

*Oxford University TSM Symposium 2007*

*Tivoli Storage Manager : Preparing the Path  
St Catherine's College, 25-27 September 2007*



| IBM OpenStorage Software

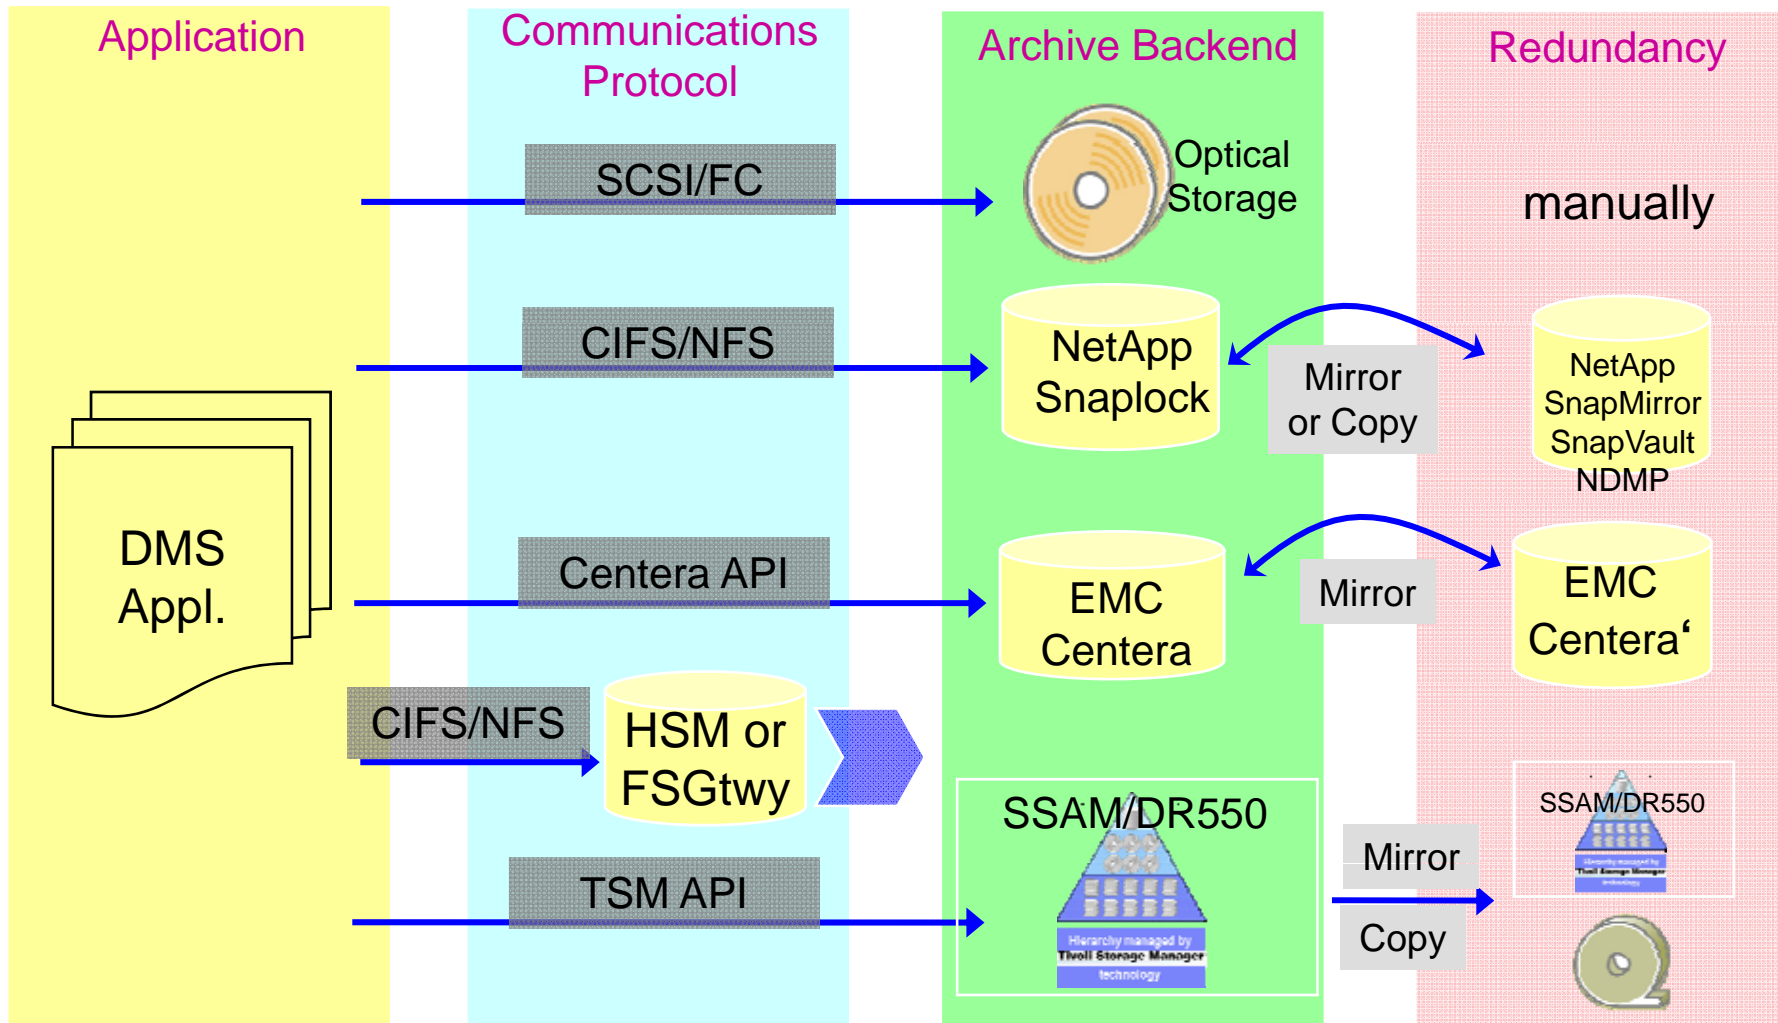
## **System Storage Archive Manager/DR550**

- Basics
- Updates
- Lessons learned with ISVs implementing the TSM-API
- Complementary attachments

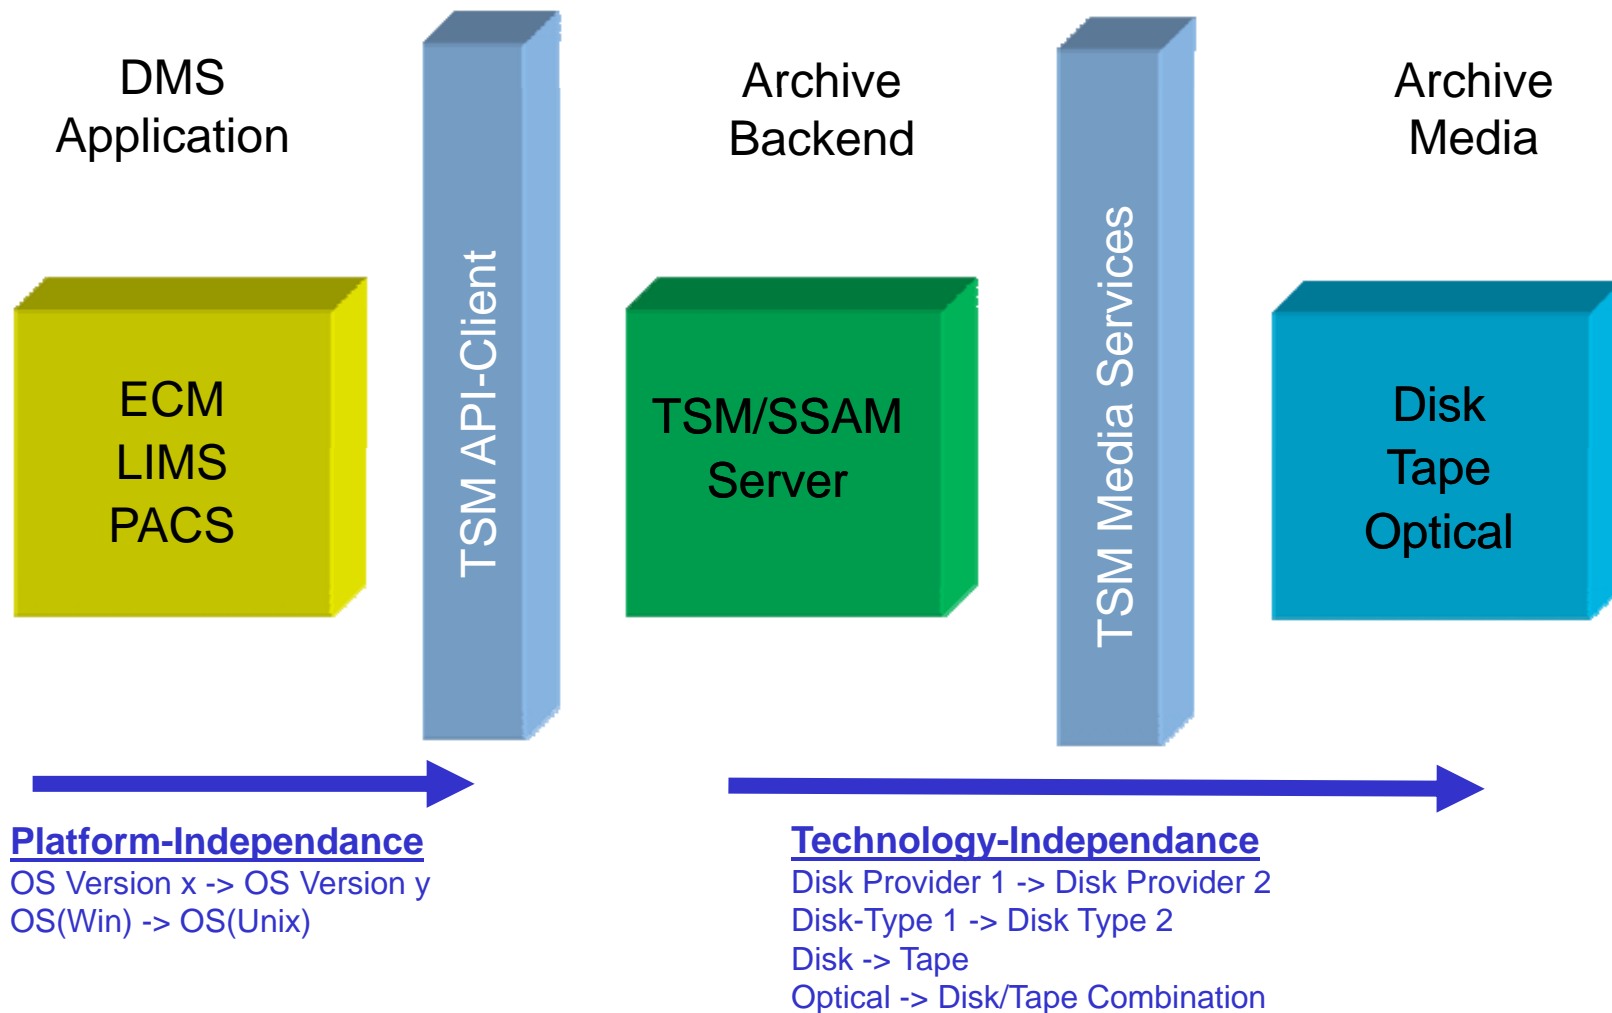
September 2007

Klemens Poschke  
IT Architect IBM Softwaregroup  
[klemens.poschke@de.ibm.com](mailto:klemens.poschke@de.ibm.com)

## Different Archive-Backend Architectures in the market



## TSM/SSAM solves compatibility issues for long-term archiving

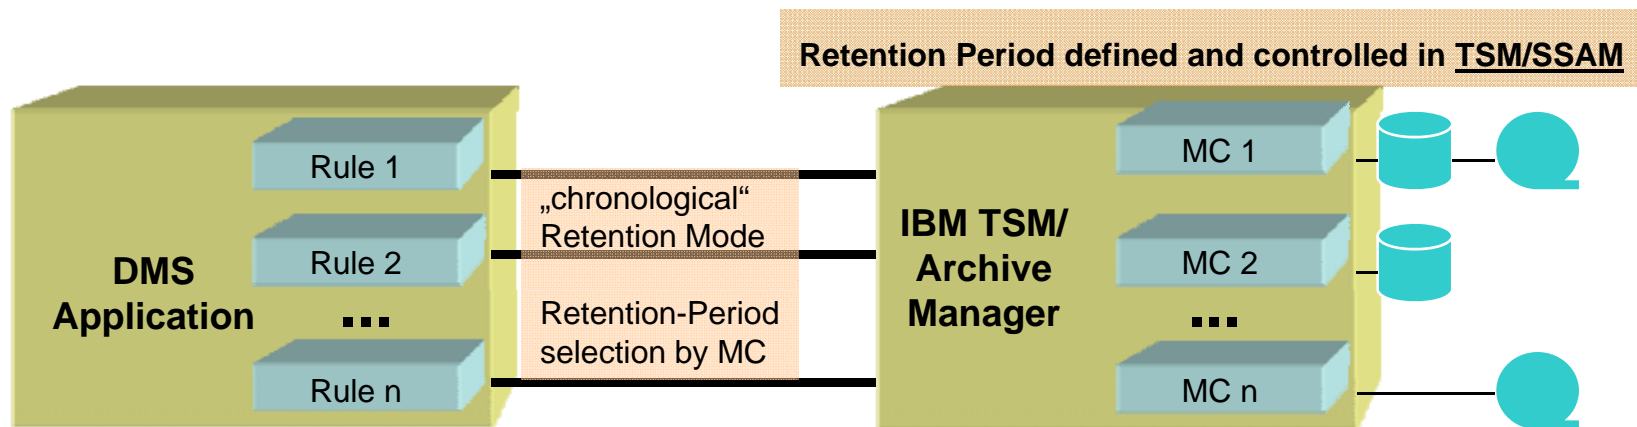


## How Does SSAM Differ from TSM EE?

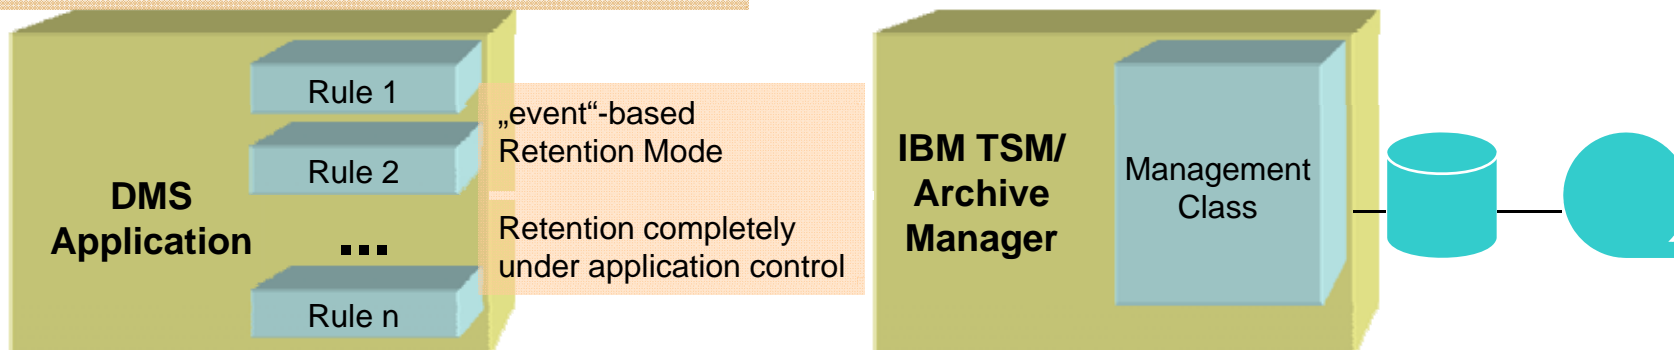
Function	TSM Extend Edition	Archive Manager
Install	TSM EE	TSM EE + "set archiveretentionprotection on"
Devices supported	Over 400	same + Centera in compliance mode and NetApp Snaplock
Server to Server backup	Yes	No
Library Sharing	Yes	Yes
Client data	Backup-, Archive-, HSM-Client	Archive Client
Open API calls	Backup , Archive	Archive
Import/Export data	Yes	Import No, Export w/o keeping retention values
Delete object, filespace, node	Yes	No
Lower archive retention period	Yes	No
Licensing	Per Processor	Per Terabyte stored in primary StgPools

# Retention-Methods

## Backend-controlled vs. Application-controlled

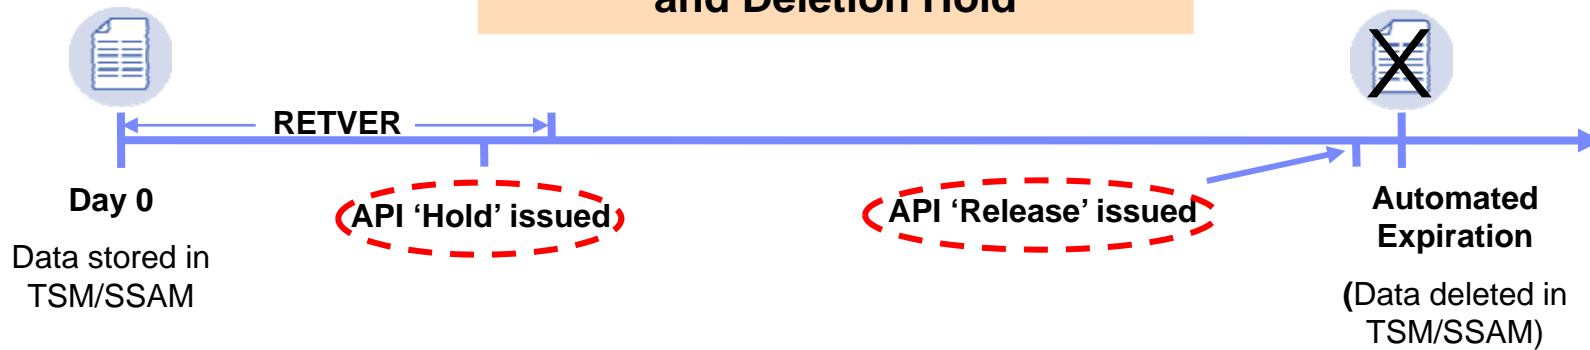


**Retention Period defined and controlled in Application**

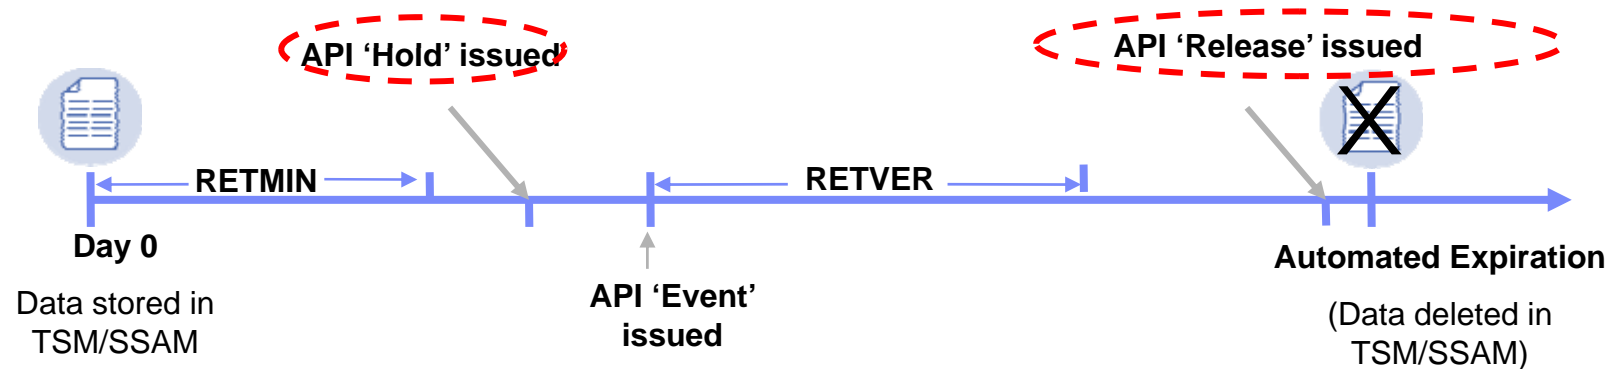


# Data Retention Policy Summary

## Chronological Retention Policy and Deletion Hold



## Event-Based Retention Policy and Deletion Hold





| IBM OpenStorage Software

# TSM/SSAM V 5.4 new features and device support

## Enhanced Security - Application Managed Encryption

- **TSM provides simple key management**
  - Available for TSM API applications; e.g. DP Exchange, DP SQL,...
  - TSM API Client generates a random encryption key password
    - Per data object selected by “include. encrypt” and “exclude. encrypt” criteria
  - TSM API Client passes encrypted data object and encryption key password to TSM Server across “the wire”
    - TSM Client-Server “Session Key” used to encrypt/decrypt
    - AES128 encryption used if BOTH Client and Server support it
  - ***DSM\_ENCRYPT\_CLIENTENCRKEY*** – Key is stored in TSM/SSAM Server
  - ***DSM\_ENCRYPT\_USER*** – Key is maintained by application
  -
- **Future – Backup-Archive client simple key management support (TSM 5.5)**



## TSM and Tape Encryption with TS1120


- TSM will have full support of TS1120 drive encryption beginning with version 5.3.4 on AIX, Linux, Solaris, and Windows. This support will require device driver upgrades and firmware upgrade on TS1120
- TSM on zOS will not support the application method. It will only support the system method of encryption.
- TSM requires all drives in the logical library (partition) be configured to use the same key management method (Application, System or Library)
- Volume Queries will show which key manager/method was used on a cartridge (Application, System or Library).
- Limited to TSM backup/archive data only
  - Does not encrypt Backupsets, TSM DB Backup or TSM Export tapes

```
DEF DEVCL ... TYPE=3592x ...DRIVEEncryption ON | ALLOW | OFF
```

ON = *Application Method*

ALLOW = *LIBRARY or SYSTEM Method*

## Shredding of Data Stored on Disk



- Allows disk storage pools to be designated as “shreddable”
  - Random-access disk only (not “filepool” or tape or other media types)
  - Only Primary Pools
  - up to 10 sequential overwrite processes
- When a data object is moved or deleted from a shreddable pool, TSM server overwrites the object
- Sensitive data objects are destroyed when deleted/moved, preventing undesirable data discovery

**AdminCmd:** *SHRED DATA ...*

**STGPool Definition:** *DEF/UPD STGP ... SHRED=*n* (0-10)*

**STGPool Operation:** *BACKUP STGP or MOVE DATA with SHREDTONOSHRED Parameter*


**Data Movement:** *EXPORT NODE/SERVER or GENERATE BACKUPSET with ALLOWSHREDDABLE Parameter*

**TSM Server Option:** *SHREDding AUTOMATIC | MANUAL*



| IBM OpenStorage Software

## DR550 update



## IBM System Storage DR550 Family

### ■ DR550

- Single DR550 :
  - 500GB drives: 8 , 16 TB
  - 750GB drives: 12, 24 TB (new)
- Dual DR550:
  - 500GB drives: 8, 16, 32 TB (dropped 56 and 112TB )
  - 750GB drives: 12, 24, 48 TB (new)
- Choice of Ethernet connections
- Synchronous or asynchronous replication options
- Tape ready
  - ✓ Optional WORM (TS1120 or LTO-3)



**IBM DR550**



**IBM DR550 Express**

### ■ DR550 Express

- 500GB drives: 1.1, 5.1 , 9.1 TB (same as before)
- 750GB drives: 1.1, 7.1, 13.1 TB (new)
- Choice of Ethernet connections
- Optional 25U rack with lockable doors
- Tape ready
  - ✓ Optional WORM (TS1120 or LTO-3)

## New major features for DR550

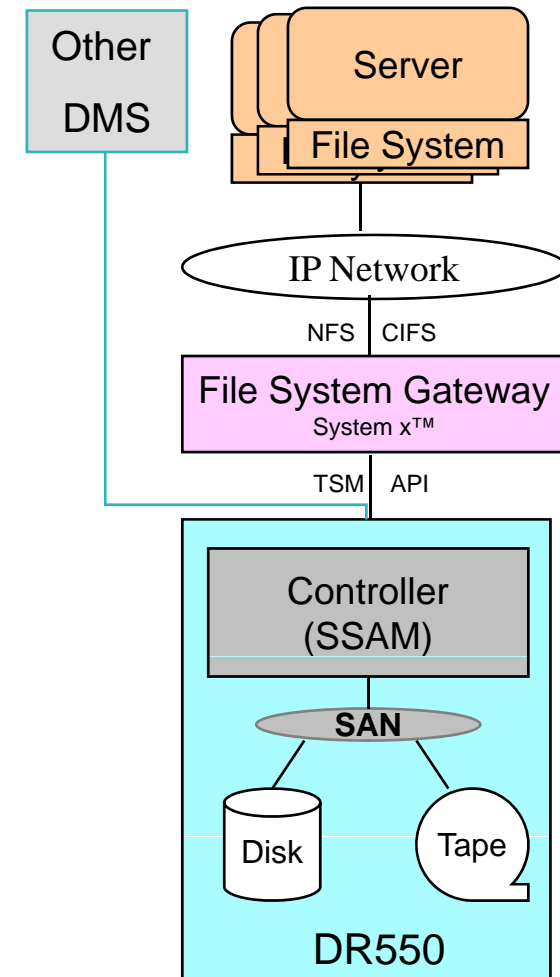

- DR550 V3.5 (Jan07)
  - System Storage Archive Manager V5.4
  - Data shredding to overwrite expired objects
  - Tape encryption key management support
  
- DR550 V4.0 (Jun2007)
  - DR550 File System Gateway
    - NFS and CIFS network file access to the DR550



TS1120 with data encryption

## DR550 Gateway Overview

- **DR550 File System Gateway**
  - System x™ with SLES 10 and special Software preconfigured (OEM from ByCast™)
  - Providing NFS and CIFS interface
  
- **Compliance Implementation**
  - Retention period based on file-path-pattern
  - WORM protection of “stubs” in Gateway
  
- **Advanced Functions**
  - Backup and restore of metadata
  - Caching for frequently needed files
  - Pre-Staging of entire directories
  - High availability option with two gateways and dual node DR550



## DR550 Certification by Certified Public Accountants

- ✓ Germany (KPMG)
- ✓ Austria (KPMG)
- ✓ Switzerland (KPMG)
- in work (KPMG):
  - Danmark
  - Norway
  - Finnland
  - Iceland
  - Sweden
  - France (not by KPMG)
- ✓ USA (Cohasset)
- planned in Europe (KPMG):
  - Spain
  - UK
  - Poland



Prüfungsbericht

über die Prüfung des IBM TotalStorage Data  
Retention 550 Speichersystems (IBM DR550)  
in der Version 3.0

IBM Deutschland GmbH, Stuttgart

Mai 2006



IBM OpenStorage Software

ISV-Part





## ISV support for IBM SSAM/IBM DR550

### ■ ISV's supporting IBM DR550 (by Aug2007)

- Princeton Softech Optim 6.2
- BrainTribe (Formerly Comprendium)
- Bycast Inc.
- CaminoSoft
- Ceyoniq
- d.velop AG
- Easy Software
- Enigma Data Systems
- Gamma Systems
- Heilig & Schubert (H&S)
- Hummingbird (now OpenText)
- Hyland Software (OnBase)
- Hyperwave
- PBS
- Waters
- Betasystems (cert. pending)
- Interwoven via Micro Strategies Connector
- MBS Technologies
- OpenText / iXOS
- Saperion
- SER Solutions
- Triade
- Symantec / KVS EV
- Windream
- Zantaz
- Solix Technologies
- Stellar Technologies
- RJS Software Systems
- Lighthouse Global Technologies

Ready for

e-business software

Business Partner products that have been tested and validated to meet key integration standards defined by IBM.

## Experiences with ISV implementations

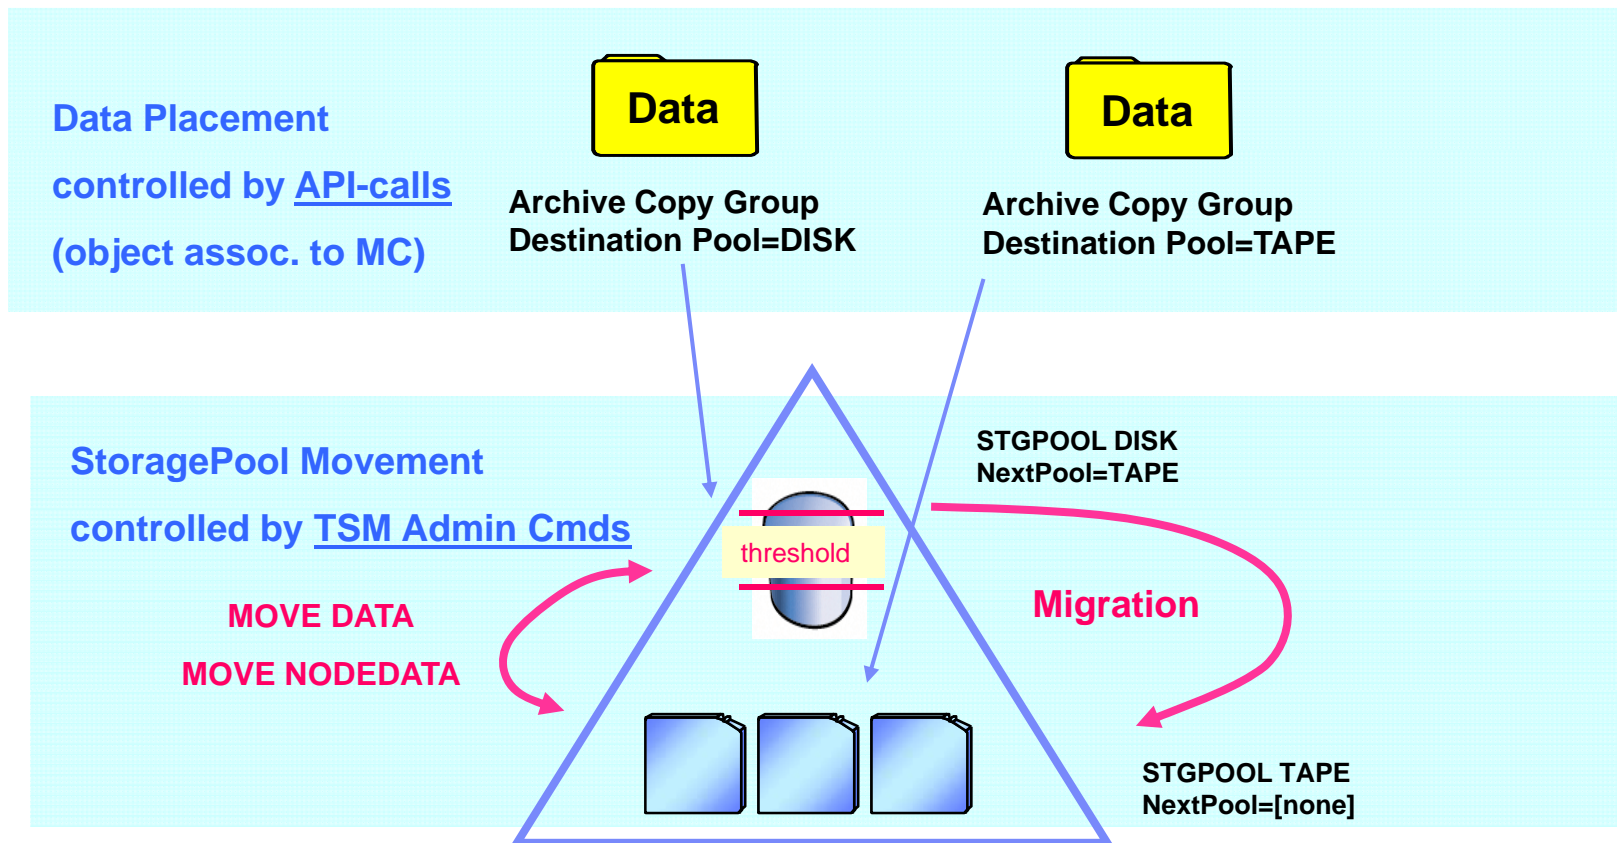
- **DMS-ISVs love the SSAM approach to separate application intelligence from archive backend storage via API calls**
  - no effort to code & test storage device drivers
  - minor modification of their current approach how to write/read data
  - easy mapping of the *platter/folder/data* structure to TSM object addressing scheme *FS/HL/LL*
  - easy technology migration, not bound to specific providers on the backend side
  - cross-platform access of TSM-API written data (restrictions on UNICODE)
  - easy integration in customer's TSM backup infrastructure and co-existence of different DMS applications on the same archive back-end
  - complement their offering with vaulting capabilities for free

## Experiences with ISV implementations

### ■ However:

- simple substituting WRITE/READ commands with TSM API calls does not fully exploit TSM capabilities
- real improvements if applying architectural changes:
  - ISO image vs. single object
  - performance improvements through multithreading, multi-object transaction, use of ASNODE option
  - Object-Ids might not be unique over the object life-cycle (e.g export/import)
  - TSM provides options to store free text together with the object (Archive description, ObjInfo) for search/index operations
  - exploitation of new „event“-based retention and „hold/release“ functions
  - exploit the tiered-storage capabilities of TSM

# TSM Object Placement and Storagepool Movement



**MIGRATION = autom. moves all data from stgpool x to stgpool y within borders (one-way operation)**

**MOVE DATA = moves all data of all nodes from stgpool x to stgpool y (back/forth operation)**

**MOVE NODEDATA = moves all data of selected node from stgpool x to stgpool y (back/forth operation)**

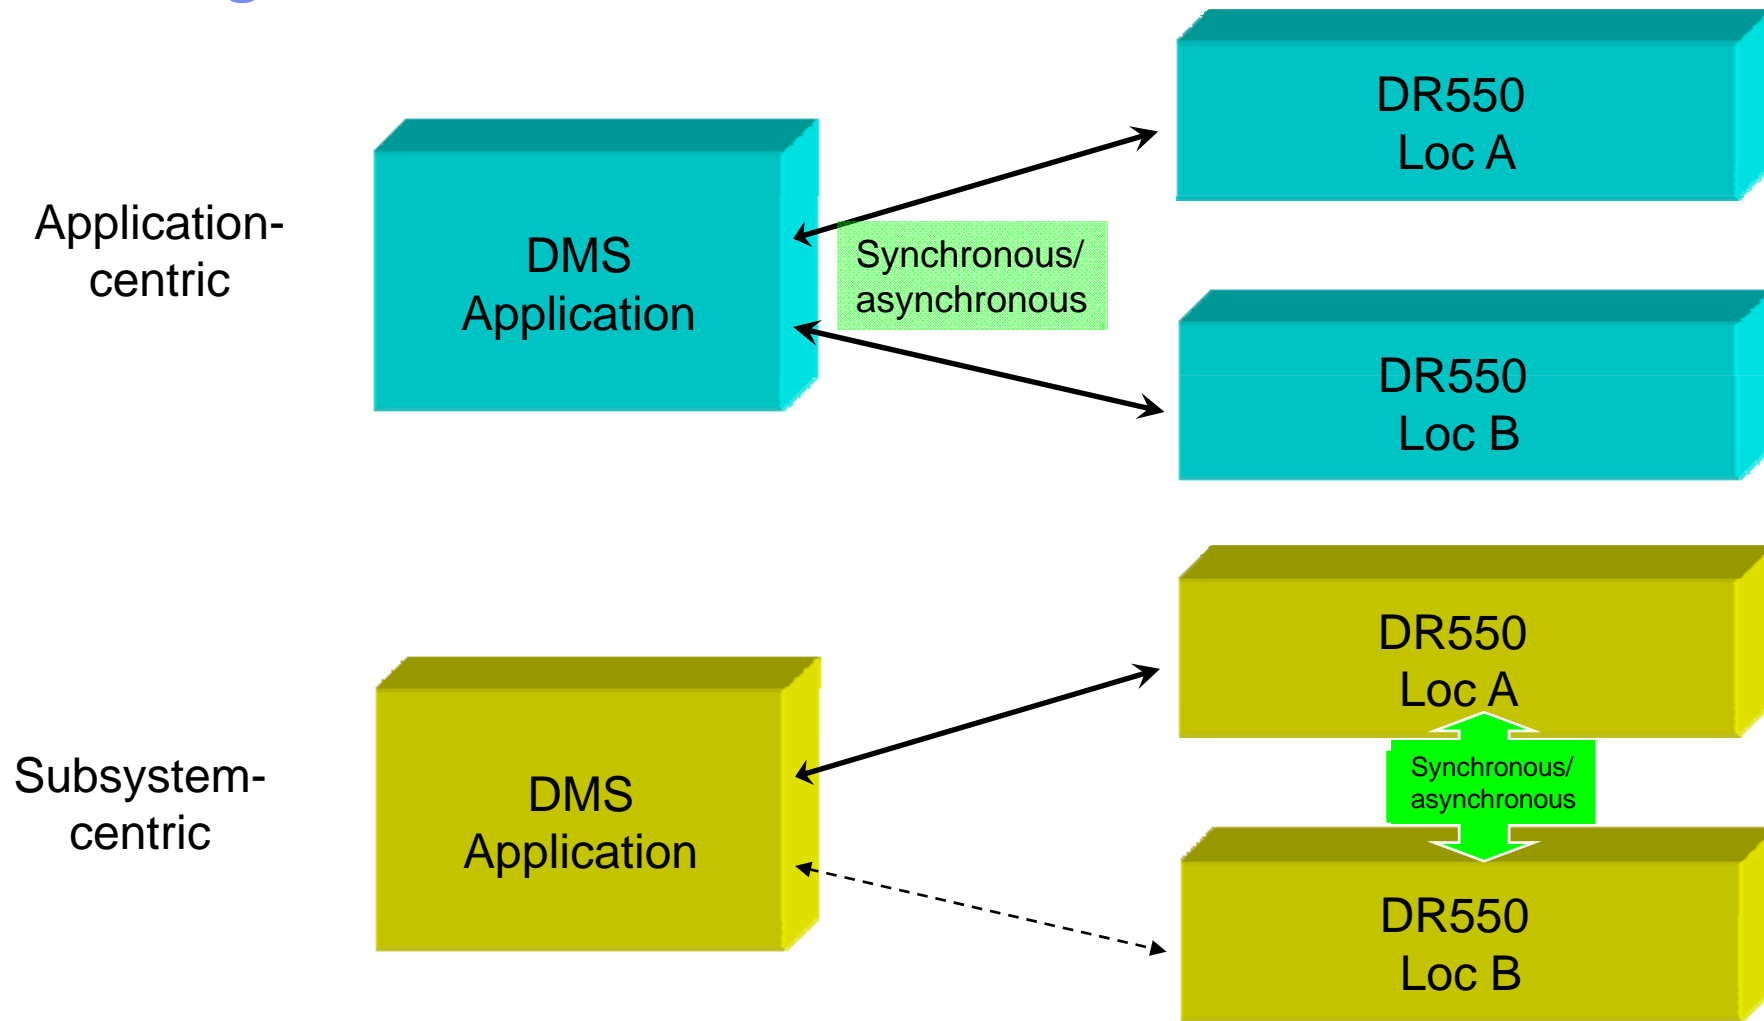
## Experiences with ISV implementations

- **Tiered-Storage architecture (disk->tape) of TSM/SSAM is currently not supported by some ISV applications**
  - full access to object placement/movement under control of the TSM-API is requested
  - wait/deadlock conditions possible on sequential media (tape) if concurrent access of many users (typically in ISO-image based architectures)
  - no TSM admin interface available with direct access from the API client to query object placement details in TSM storage hierarchy
  - these enhanced functions are needed to optimize parallel user access controlled by the DMS application

## Planning considerations for ISV application integration

- **In a TSM world typically a database is separated into two instances if TSM DB-size affects throughput – via export/import processes**
- **SSAM/DR550-Server can not be divided by export/import processes, this requires a careful sizing before installation**
  - Maximal SSAM DB-size must be planned  
(rule-of-thumb: 1 archived object with 1 media copy approx. 0.7kB in SSAM-DB)
  - recommended DB-sizes in archive environments up to 300 GB (~ 450 Mio Objects)
  - instance separation and load balancing of SSAM is only possible via migration processes initiated/controlled by the application
  - platform changes of SSAM-server only via application migration
  - horizontal scalability has to be supported by the DMS-application e.g. via extending the API object addressing scheme  
[SSAM\_Server/Filespace\_name/HL\_name/LL\_name]

## Mirroring alternatives with SSAM/DR550



**In multi-DMS application environments incompatible mirroring capabilities must be adjusted to the most common level → DR550 subsystem mirroring**




| IBM OpenStorage Software


## TSM/SSAM support consideration for EMC Centera and NetApp Snaplock



## TSM/SSAM and Centera



- With TSM EE there is no retention protection
- only support of Centera in „non-compliance-mode“
- retention controlled by Centera global settings „no“ or „forever“
- means: either no protection or no deletion



- With SSAM there is retention protection, (Centera in „compliance“ mode)
- Centera does not automatically expire data, but allows deletion after expiration
- Binding of retention attributes to individual objects

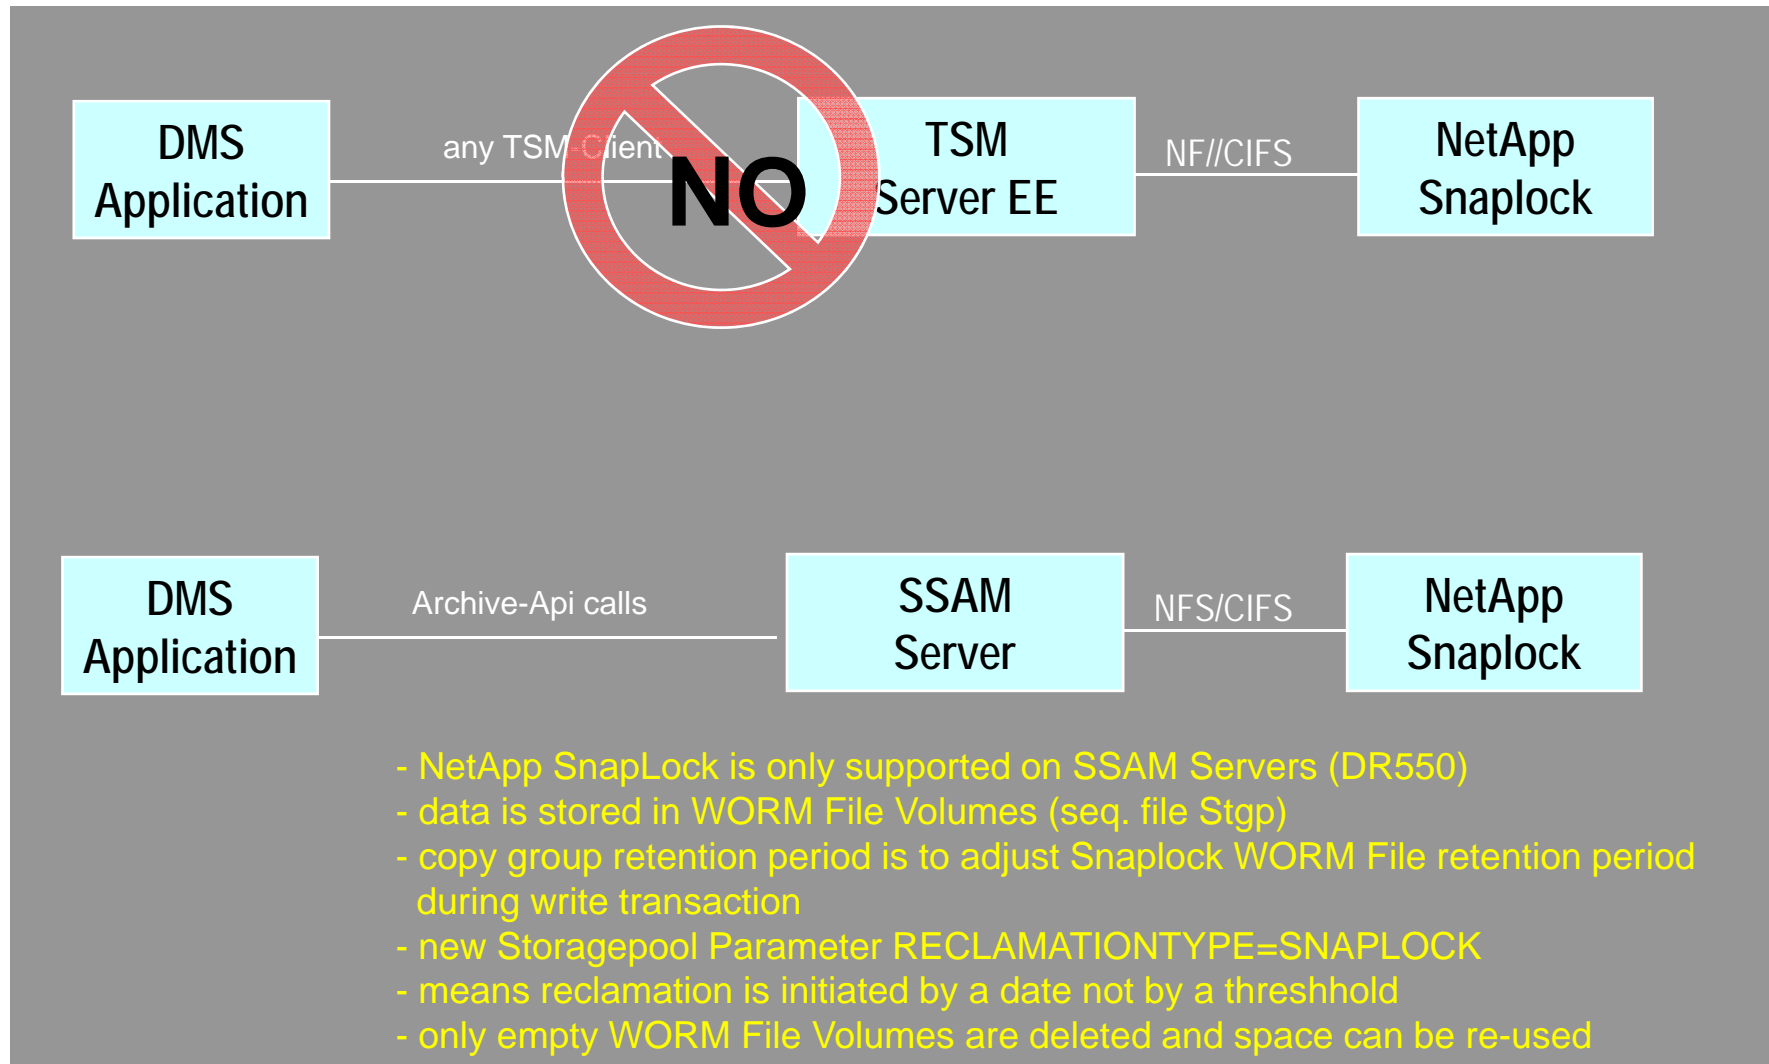
## TSM – Centera Deviceclass: Unsupported Operations

```
DEFine DEVclass device_class_name DEVType = CENTERA .....
```

- „migration“
- „reclamation“
- „move data“ into or out of a Centera storage pool.
- „move nodedata“ into or out of a Centera storage pool.
- „backup storagepool“ of Centera storage pools.
- „restore volume“ of Centera storage pool volumes.
- Defining Centera volumes
- Using a Centera device class to **back up a database**.
- Using a Centera device class for **database loading or unloading**.
- Using a Centera device class as the target of volume history, device configuration, trace logs, error logs, or query output files.



# TSMSSAM and NetApp Snaplock



## SSAM/SnapLock integration


- SSAM uses SnapLock Volumes as „container“ in SSAM „file-pool“
- A Snaplock volume is filled by „write append“ until completely filled – ensures no overwriting and no deletion of objects by admin. Full volumes are set R/O in SnapLock filesystem.
- A SnapLock volume can only be deleted if logically empty (e.g. after SSAM reclamation)
- important parameter in SSAM/SnapLock storagepool definitions is:  
***RECLAMATIONTYPE=SNAPLOCK*** (*reclamation by time*)
- fragmented volumes (fragmentation by objects being eligible for deletion because retention period is over) will be reclaimed by SSAM, objects to be retained are copied into a new volume and the previous volume is deleted in SnapLock Filesystem
  - *minimum retention period = 30 days (default, >=60 days recommended)*
  - *maximum retention period = 30 years (default)*
- Snaplock uses NFS/CIFS „last access date“ attribute for storing retention date with the object



## Info sources

- SSAM
- <http://www-306.ibm.com/software/tivoli/products/storage-mgr-data-reten/>
- DR550 Home page:
- <http://www-1.ibm.com/servers/storage/disk/dr/index.html>
- DR550 Users Guide:
- <http://www-1.ibm.com/servers/storage/pguide/index.html>
- DR550 Redbook:
- <http://www.redbooks.ibm.com/redpieces/abstracts/sg247091.html>
- TSM API Documentation
- <http://publib.boulder.ibm.com/infocenter/tivihelp/index.jsp>
- TSM API Code – free download, including sample program with source code
- <ftp://ftp.software.ibm.com/storage/tivoli-storage-management/patches/client/>

# Questions?



Thanks for attending....