Tivoli Storage Manager Scalability Enhancements

Dave Cannon
Tivoli Storage Management Development
Oxford University TSM Symposium
September 2001

Agenda

- Recent enhancements
- Planned enhancements
- Potential future enhancements

Scalability Contributors

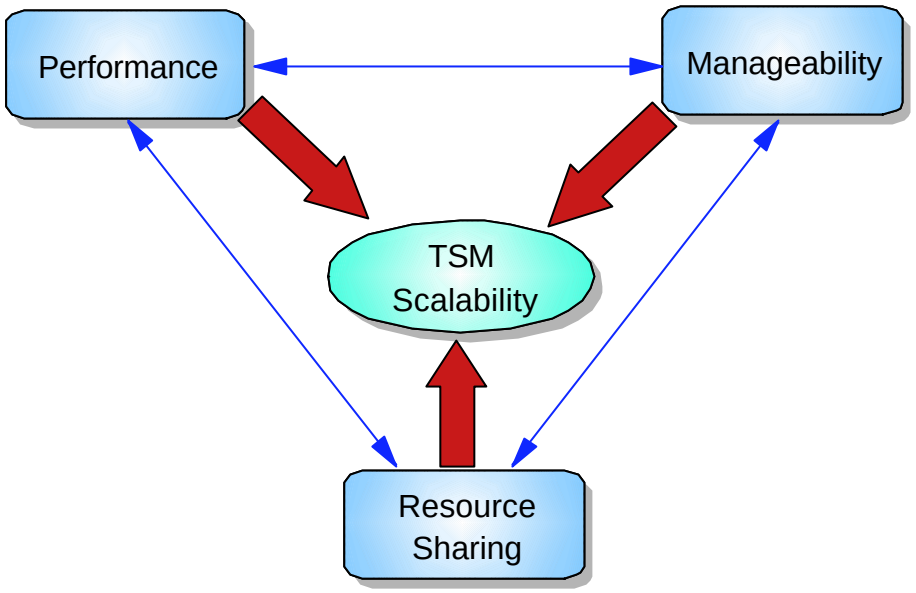

Tivoli

Recent Enhancements

Subfile backup/restore
Tivoli Data Protection for EMC
Tivoli Data Protection for ESS
LAN-free data movement
Journal-based incremental backup
3590 performance for S/390

TDS for Storage Management Analysis
Increased recovery log size
Volume-specific space reclamation
TSM Management Console

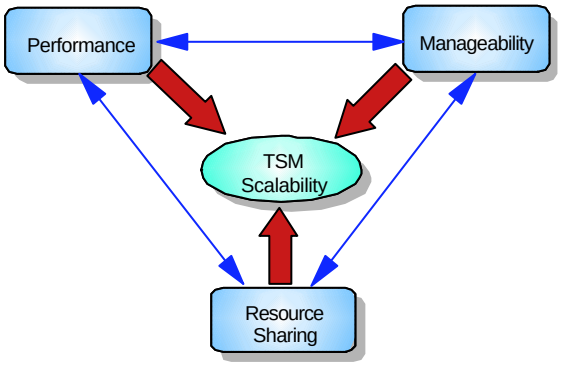

Tape library sharing
LAN-free disk sharing with SANergy
Unicode support

Tivoli

Planned Enhancements

These enhancements are planned

- Move Nodedata command (V5.1)
- Export/import enhancements (V5.1+)
- Simultaneous writes to copy pools (V5.1)
- Windows 2000 image backup (V5.1)
- Multi-session restore (V5.1)
- New concepts (beginning in V4.2.1)
- TDP for NDMP (V4.2.1)
- SCSI-3 extended copy (V5.1)

Plans, schedules, and functional content are subject to change

Tivoli

Move Nodedata Command

Consolidation

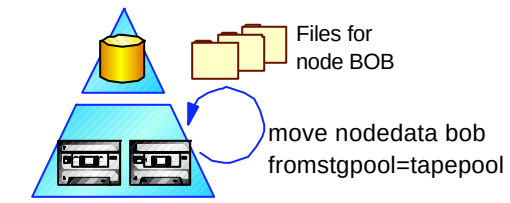

Movement to Disk

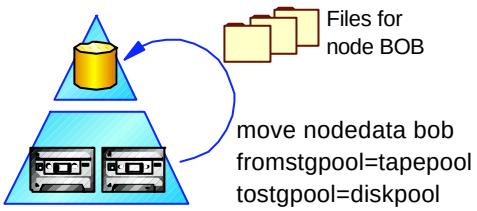

- Moves files belonging to specified node(s) and residing in a specified sequential-access storage pool
- Prepares for rapid client restore by
  - Consolidating node data within a sequential-access pool
  - Moving data to disk for easy access

Planned: Move Nodedata

Tivoli

File-Filtering Options

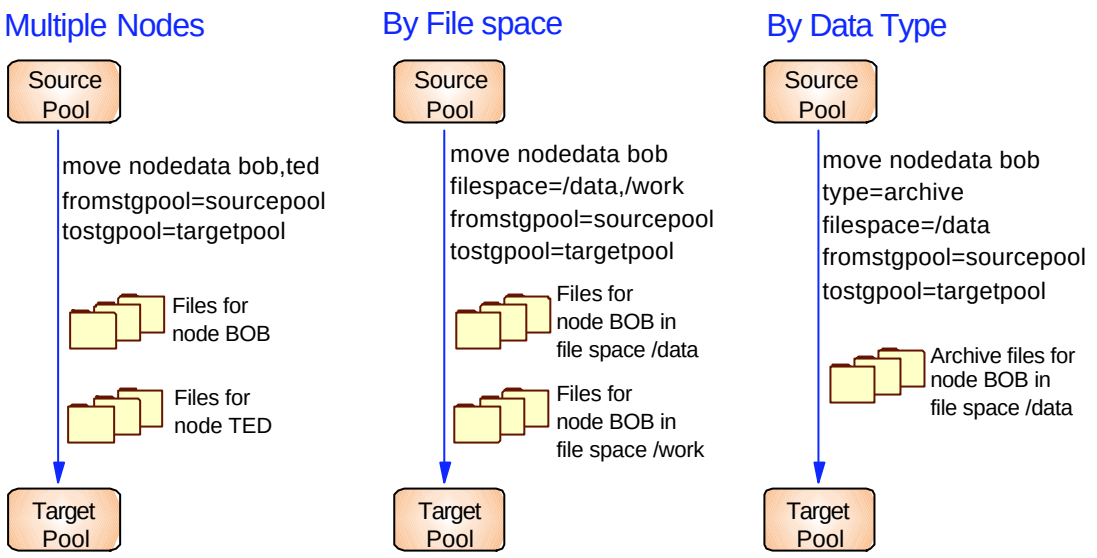

If multiple nodes are specified, all file spaces are moved

Data-Transfer Options

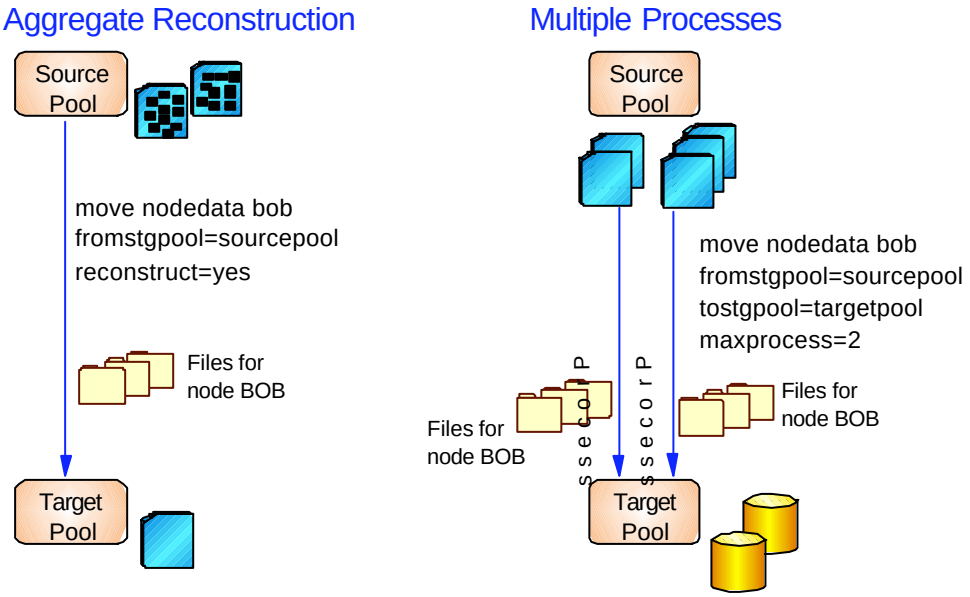

Server-to-Server Export/Import

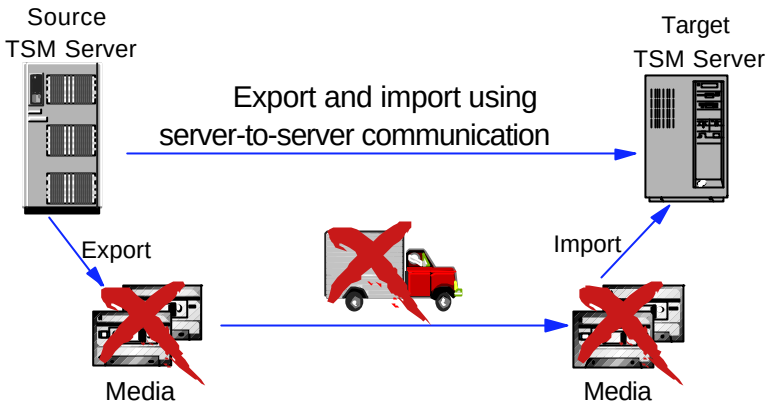

- Export and import are performed in a single operation
- Eliminates need for common media
- Eliminates copies to and from export media
- Avoids management and transportation of media
- Useful for splitting servers and for duplicating servers to achieve disaster recovery protection

Planned: Export/Import

Tivoli

Server-to-Server Export/Import (cont.)

- Initiated with export command on source server
  - Export Server
  - Export Node
  - Export Admin
  - Export Policy
- New export parameters
  - Toserver
  - Previewimport
  - Replacedefs
  - Dates

Planned: Export/Import

Tivoli

Merging of File Spaces During Import

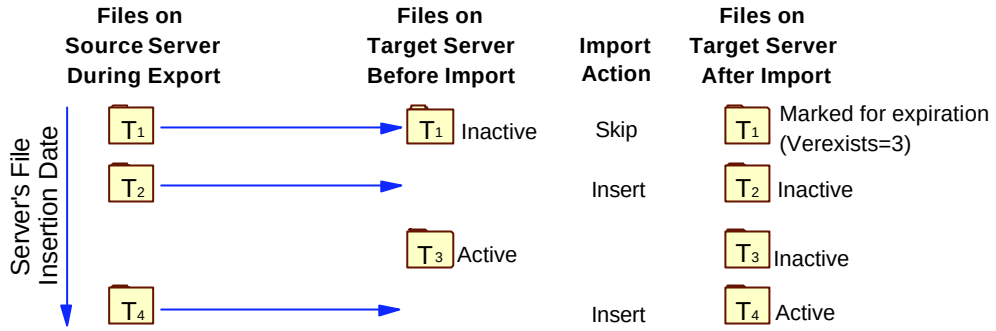

- Mergefilespaces parameter indicates whether import will merge files into existing file spaces on target server or generate new file spaces
- Allows transfer of backup and archive data in stages
- Provides restart capability

Planned: Export/Import

Tivoli

Incremental Export

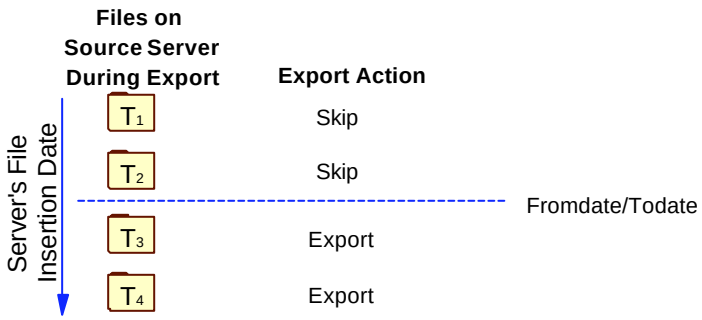

- Fromdate and Todate specify earliest insertion date of files to be exported
- Coupled with file space merging, allows ongoing duplication of data between two servers

Planned: Export/Import

Tivoli

Simultaneous Writes to Copy Pools

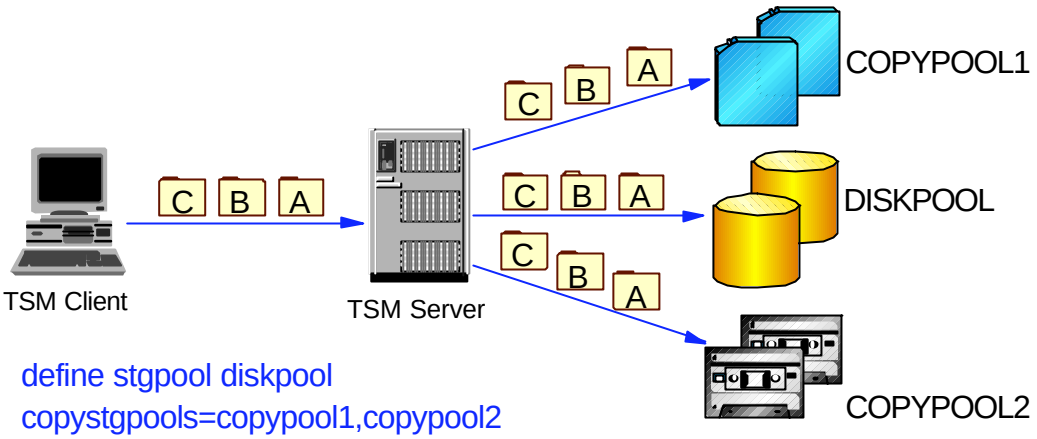

- Simultaneous writes to primary pool and copy pool(s) during client backup, archive, and space-management operations
- Target pools can have different device classes
- Should be used in conjunction with incremental storage pool backup

Simultaneous Writes: Error Handling

- If an output error occurs writing to the primary pool, all writes fail and the transaction rolls back
- If an output error occurs writing to a copy pool
  - Default behavior is to discontinue writing to this copy pool for remainder of the session, but continue storing files into primary pool and any other designated copy pools
  - Optional behavior is to fail all writes and end session (all or nothing)

Windows 2000 Image Backup

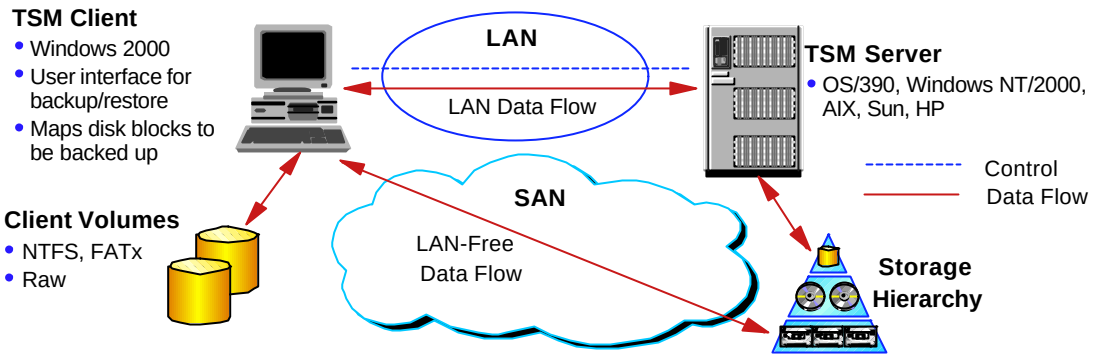

- Optimized for backup/restore of entire file system
- Fast, file-system image backup
  - Minimal overhead for TSM server database
  - Online-backup causes minimal disruption to applications
- Fast restore because no overhead for file creation
- Backup/restore can be over LAN or LAN-free

Tivoli

Windows 2000 Image Backup (cont.)

Online Backup

- Requires Logical Volume Snapshot Agent (LVSA)
- Uses virtual snapshot to create PIT image
- Steps
  1. Quiesce applications (pre-snapshot command)
  2. Begin snapshot
  3. Resume application processing (post-snapshot command)
  4. Back up volume image
  5. End snapshot

Offline Backup

- Does not require LVSA
- Steps
  1. Quiesce applications
  2. Lock volume
  3. Back up volume image
  4. Unlock volume
  5. Resume application processing

Backup of Used Blocks

- In-use blocks are sent to server (possibly interleaved with occasional unused blocks for efficiency)
- Requires file system (NTFS, FATx)

Backup of All Blocks

- All blocks in file system are sent to server
- Used for raw volumes

Tivoli

Multi-Session Restore

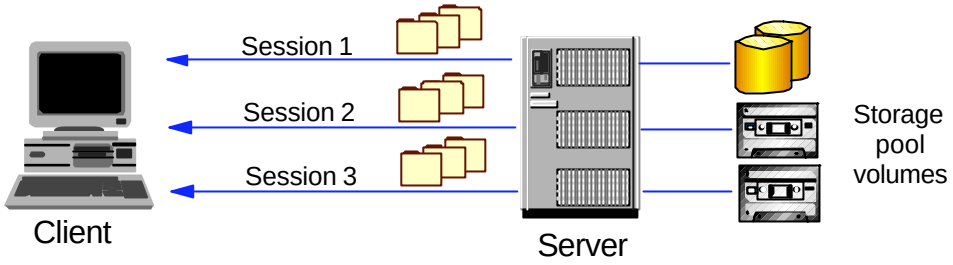

- LAN-based restores using multiple client-server sessions for improved throughput
- Limited by
  - Number of sequential-access volumes with data to be restored
  - Mount points
  - Client's resourceutilization option

Planned: Multi-Session Restore

Tivoli

Concept: Outboard Data Mover

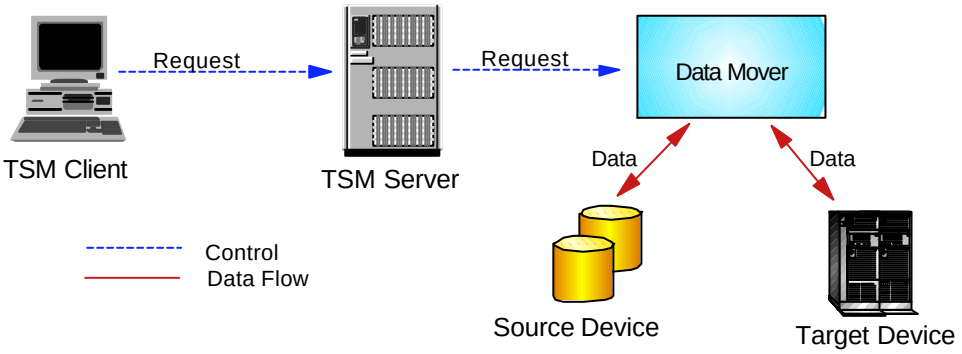

- TSM data mover
  - Named device, external to TSM client or server
  - Accepts request from TSM to transfer data
- Reduces CPU cycles on TSM client and server
- Avoids data movement over the LAN
- Examples: NAS device, SCSI-3 device

Planned: Concepts

Tivoli

Concept: Data Format

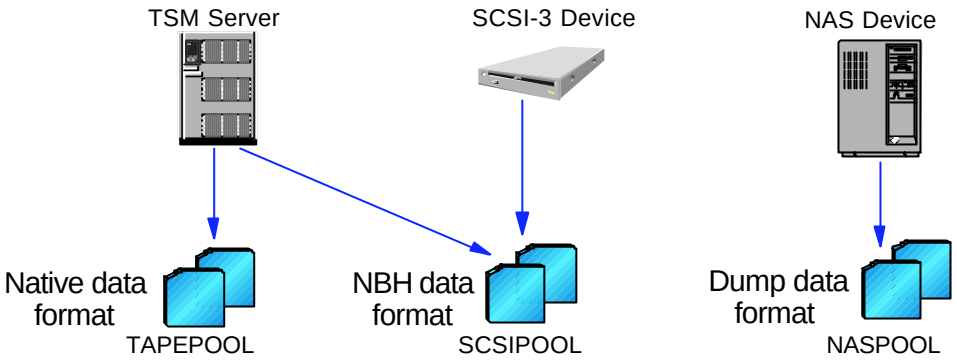

- TSM server stores data in its own "native" format or in NBH format
- Outboard data movers may store data in other formats
- Each storage pool and data mover will have a designated data format
- Certain operations may be restricted for non-native storage pools

Planned: Concepts

Tivoli

Concept: Path

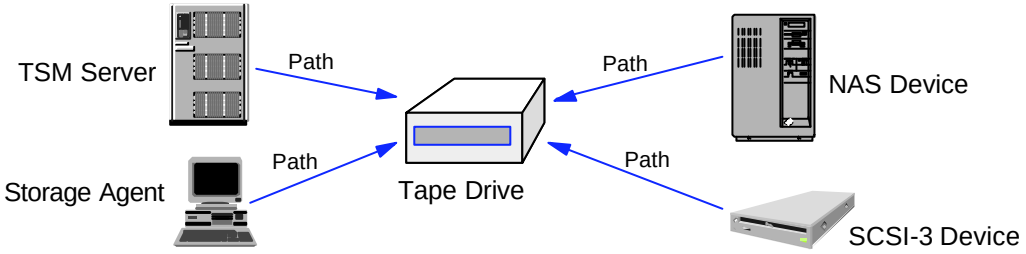

- A TSM path consists of
  - Source and target
  - Method by which source can access target
- Paths will replace
  - Device parameter on library and drive definitions
  - Drivemapping definitions for storage agents
- Allows sharing of the target device for improved resource utilization and scalability

Planned: Concepts

Tivoli

TDP for NDMP Topology

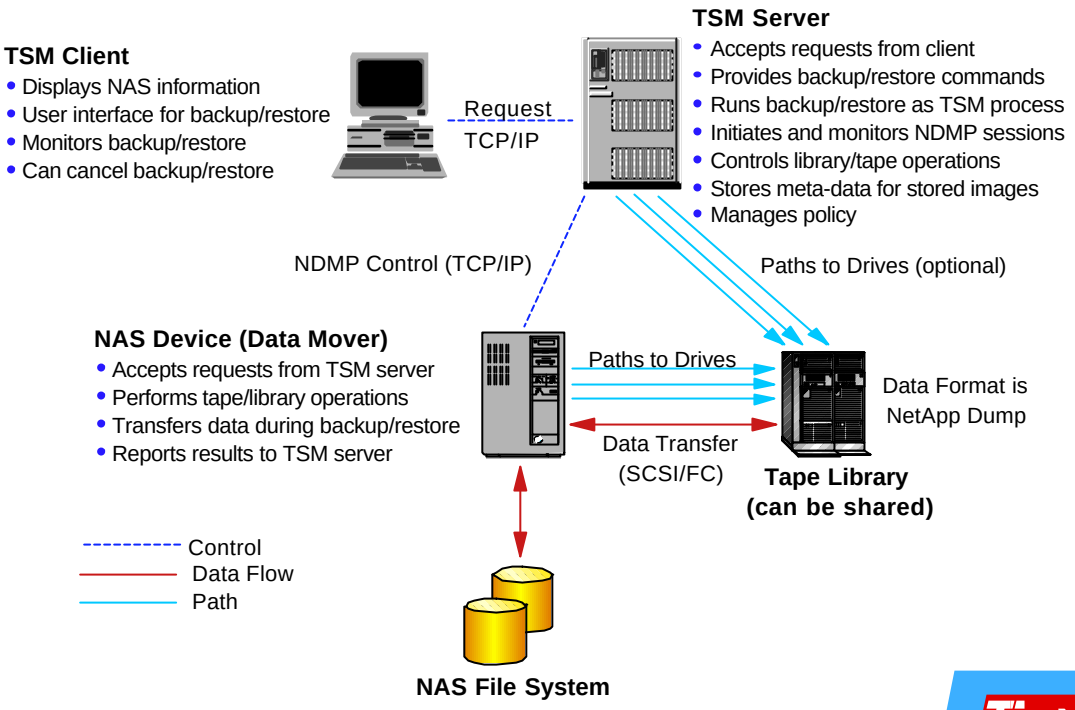

Planned: TDP for NDMP

Tivoli

TDP for NDMP Function

- NDMP-controlled backup of Network Appliance file servers with Data ONTAP 6.1.1 or higher
  - Full file-system image
  - Differential file-system image (files that have changed since last full backup)
- NDMP-controlled restore
  - Full file-system image
  - Full file-system image plus one differential file-system image
- Policy-based management of file-system images
- Data flow for backup/restore is LAN-free and outboard of TSM client and server
- Parallel backup/restore operations when multiple NAS file systems are processed

Planned: TDP for NDMP

Tivoli

TDP for NDMP Function (cont.)

- SCSI-attached libraries controlled via
  - Direct attachment to TSM server
  - Passing of SCSI commands through NAS device
- Sharing of tape drives
- Windows NT/2000 servers
- UNIX servers planned for 2002
- Choice of user interfaces for initiating, monitoring, and canceling backup and restore operations
  - Server console or administrative command-line client
  - Administrative web interface
  - Windows NT/2000, AIX, or 32-bit Sun Solaris client
  - Web client
- Scheduling of backup/restore operations using the administrative command scheduler

Planned: TDP for NDMP

Tivoli

SCSI-3 Extended Copy

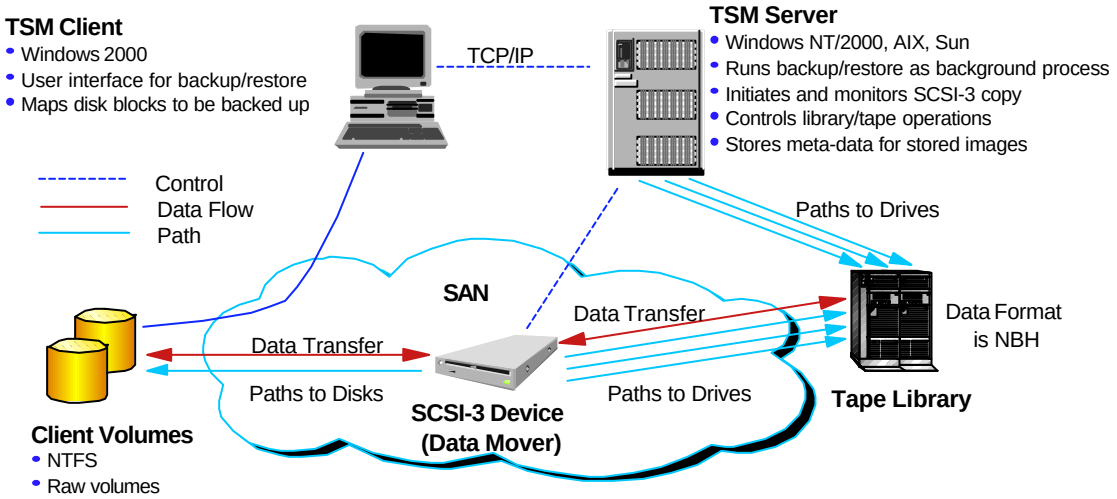

- Features and benefits are similar to those of Windows 2000 image backup
- Additionally, CPU cycles are moved from TSM client and server, which has potential to improve performance

Planned: SCSI-3

Tivoli

Potential Future Enhancements

Examples of items that have been considered for possible future implementation, but are not currently in plan

- Collocation enhancements
- Parallel migration
- Recovery log utilization
- Database reorganization
- Use of ESS/Timefinder for database backup
- TDP for NDMP extensions

Tivoli

Some Observations on Collocation

- Volume Capacity << Size of Collocation Unit
  - No or insufficient collocation
  - Scattering of data across many volumes
  - Long restores
- Volume Capacity ~ Size of Collocation Unit
  - Efficient tape utilization
  - Efficient restores
- Volume Capacity >> Size of Collocation Unit
  - Collocation too granular
  - Inefficient use of tapes and library slots, if tapes are dedicated to a node OR
  - Random mixing of nodes on each tape can lead to inefficient internal data-transfer operations to maintain collocation
    - Excessive mounting of target tapes
    - Multiple passes of input tapes
  - As tape capacities increase, this case may become more common

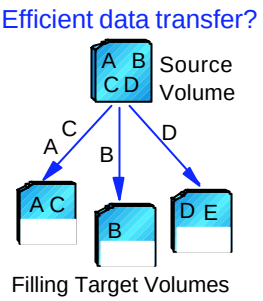

Tivoli

Potential: Collocation

Collocation Groups

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

```
define collocgroup a
update node ted collocgroup=a
update node sue collocgroup=a
update node mary collocgroup=a
```

Group A:
Node TED
Node SUE
Node MARY

Group B:
Node BOB
Node JOE
Node ANN

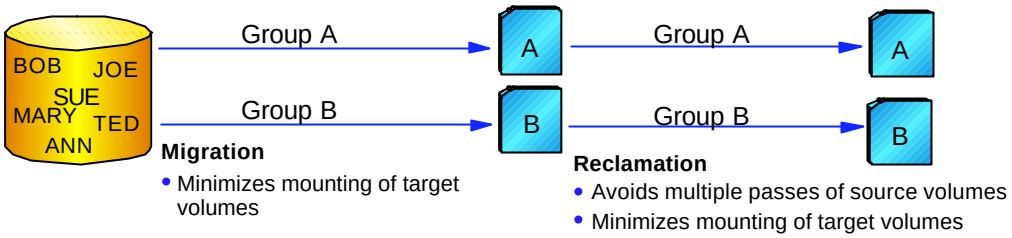

Migration
- Minimizes mounting of target volumes

Reclamation
- Avoids multiple passes of source volumes
- Minimizes mounting of target volumes

- For move/copy operations to sequential media, data transfer performed by group
- Avoids multiple passes of source volumes and minimizes volume mounts
- Reduces number of tapes required for effective collocation
- Increases feasibility of collocated copy pools for offsite storage

Potential: Collocation

Tivoli

Parallel Migration

Today: Each process migrates data for a different set of nodes

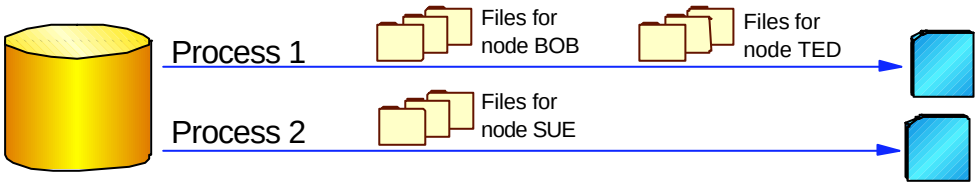

Enhancement: Allow multiple processes to migrate data for the same node

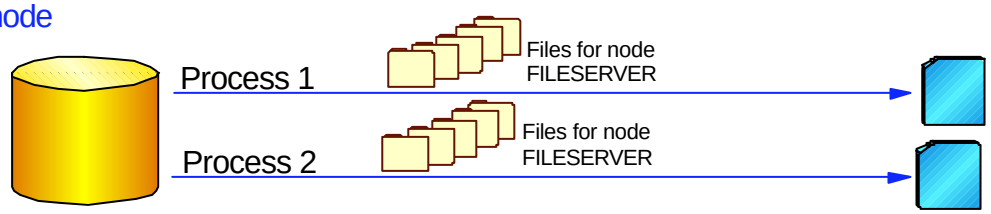

- Allows parallel processing of data for very large nodes
- Ensures participation by all processes until migration is complete

Potential: Parallel Migration

Tivoli

Recovery Log Utilization

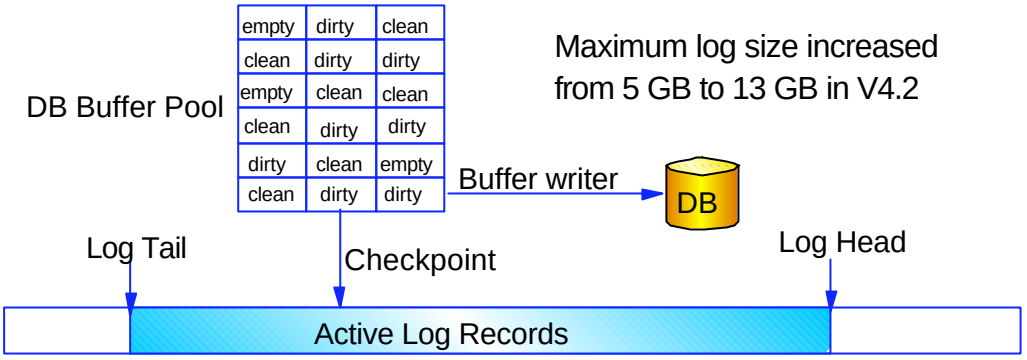

Potential future enhancements

- Provide tool to diagnose the reason for pinned recovery log
- Reduce log space occupied by checkpoint records
  - Reduce checkpoint frequency
  - Avoid accumulation of checkpoint records in log
  - Improve performance of buffer writer

Database Reorganization

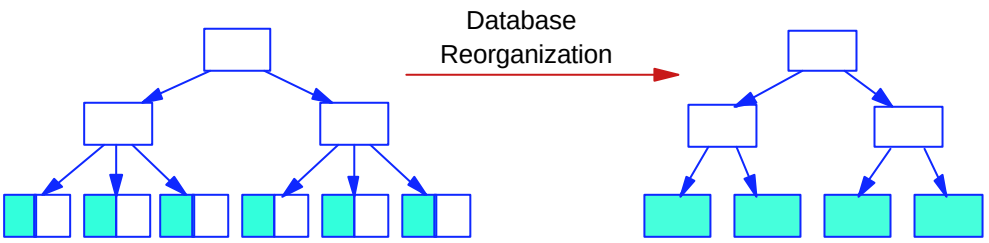

- Benefits of database reorganization using current unload/load
  - Reduces database space utilization
  - Minimizes fragmented pages
  - Improves performance for database scans
- Potential enhancements
  - Tool to determine when database reorganization is needed
  - Online, non-disruptive database reorganization

Use of ESS/Timefinder for DB Backup

IBM Enterprise Storage Server

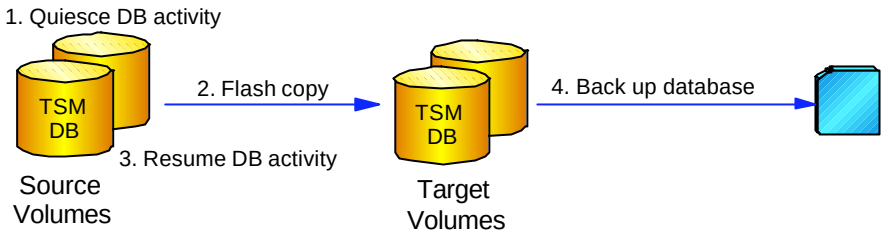

EMC Timefinder

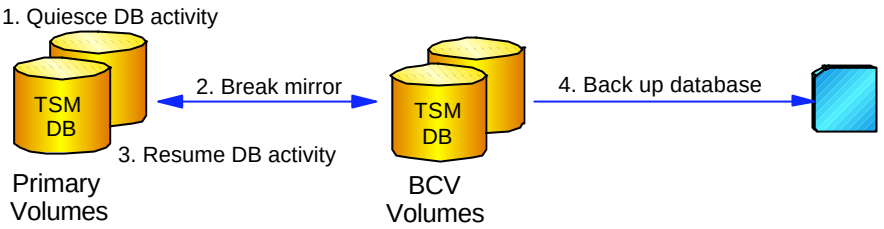

Database backup with minimal impact to TSM server

Potential: DB Backup

Tivoli

TDP for NDMP: Single-File Restore

- Restore of individual files or directories from file-system image
- Basic restore
  - Administrator specifies file name and directory
  - File server scans backup to locate specified file
- Direct-access restore
  - TSM collects and stores file information during backup of file-system image
  - Client GUI displays image contents and provides interface for specification of file(s) to be restored
  - File server positions to and restores selected files

Potential: TDP for NDMP

Tivoli

TDP for NDMP: 3-Way Configuration

- NAS device 1 does not need access to tape drives
- Exploits tape libraries with integrated NDMP-server support
- Data flow over the LAN

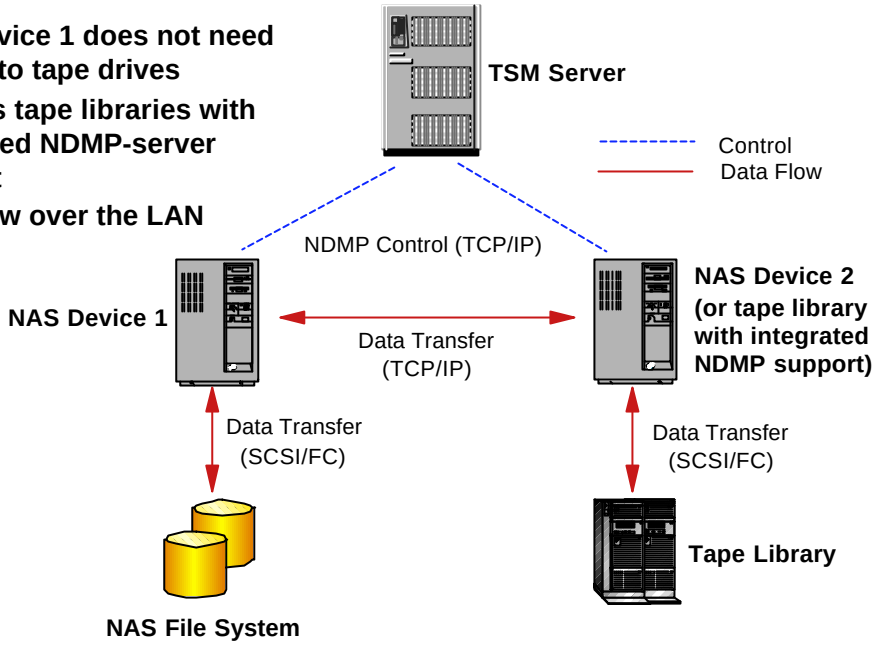

Potential: TDP for NDMP

Tivoli

TDP for NDMP: Filer to Server

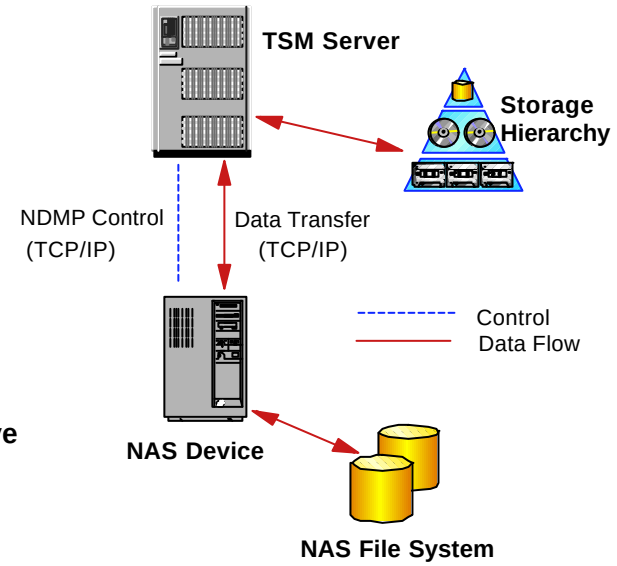

- NAS device does not need access to tape drives
- NAS data stored in TSM's storage hierarchy with native data format
- Data flow over the LAN

Potential: TDP for NDMP

Tivoli

Other TDP for NDMP Enhancements

- Additional NAS vendors
  - IBM
  - EMC Celerra
  - EMC CLARiiON IP4700
  - Auspex
- Additional tape libraries (V4.2.1 supports SCSI only)
  - 3494
  - ACSLS
- Backup of NAS images to copy storage pools (V4.2.1 supports duplication by backing up to multiple primary pools under different node names)

Potential: TDP for NDMP

Tivoli

Summary

- Recent releases of Tivoli Storage Manager have improved scalability through enhancements to
  - Performance
  - Manageability
  - Resource sharing
- Future product releases will further improve scalability by offering new/improved functions and by exploiting emerging technologies

Tivoli