



IBM Software Group

Technical Perspective on Future Enhancements

Dave Cannon
IBM Tivoli Storage Management Development
Oxford University TSM Symposium 2003



© 2003 IBM Corporation

IBM Software Group | Tivoli software



This presentation describes future plans and strategies for IBM Tivoli Storage Management software products. IBM reserves the right to change product plans according to the Company's business needs and requirements. These are not commitments to deliver the described features within a given time period. We reserve the right to change product features and time periods to accommodate IBM's business objectives.



Topics

- Management of data in server storage
- Optimizing recovery of client data
- Database management

Key

- Candidate – more likely to be delivered
- Consideration – less likely to be delivered

Management of Data in Server Storage

- Collocation groups
- Multi-process reclamation
- Simultaneous migration and storage pool backup

Collocation Considerations



- No collocation
 - Fragmentation of data for each node
 - Long restores

- Node size similar to volume capacity
 - Efficient tape utilization
 - Fast restores

- Node size much less than volume capacity
 - Inefficient use of tapes and library slots, if tapes are dedicated to a node OR
 - Random mixing of nodes on each tape leading to inefficient internal data-transfer operations to maintain collocation
 - Excessive mounting of target tapes
 - Multiple passes of input tapes

Collocation Groups

Define groups of nodes whose data will be collocated together on sequential media



Candidate for 2H2004



Benefits of Collocation Groups

- Collocation of small nodes without requiring that a tape and library slot be dedicated to each node
- Optimal recovery
 - Efficient collocation of small nodes
 - Increased efficiency for multi-session restore by spreading data for a node over multiple volumes
 - Possible collocation of copy storage pools for offsite storage
- Improved efficiency for internal data-transfer operations by transferring all nodes in the group together
 - Minimizes mounts of target volumes
 - For sequential-to-sequential transfer (e.g., reclamation), minimizes database scanning and reduces tape passes

Collocation Groups and Storage Pool Collocation

	Storage pool collocation attribute	Node not defined in a collocation group	Node defined in a collocation group
Collocation granularity ↓	NO	No collocation	No collocation
	GROUP	Collocation by node	Collocation by group of nodes
	NODE	Collocation by node	Collocation by node
	FILESPACE	Collocation by file space	Collocation by file space

Managing Collocation Groups

- DEFINE COLLOGROUP
- UPDATE COLLOGROUP
- DELETE COLLOGROUP

- DEFINE COLLOCMEMBER
- DELETE COLLOCMEMBER

- QUERY COLLOGROUP
- QUERY NODE

- QUERY NODEDATA
 - Filters by
 - Node or collocation group
 - Storage pool
 - Volume name
 - Output
 - Node name
 - Volume name
 - Occupancy
 - Storage pool

- MOVE NODEDATA (consolidate data for a group)

Multi-process Reclamation


- Support multiple concurrent reclamation processes for same sequential-access storage pool

- Maximum number of processes configurable

- Would allow optimal use of tape drives

- Candidate for 2H2004

Simultaneous Migration and Storage Pool Backup



- Can reduce backup window as compared to simultaneous write to primary and copy storage pools during backup operation
- Overlap windows for migration and storage pool backup

Optimizing Recovery of Client Data

- Disk storage enhancements
- Collocation of active data
- Restore of individual files from file system image



Disk Storage Enhancements

- Disk-to-disk backup
 - Recent industry attention driven by availability of high-capacity, inexpensive disk subsystems
 - TSM (formerly ADSM) has supported disk storage for 10 years
 - Ongoing product enhancements to optimize disk storage with TSM

- Potential enhancements for sequential-access disk pools
 - Allow storage pools to span multiple file systems
 - Support parallel migration
 - Minimize fragmentation of volumes to improve performance
 - Support direct I/O for performance

- Enhance random-access disk pools to support reconstruction of aggregates


- Consideration for 2H2004



Disk Storage Enhancements (continued)

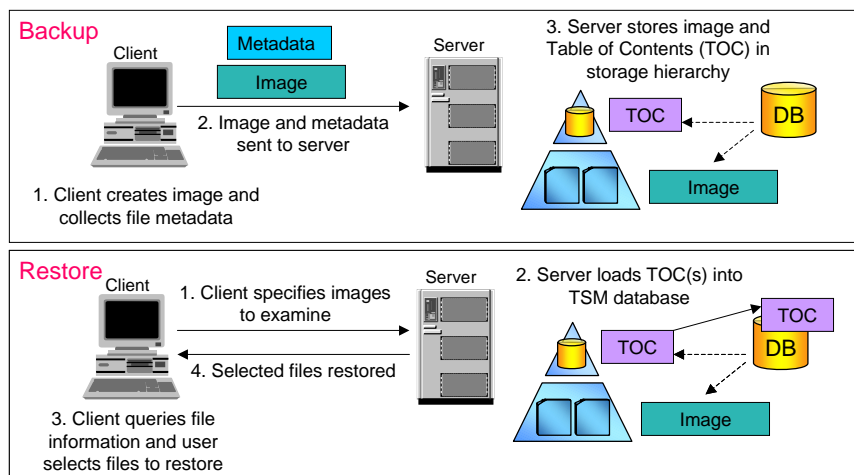
	Random-access	Sequential-access
Space allocation/tracking	Random 4KB blocks	Sequential within volume
TSM caching	Optional (backup overhead)	Not supported
Space recovery (no cache)	When file is deleted/moved	When volume is reclaimed
Recovery of cache space	When space is needed	Not applicable
Aggregate reconstruction	Not supported	During volume reclamation
Concurrent volume access	Yes	Not supported
Multi-session client restore	One session for all volumes	One session per volume
Target for LAN-free backup	Not supported	Yes, via SANergy
Can be used for copy pools	No	Yes
Migration/stg pool backup	By node and file space	By volume
Parallel migration processes	Yes	No
Migration to collocated pool	Very efficient	Moderately efficient
Pools can span file systems	Yes	No
Database regression	Must audit all volumes	Reuse delay avoids audit

Collocation of Active Data



- Possible benefits
 - Optimized access to active versions
 - Reduced size of disk pools if only actives are stored
 - Reduced data movement in preparation for restore of active data
 - Ability to store only active files in copy storage pool
- Challenge is to avoid compromising performance for internal data-movement operations
- Consideration for 2H2004

Restore of Individual Files from File System Image



Database Management

- Analysis tool for database reorganization
- Replacement of proprietary TSM database with DB2

Analysis Tool for Database Reorganization

- Reorganization of TSM database consolidates fragmented pages
 - Reduces size of database and can improve performance
 - Offline operation
- New database tool will
 - Initiate long-running scan of all database tables
 - Output estimate of recoverable space from database reorganization
- Recommendation: Perform estimate infrequently and only when server activity is low
- Function will be included in 5.2.2 server (4Q03) but not documented
- To evaluate, contact TSM support
- Candidate for 2H2004

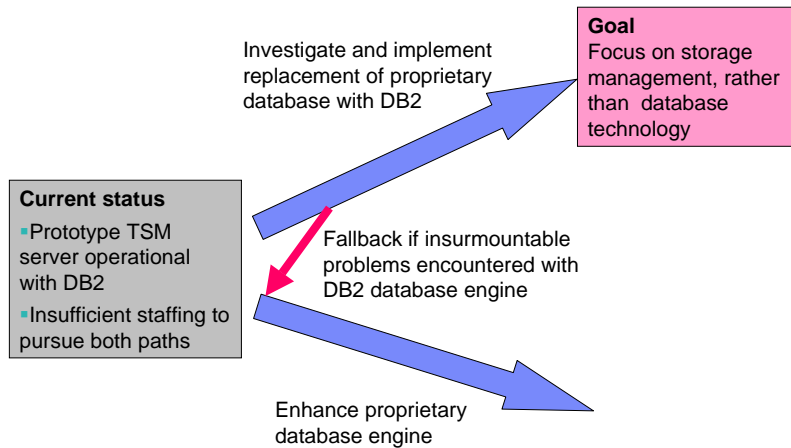
Past History: Why a Proprietary TSM Database?

- Portability to diverse server platforms
 - Proprietary TSM database portable to any server platform
 - In early days of ADSM, DB2 had only limited platform support
- TSM database optimized for storage-management operations
 - Database schema
 - Customized, multilevel locking
 - ARIES technology reduces overhead for logging and recovery
- Bit vectors stored in database for tracking block allocation in random-access storage

Why Consider DB2 as TSM Database Now?

- DB2 has evolved
 - Available on many platforms
 - Includes ARIES technology
- Use of DB2 would allow exploitation of capabilities such as
 - Online database reorganization
 - Referential-integrity checking to eliminate or minimize need for database audit
 - Complete SQL function including UNION, INTERSECT, and correlated sub-queries not available with proprietary database
 - Ability to optimize SQL query performance by defining secondary indices
 - Distributed database for operations such as LAN-free and library sharing
- Feasibility of bit vectors still a question with DB2, but would not be necessary if random-access disk is replaced with sequential disk

Which Way?



Customer Expectations for Transition to DB2

- No extra expense for database product license
- TSM administration should not become more complex (but may change)
- If conversion of existing database entries is required, this should be
 - Automated
 - Not require an unreasonable amount of TSM downtime
- Performance similar to that of current proprietary database engine

Further Investigation

- Performance with DB2 database engine
- Migration from existing proprietary database to DB2
- Backup and recovery of the database
- Elimination of database audit by using referential-integrity rules
- Bit-vector support for random-access disk storage pools

Alternative Database Migration Approaches

