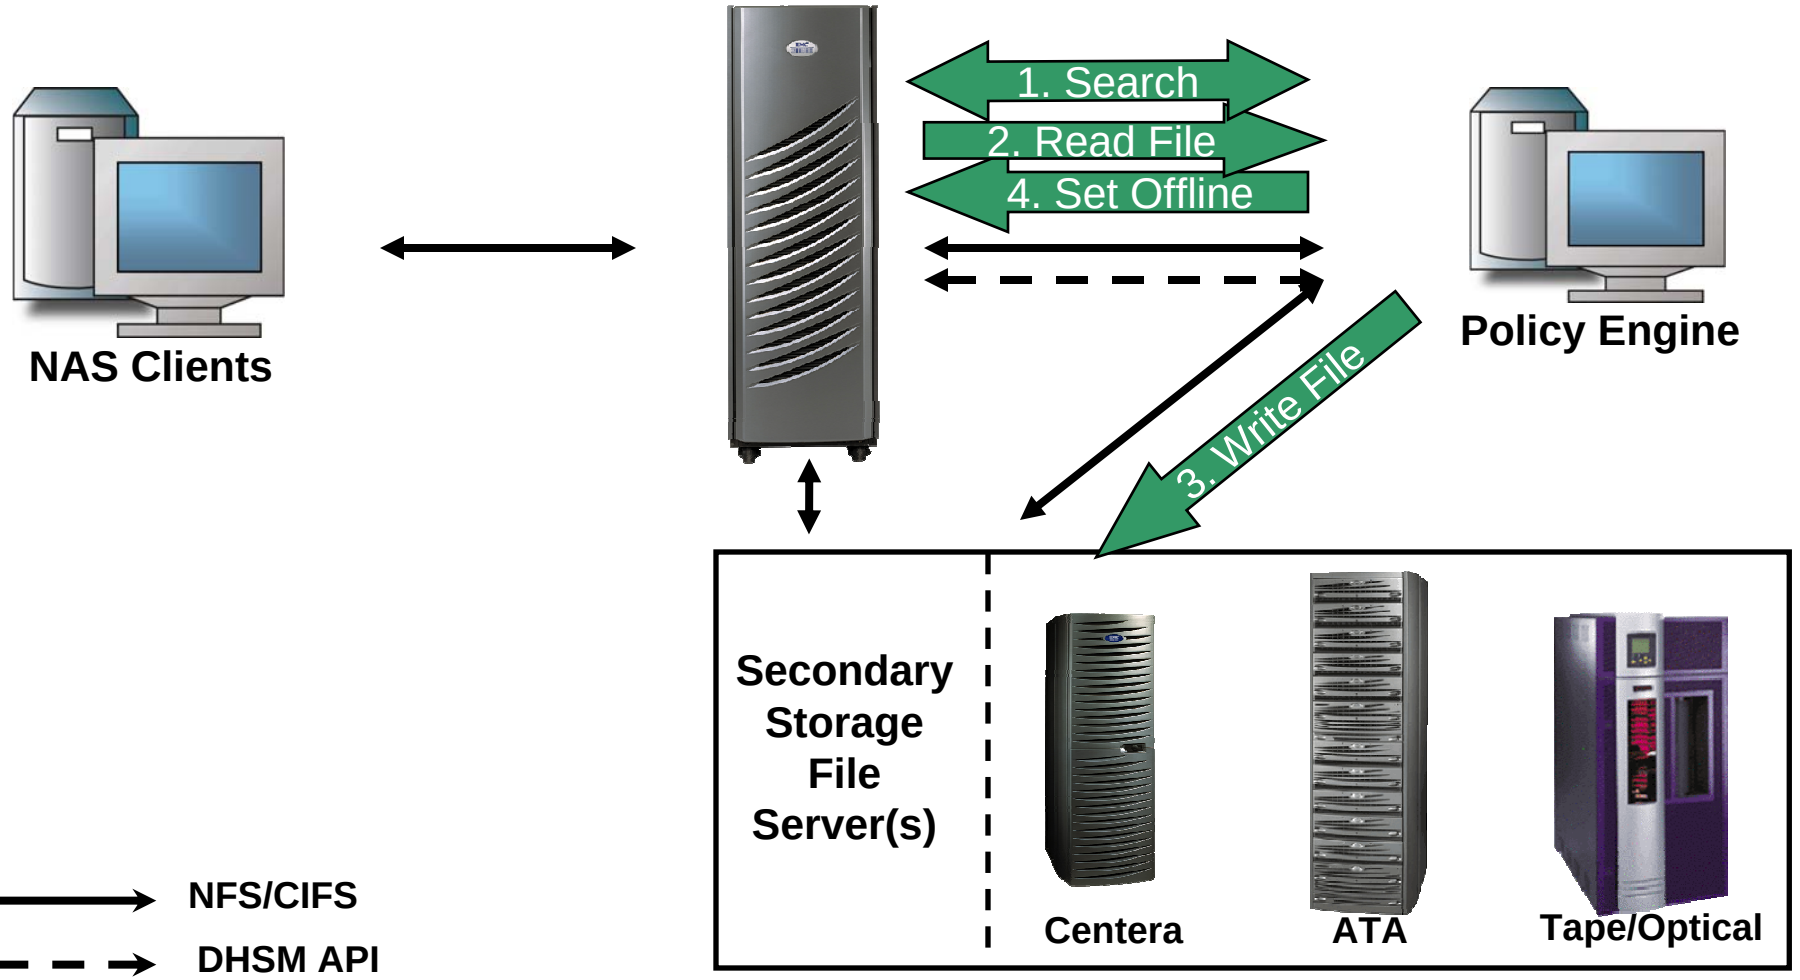

DHSM – Read Path

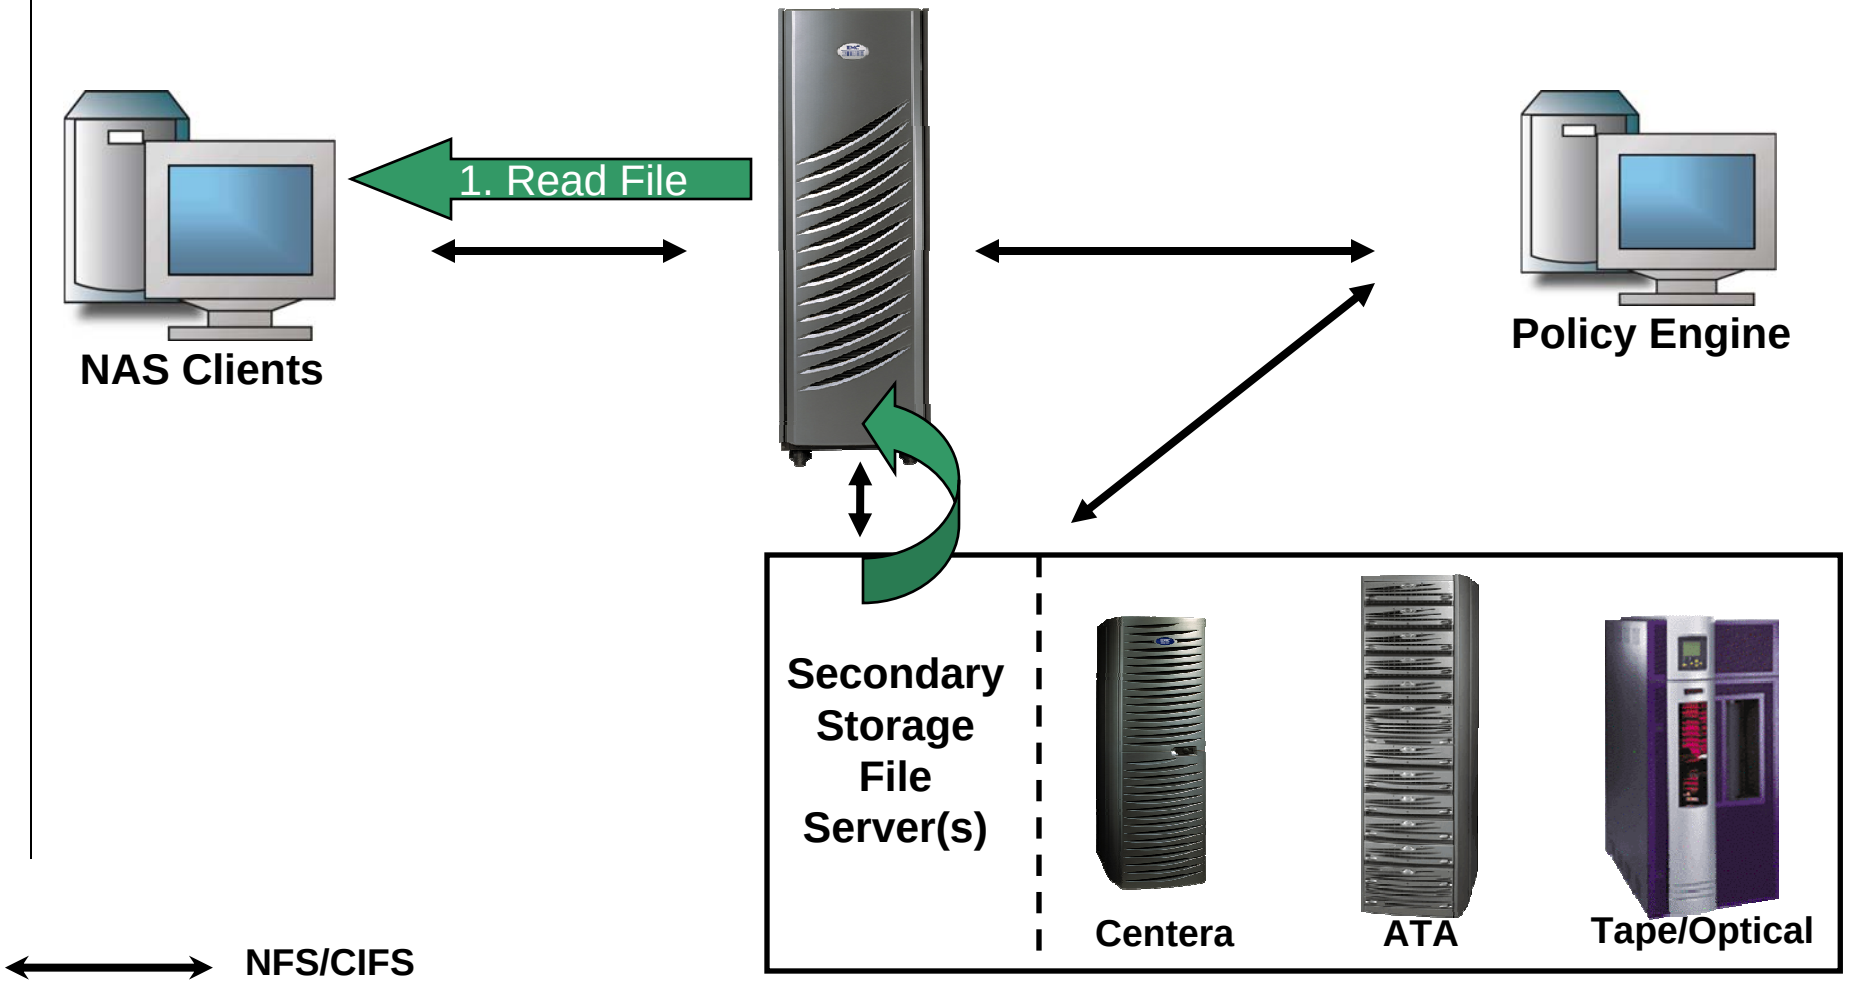

DHSM API

- XML over HTTP
- Two particular calls
  - DHSM_SET_OFFLINE_ATTRS – set a file offline or modify its attributes
  - DHSM_GET_ATTRS – query a file's attributes, including offline status.

Offline Files

- All attributes and metadata reside on primary store
- Offline Inode
  - Opaque Data
  - Absolute pathname
  - Migration Method
  - Verifier
- Validation
  - Validates that secondary file is in sync with primary offline file
  - Modification time/file length validation
  - Occurs prior to offline I/O.

Software Partner Status

- DHSM API Development Kit available since November 2003

- Actively seeking additional API partners

DHSM – Benefits

- Data value is aligned with storage
- Can use almost any type of secondary storage
- Avoids HSM massive unintentional recall
- Transparent
  - Migrated files look the same as online files
  - Clients only access the primary storage on Celerra
- Automatic, user-driven data recall to primary storage in real time (if desired)
- Multi-protocol solution
- Multi-tier hierarchy
- Virtually unlimited file system for the Celerra

Handling Backup

- NDMP and CIFS-based backups automatically back up offline inodes on Celerra
  - Option to backup content through the Celerra as well if desired (NDMP option, CIFS Backup Operator Group integration)
  - Allow offline backups and offline restores
- Significant reduction in primary backup window
  - Secondary storage can be backed up less frequently

CIFS Specific Enhancements

- CIFS Offline Attribute
  - Generated by DART CIFS server if file is offline
  - CIFS clients know if a file is online/offline
  - Increase timeout of the client

- CIFS Offline Notification
  - Popup sent to CIFS client to warn the “human” user
  - Timeout is a parameter
  - Customizable message

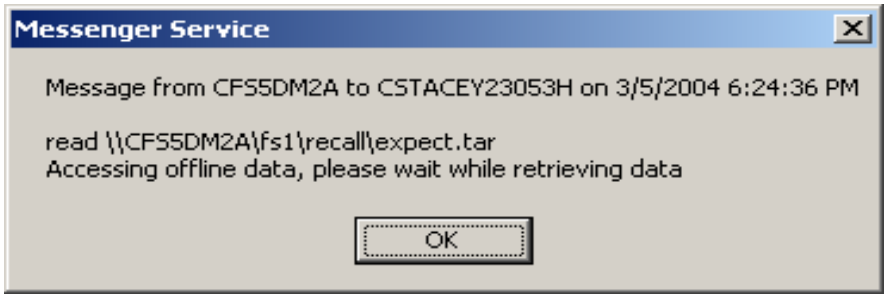

Celerra Distributed HSM Summary

- Enabling technology for building an Information Lifecycle Management solution with the Celerra
  - Policy-driven
  - Distributed
  - Open
  - Migration and management at the file level
  - Data location transparency
  - Cost and feasibility of backup

Questions?

EMC²
where information lives