IBM System Storage

Storage Update

Leif Schiøler
System Storage Nordic
+45 2880 3405
Leif_schioeler@dk.ibm.com

Agenda

- DS8000 new news
- Data mobility solutions
- TS7700 (system z virtual tape) news
- TS7520 (open systems virtual tape)
- Management of enterprise tape resources (iRMM)

DS8000 new functions

IBM System Storage DS8000 Turbo – Powerful Innovation – R3

- Performance innovation that builds on DS8000 world class performance
  - **Storage pool striping (rotate extents)**
    - •Maximizes performance without special tuning
  - **AMP (adaptive multistream pre-fetching)**,
    - •Breakthrough caching technology can dramatically improve sequential performance to reduce backup times, processing for BI/DW, streaming media, batch
  - **IBM z/OS Global Mirror Multiple Reader**
    - •provides innovation for IBM System z to improve throughput for z/OS remote mirroring (XRC)

New!!

- Optimizations to simplify and increase efficiency
  - **IBM FlashCopy SE**
    - •can lower costs by significantly reducing disk capacity needed for copies. Less capacity means fewer drives, less power, more GREEN
  - **Dynamic Volume Expansion**
    - •simplifies management by enabling easier, online, volume expansion to support application data growth.
  - **IBM System Storage Productivity Center**
    - •single pane control and management – Integrates the power of TPC and the IBM DS Storage user interfaces into a single view

Enabled through Release 3 of DS8000 microcode

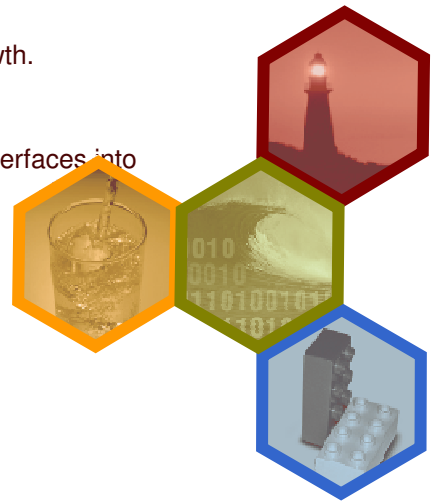

Planned Announce October 23, GA December 7

DS8000 Storage Hierarchy (background)

- **Disk**
  - Individual DDMs
- **Array Sites**
  - Logical Grouping of 8 DDMs (DS8000) or 4 DDMs (DS6000) of same speed and capacity
- **Arrays**
  - DS8000 -- one Array Site used to construct one Array (8 DDMs)
  - DS6000 – one or two 4-DDM array sites used to construct one Array (4 or 8 DDMs)
- **Ranks**
  - One Array in a Rank (DS8000--8 DDMs, DS6000—4 or 8 DDMs)
  - Rank is divided into N fixed sized Extents
- **Extent Pools**
  - 1-N Ranks form an Extent Pool
  - All Extents in a Pool are same storage type (CKD/FB); same RAID recommended
  - Associated with server0 or server1

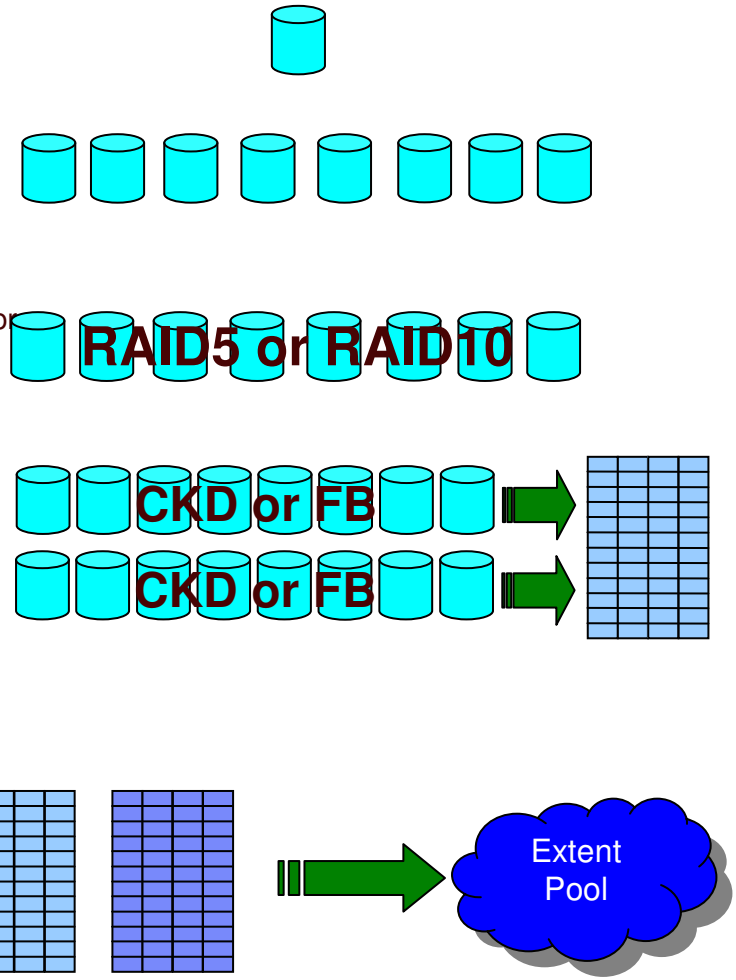

DS8000 Storage Hierarchy (more background)

- **Extent Pools**
  - –1GB (FB) or Mod1 (CKD) extents
  - –All Extents in a Pool are same storage level/type
  - –Associated with server0 or server1

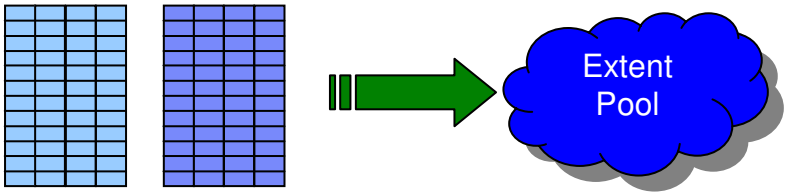

- **Volumes or LUNs**
  - –Made up of extents from 1 extent pool
  - –Associated with LSS
    - •LSS selection limited based on Extent pool server affinity
  - –Assigned to Volume Group

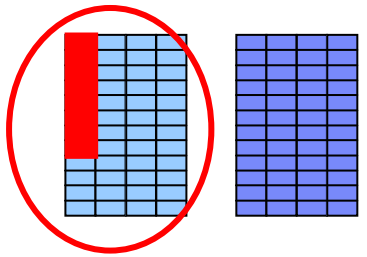

- **Volume Groups**
  - –FB LUN masking
  - –Maps LUNs to server ports or port groups
  - –1 Host port can be member of 1 volume group
  - –1 Volume can be member of >1 volume group
  - – >1 host (even different server types) can be member of 1 volume group

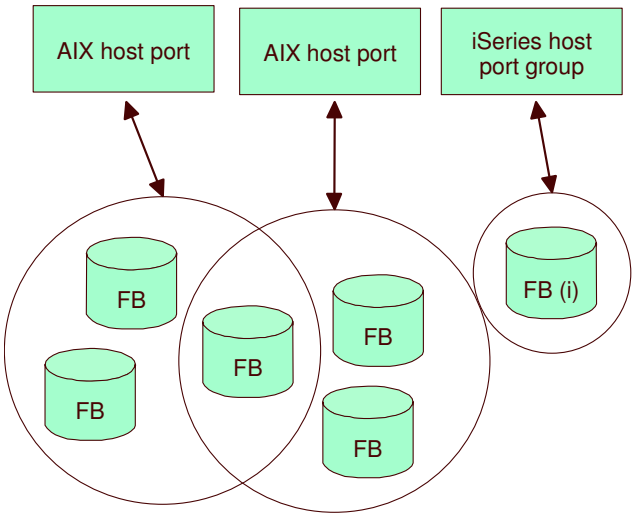

Volume Allocation Algorithms

Original method (”fill and spill”)

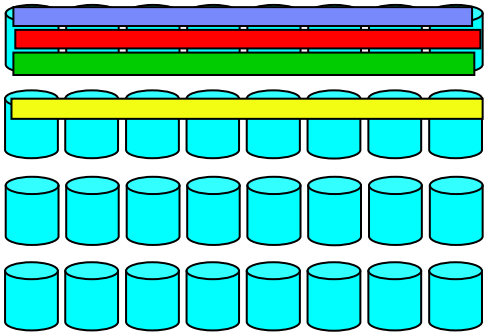

Current method (until now)
”volume round robin”

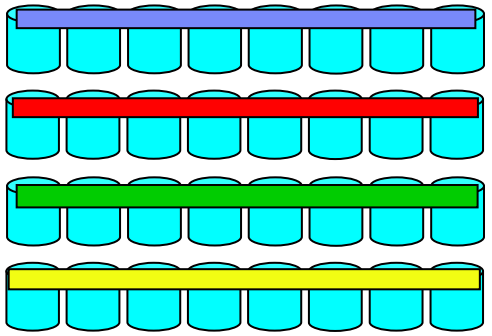

New option: extent round robin

New: Volume Spread Across RAID arrays

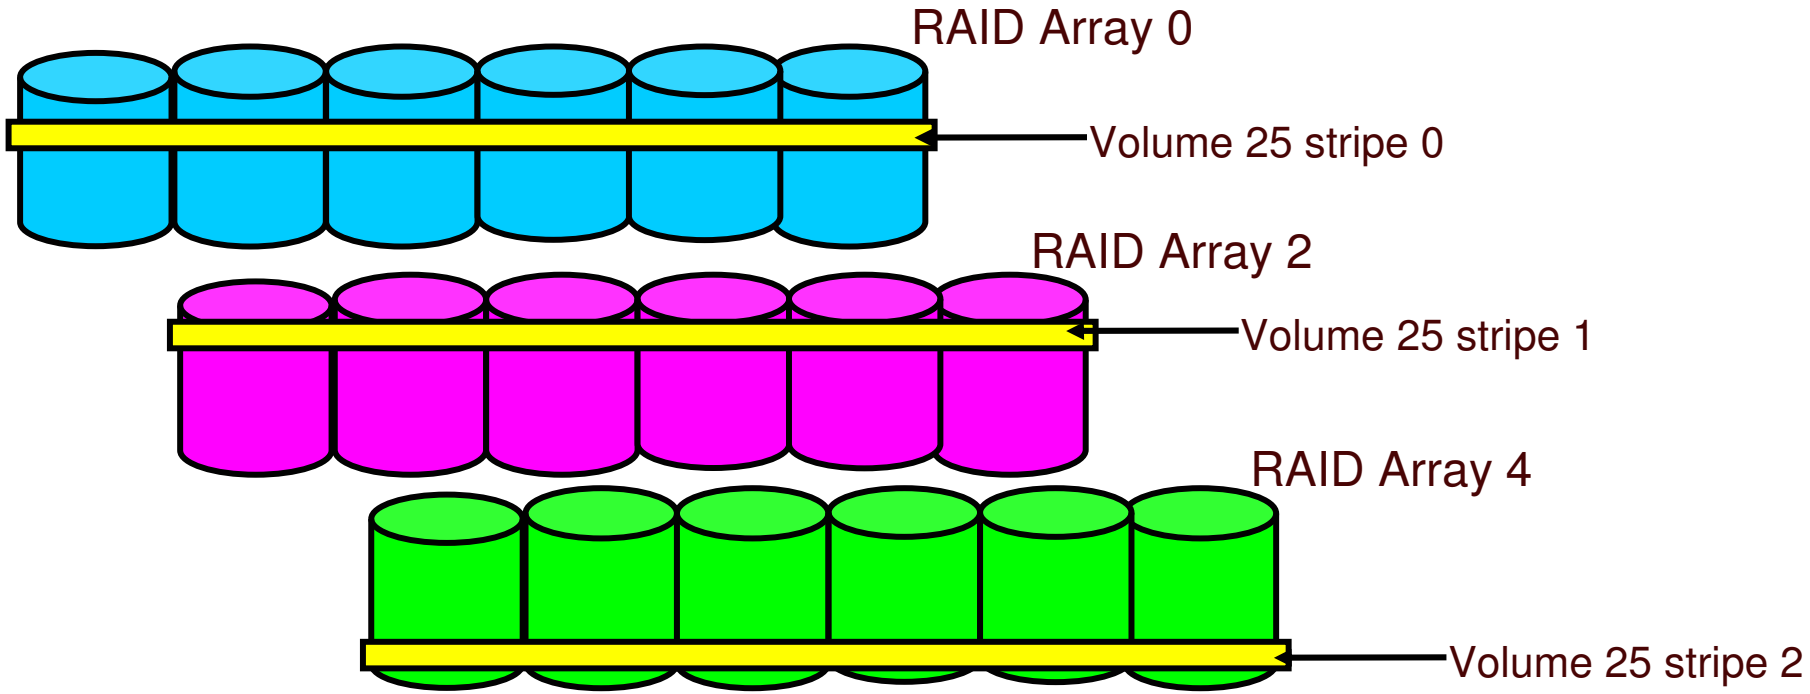

- Extents striped across multiple RAID arrays
- Improves volume throughput dramatically for some workloads
- Largely eliminates hot spots

Storage Pool Striping Details

- Existing volumes can't be converted to striped or expanded with stripes
- Arrays are filled round-robin
- Allocation method is chosen at volume creation
- An Extent Pool can contain striped as well as traditional volumes
- Recommendations
  - "Sweet spot": 4 arrays per Extent Pool
  - Separate 6+P and 7+P arrays
- Storage Pool Striping and Multi-Rank Extent Pools will become the recommended way to configure a DS8000 in general

Where better not to use Storage Pool Striping – or be at least very careful

- **Exceptions**
  - Avoid “over virtualization” (i.e. SAN Volume Controller, LVM Striping, System i, DB2 database containers, SMS Striped Datasets, etc.)
  - Configurations where it is desirable to control volume placement precisely
  - Benchmarks with artificially balanced volume activity
  - If RMF shall be used and the device addresses replace rank positions
  - MUST-ISOLATE Workloads – Customer must decide!
    - Different systems
    - Different LPARs
    - Different applications

Logical Volume Expansion

- Expand Volume in Place
  - No need to take volume off-line
  - Some OS need additional actions to see the extra space
- No Copy Services
  - All relationships must be removed prior to expansion
  - Relationships can be recreated after initialization completed
- The original allocation method is retained
- No shrinking!

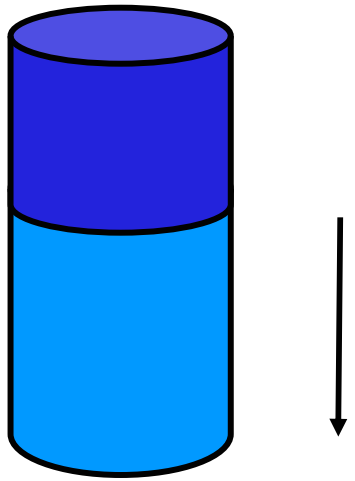

Space Efficient Volumes

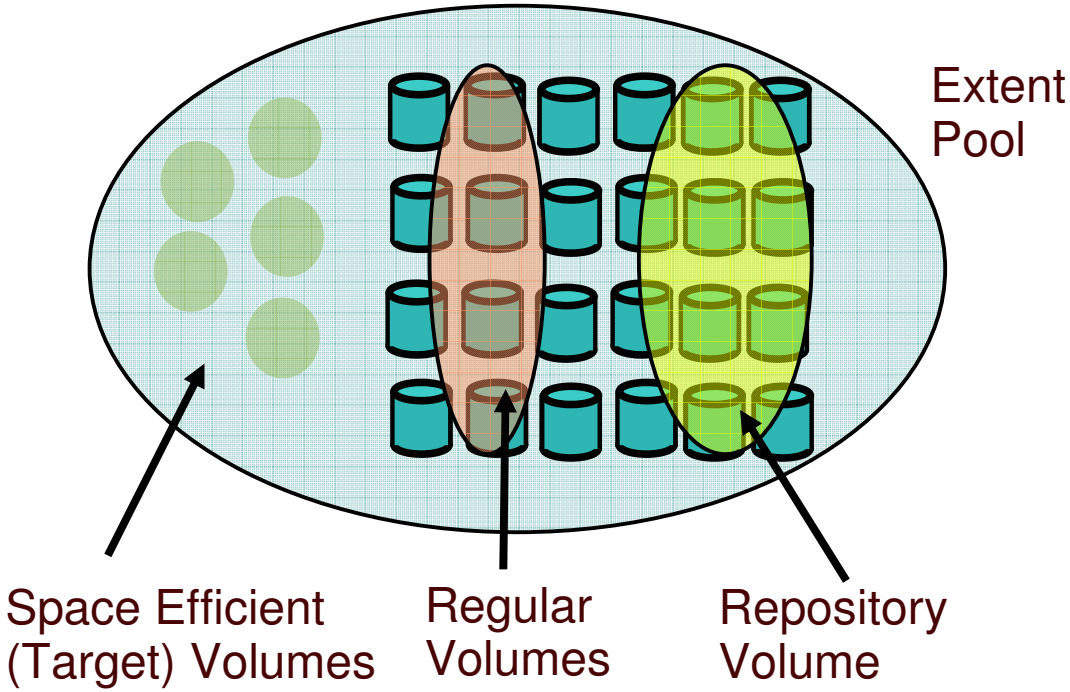

- Space Efficient Volumes have no Extents allocated
- For updates 64K increments are allocated in the Repository

Space Efficient FlashCopy

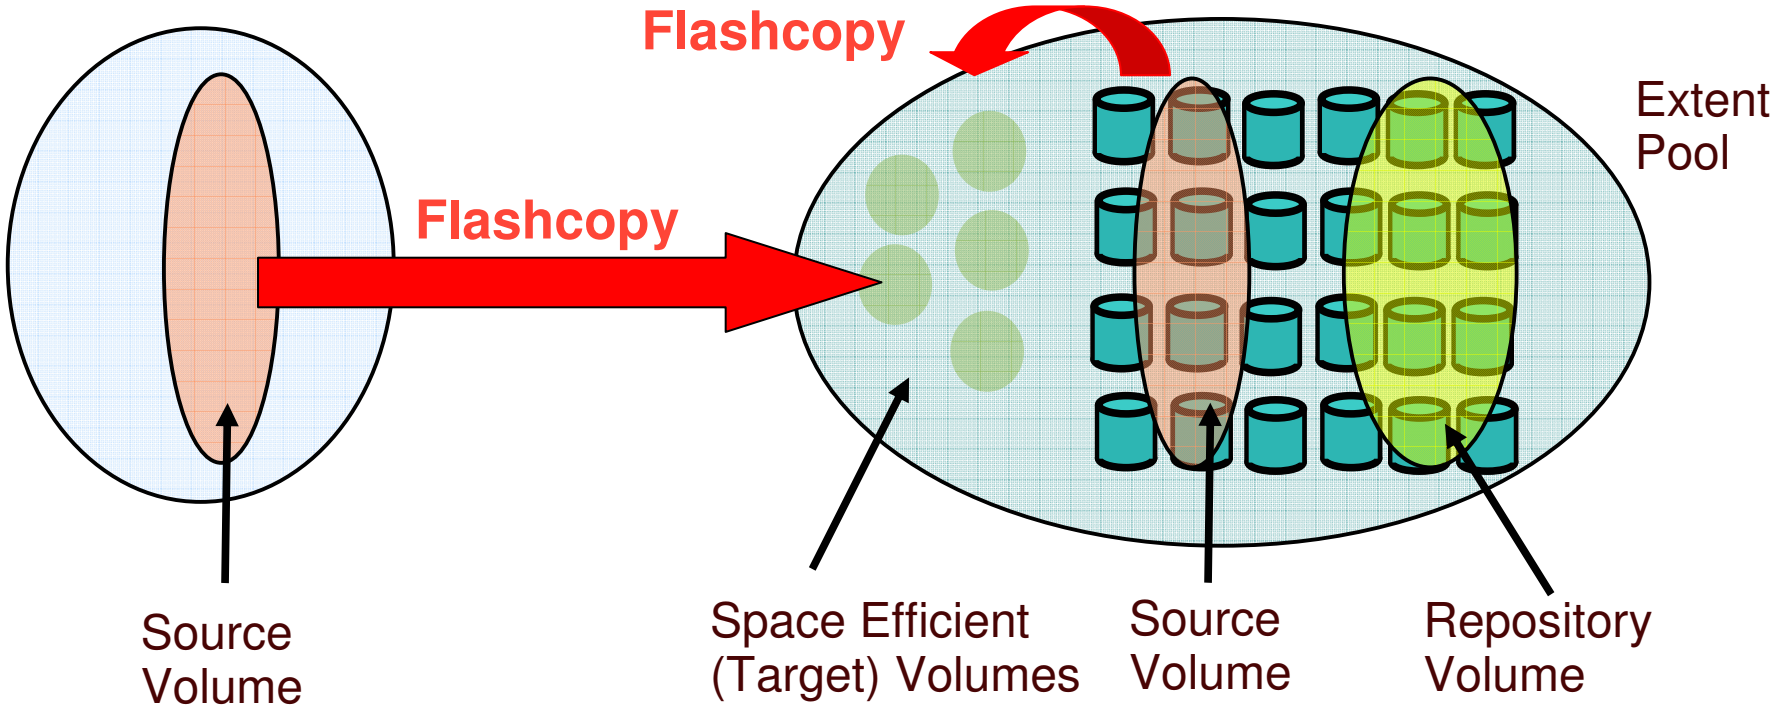

- FlashCopy relation from Source to SE Target
- No background copy allowed

Repository Volume

- One Repository Volume per Extent Pool is possible
  - No dynamic expansion in first release
  - Striped in multi-rank Extent Pools (recommended)
  - Minimum size: 16 GB
  - Maximum size: 80 TB (CKD), 100TB (FB)
  - Overflow warning threshold can be set for each repository
  - Commands available showing repository consumption of space efficient volumes
- Disk Magic Tools for Repository Sizing Guidance

Repository Out of Space

- Relationship will “fail”: Target unavailable
- In Global Mirror relations:
  - Write inhibit for the source
  - Global Copy will suspend
- Space is released when relationship is withdrawn
- Monitor utilization via DSCLI or GUI
- Errors and Warnings
  - Reported to z/OS console and via SNMP trap
  - Warning threshold can be set for each repository

Recommended Usage

- **Short Term FlashCopy relationships**
  - Tape Backup
  - Online backup (limited today)
  - (Volume C in GM relationship)
  - Additional test copy in GM relationship (“D” copy)
- **Optimized for 5% source change rate**
- **Remove relationship immediately after use**
  - Free up space in repository
  - Avoid unnecessary performance impact

Global Mirror Copy Technology

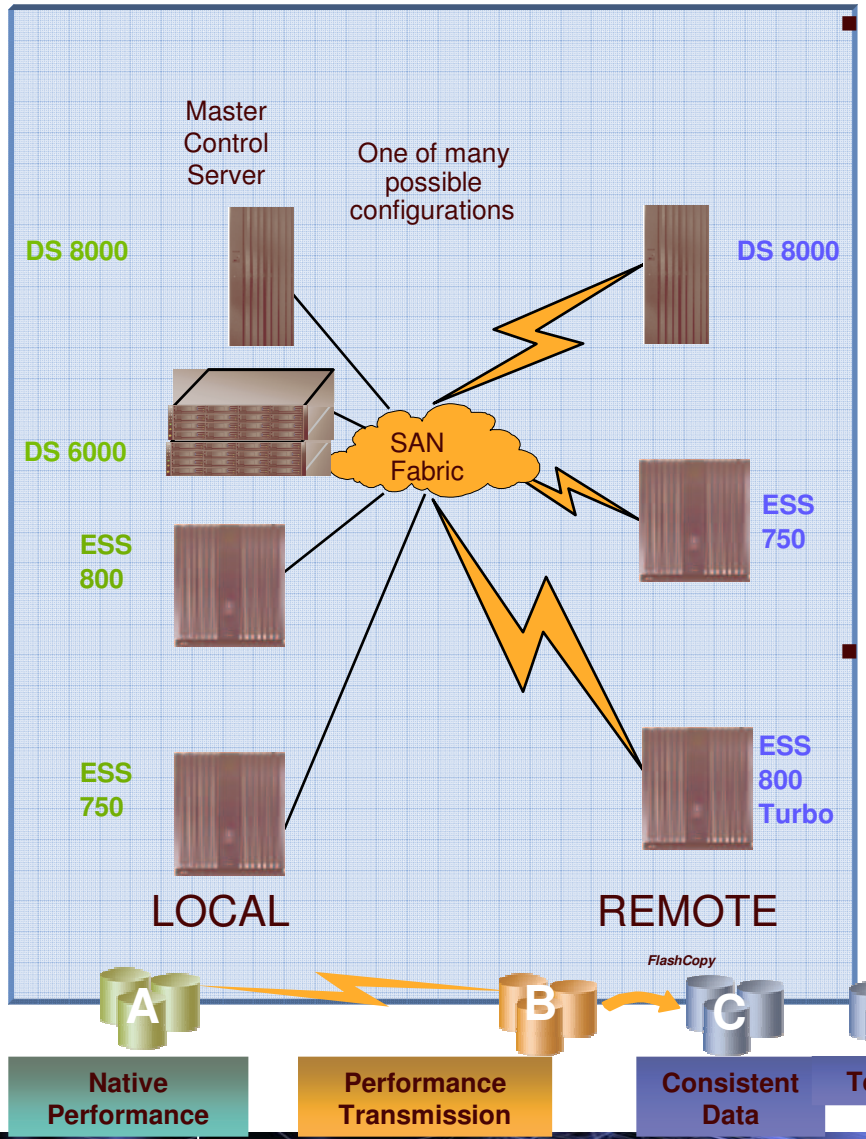

Characteristics:

- **Global Distance:** Two-site, unlimited distance, data consistency
- **Heterogeneous:** System z and/or Open systems data
- **Scalable:** Up to 15 total ESSs in Global Mirror session (with RPQ)
- **Flexible:** Many possible configurations
- **Application Performance:** No impact
- **Mirroring Performance:** Two Fibre Channel links per disk subsystem pair sufficient for most workloads

Intended Benefits

- **Autonomic:** No active external controlling software required to form consistency groups
- **Saves cost:** No server cycles required to manage data consistency
- **Lowers TCO:** designed to provide improved performance at lower cost point

Element management today

Customers find a great deal of value in managing IBM storage with TPC, but initial deployment can be challenging.

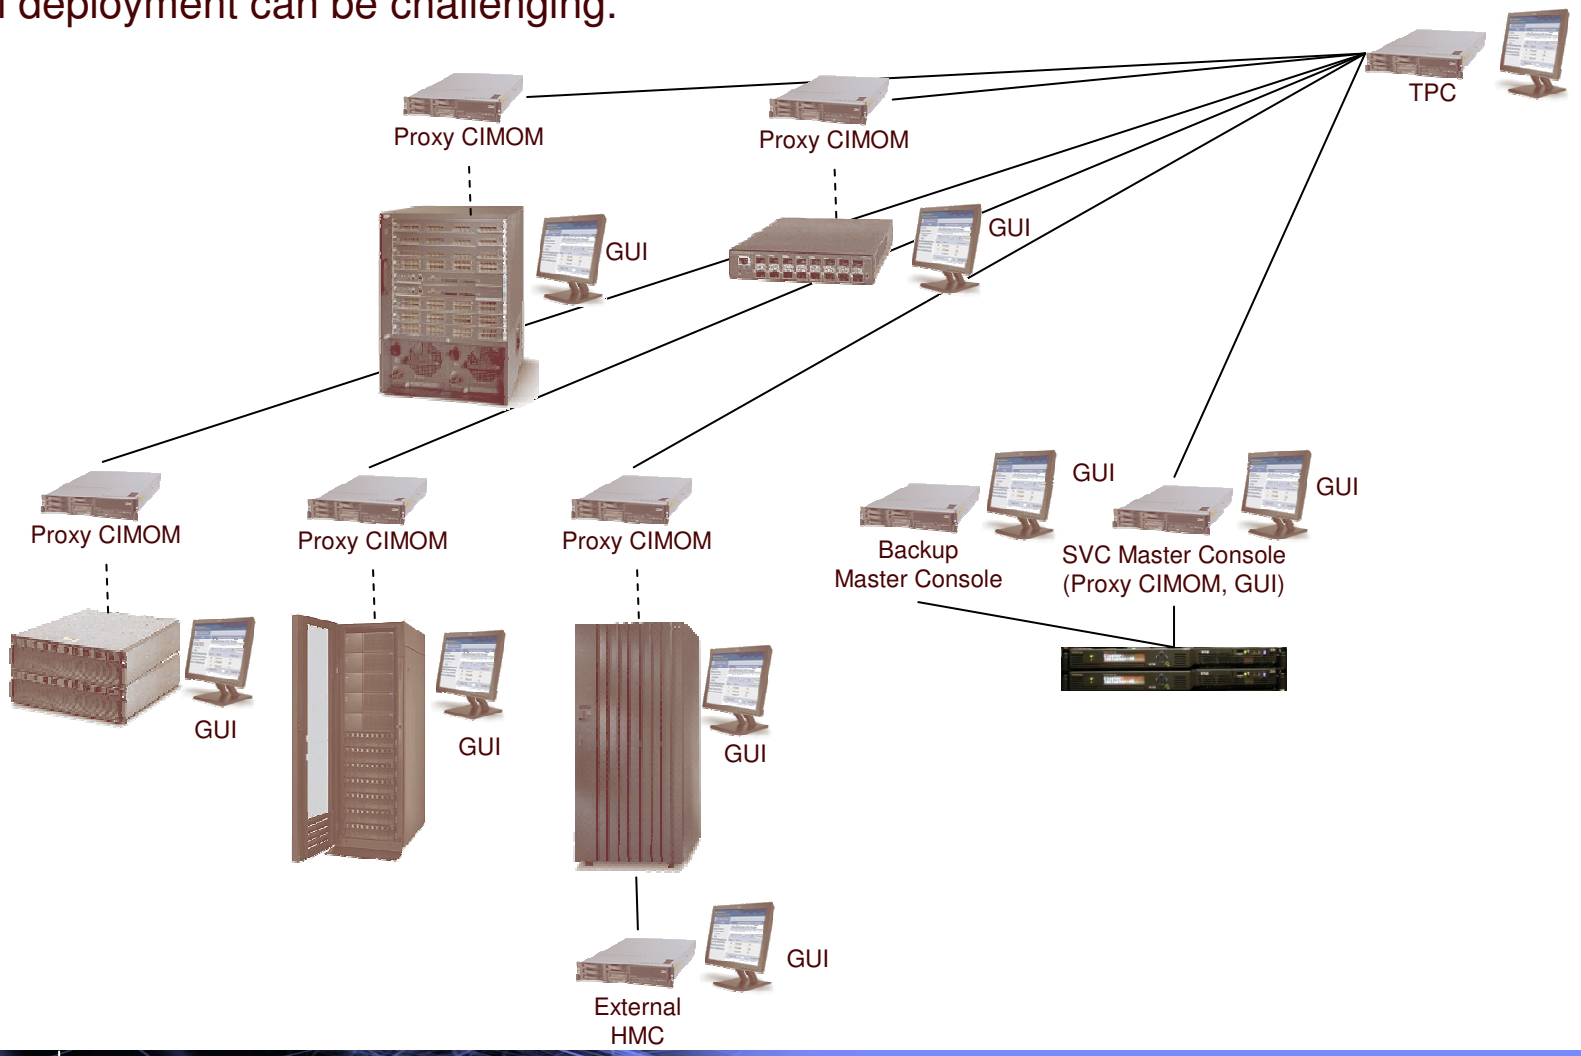

Introducing the SSPC Concept

CIMOMs become imbedded in devices, user interfaces and management functions become part of SSPC, further reducing server counts

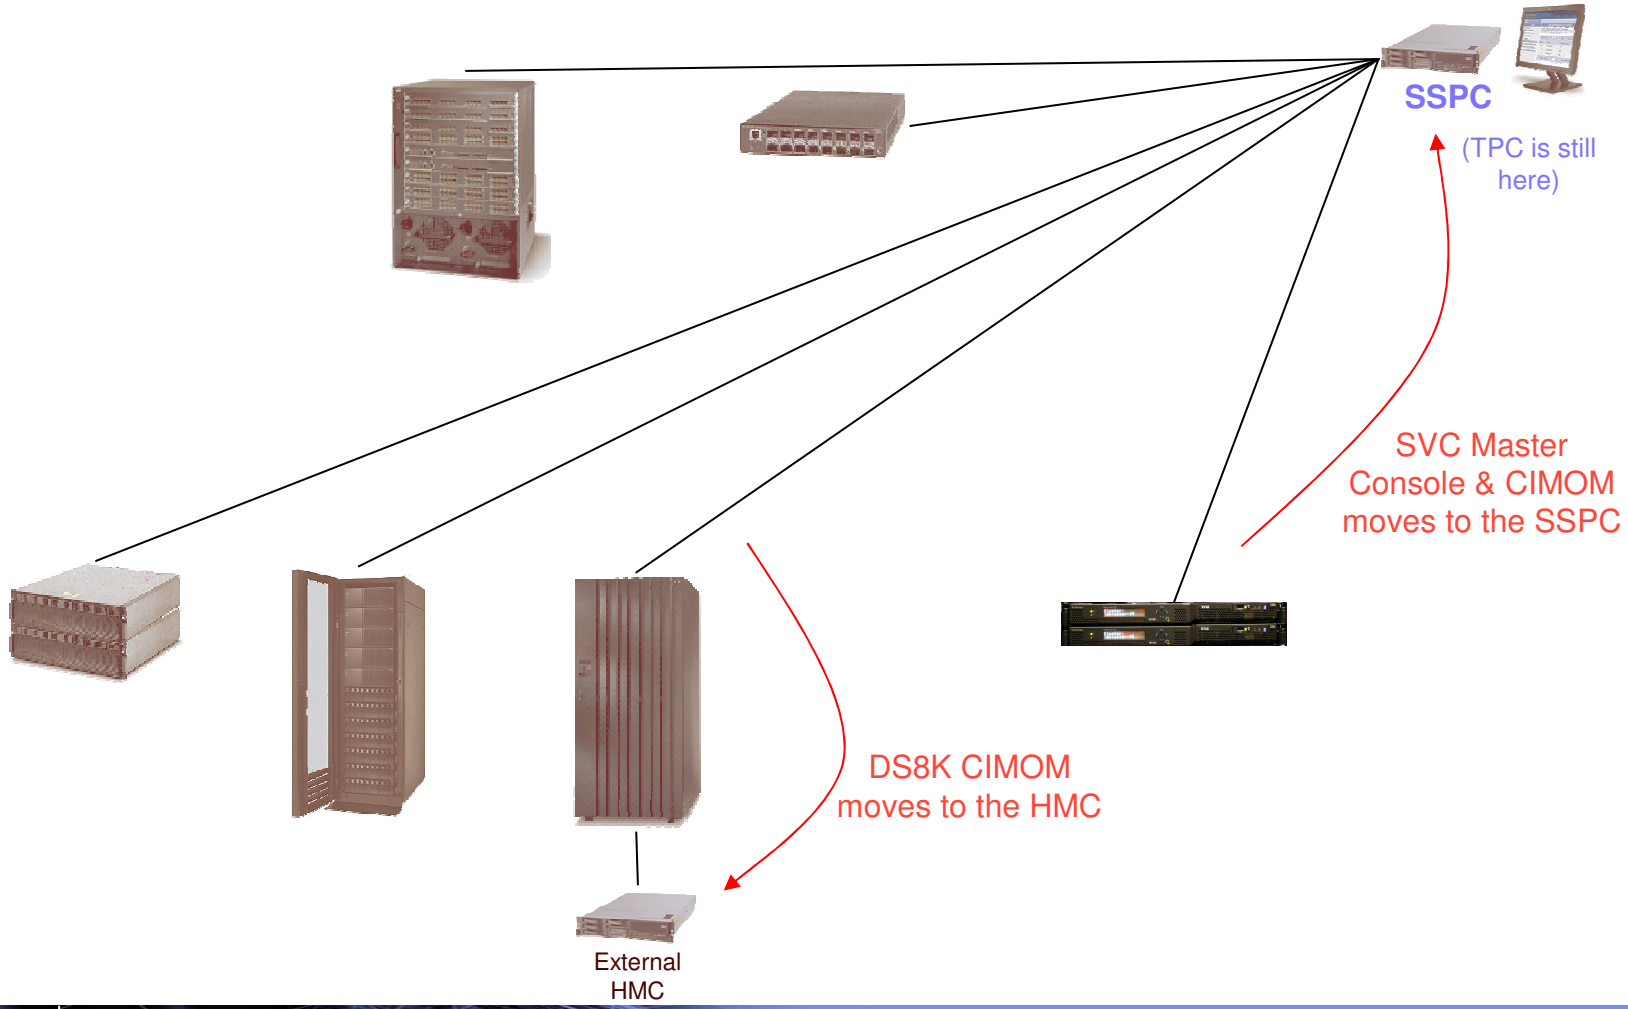

Yesterday

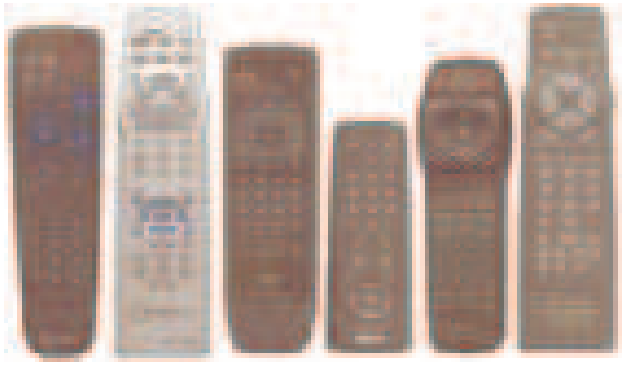

Tomorrow

Wouldn't it be nice to have a single master control that can manage everything:

Architecture Overview

- Sold as a Solution
- Central User Interface
  - Configuration
  - Operations
  - Health Console

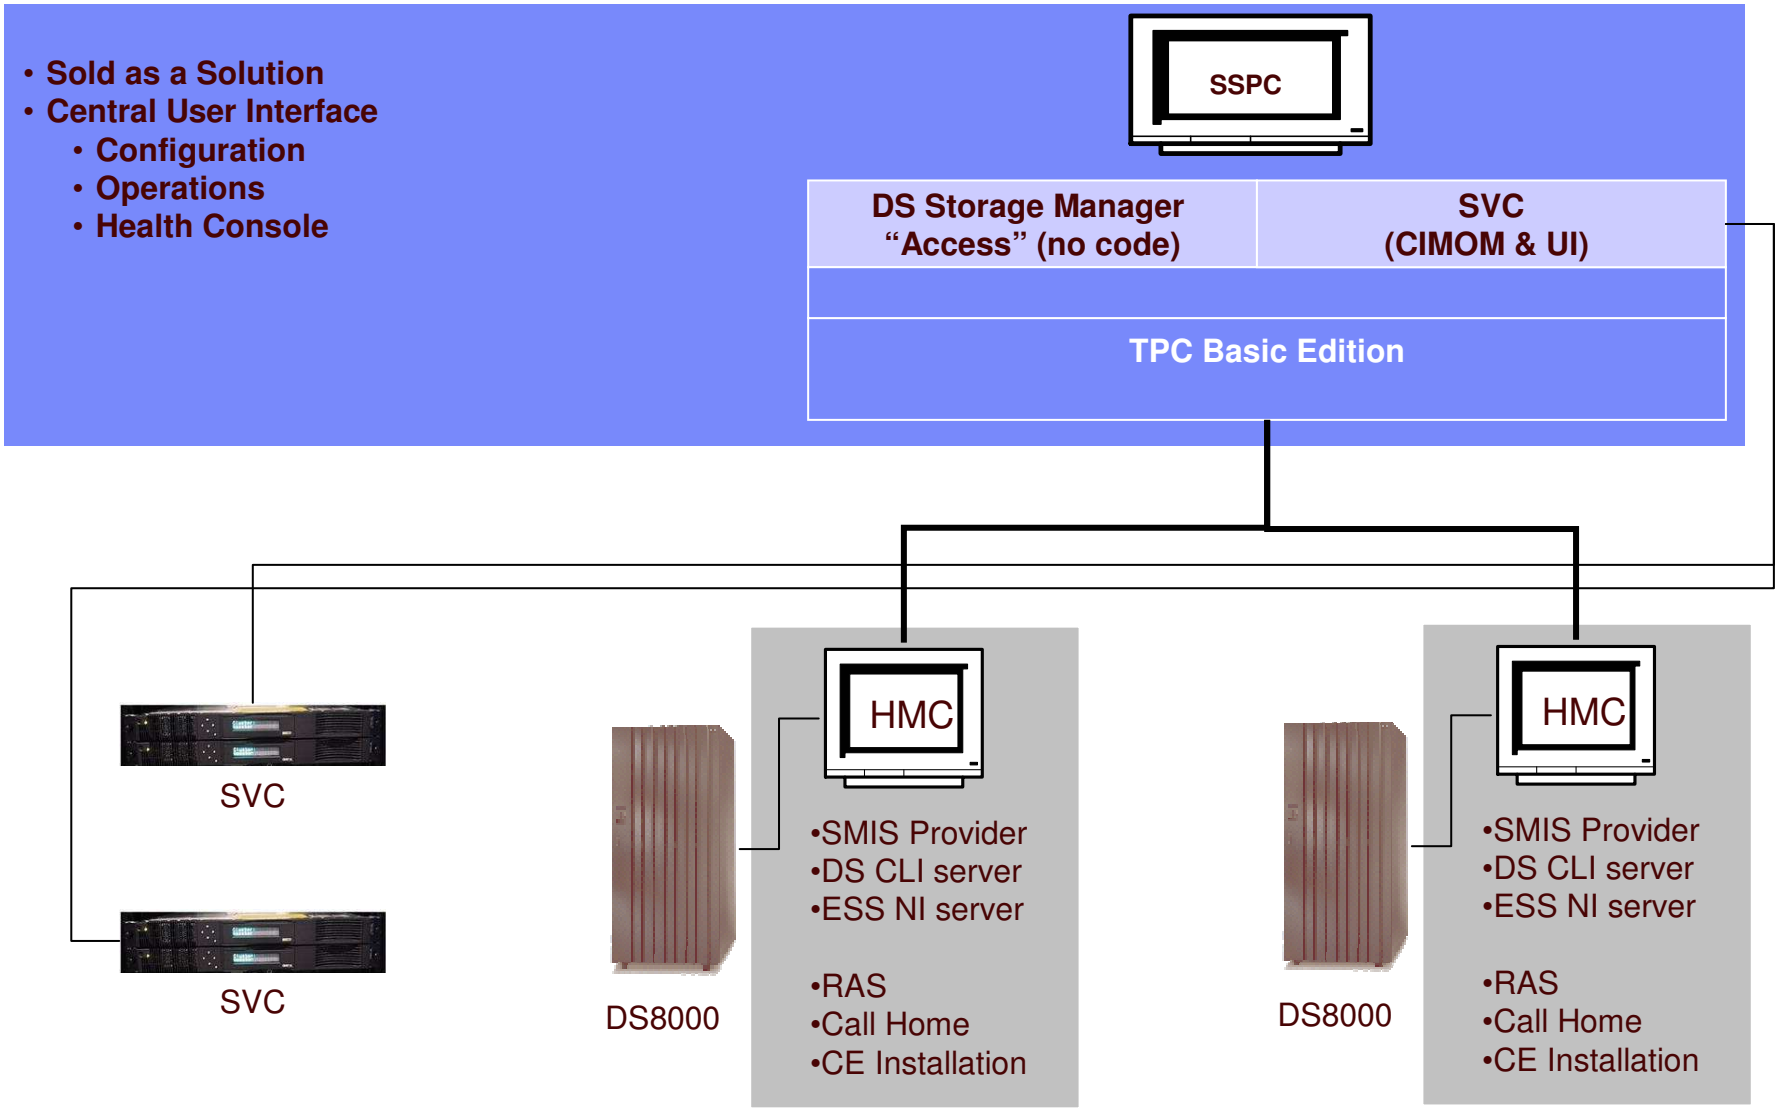

What value does the SSPC bring?

- It centralizes the disk element managers to one location and removes the need to install them on separate servers
- It provides a simple entry point to introduce storage management software (TPC) that is initially targeted at managing disk
- The TPC for Disk product can be added to enable performance management of disk, if desired
- TPC is designed with scalability in mind and can grow beyond disk to become an encompassing solution when the management of fabric and data is added
- It offers a simple migration path and hardware platform for adding higher value storage management applications like Performance Management (TPC-SE) and Replication Management (TPC-R)

Data mobility solutions

Enterprise Data Migration – Different Shapes and Sizes

Technology Refresh

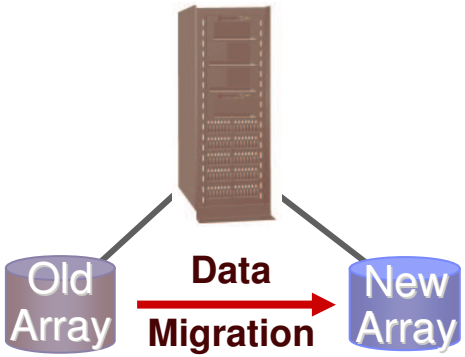

Server/Storage Consolidation

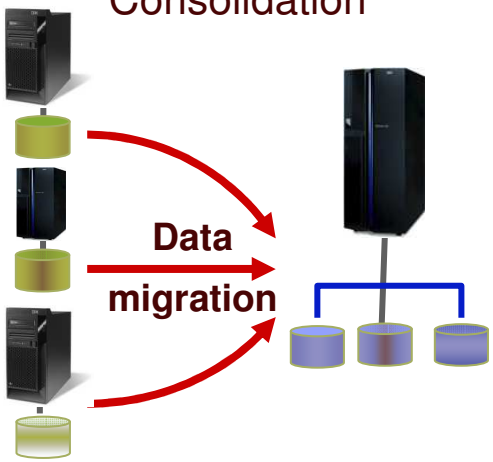

Data Center Relocation/Consolidation

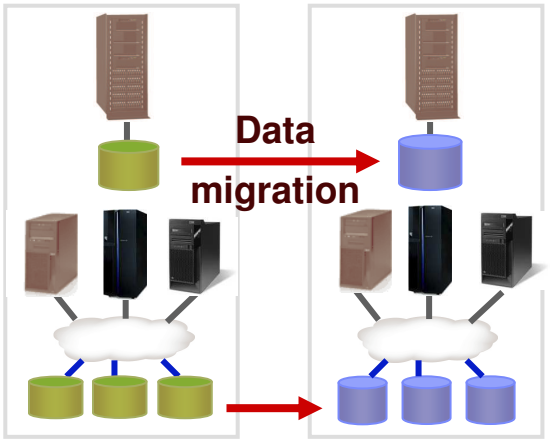

Facilitate Tiered Storage

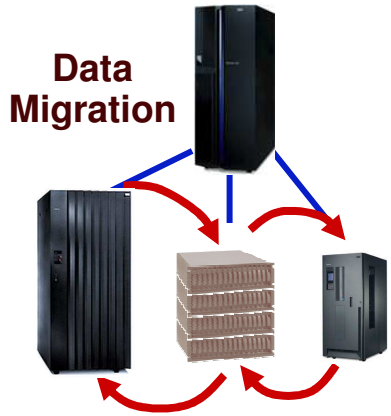

Direct Attached Storage (DAS) to Storage Area Network (SAN)

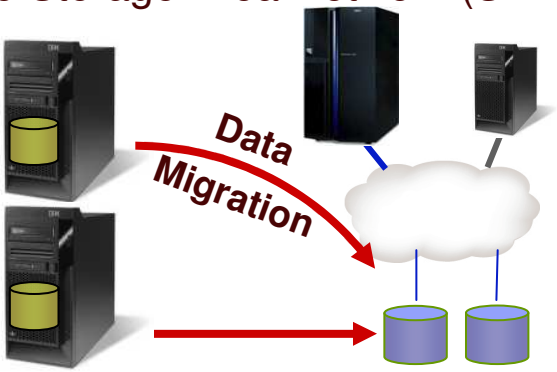

Improve Performance

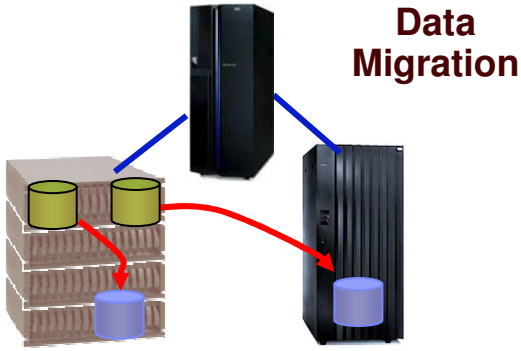

Clients must address three major issues when managing and moving data

- Magnification of daily challenges
  - Business cannot be disrupted
  - Pressure is on planned downtime
- Migration is difficult and complex
  - It can require highly specialized skills
  - It often causes unpredictable results
  - It involves disparate tools and processes
- Proprietary technology can reduce chances for success
  - Organizations need to move data with standard, responsive and reliable processes
  - Clients must contend with limited platform and multivendor choices

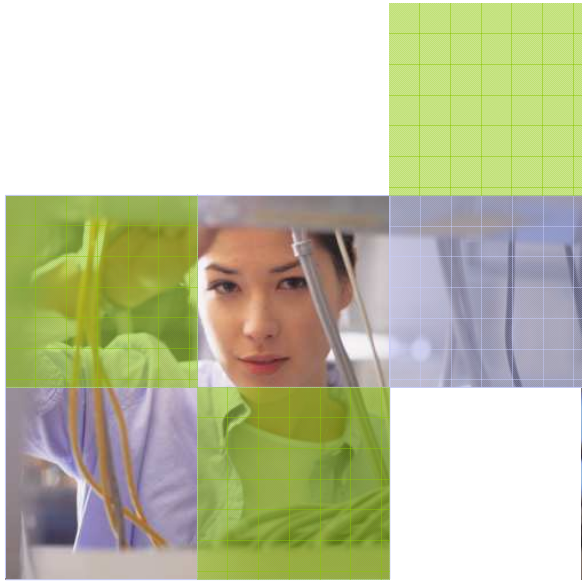

Nonstop Mobility™ Offerings

- A simple unified solution across the enterprise

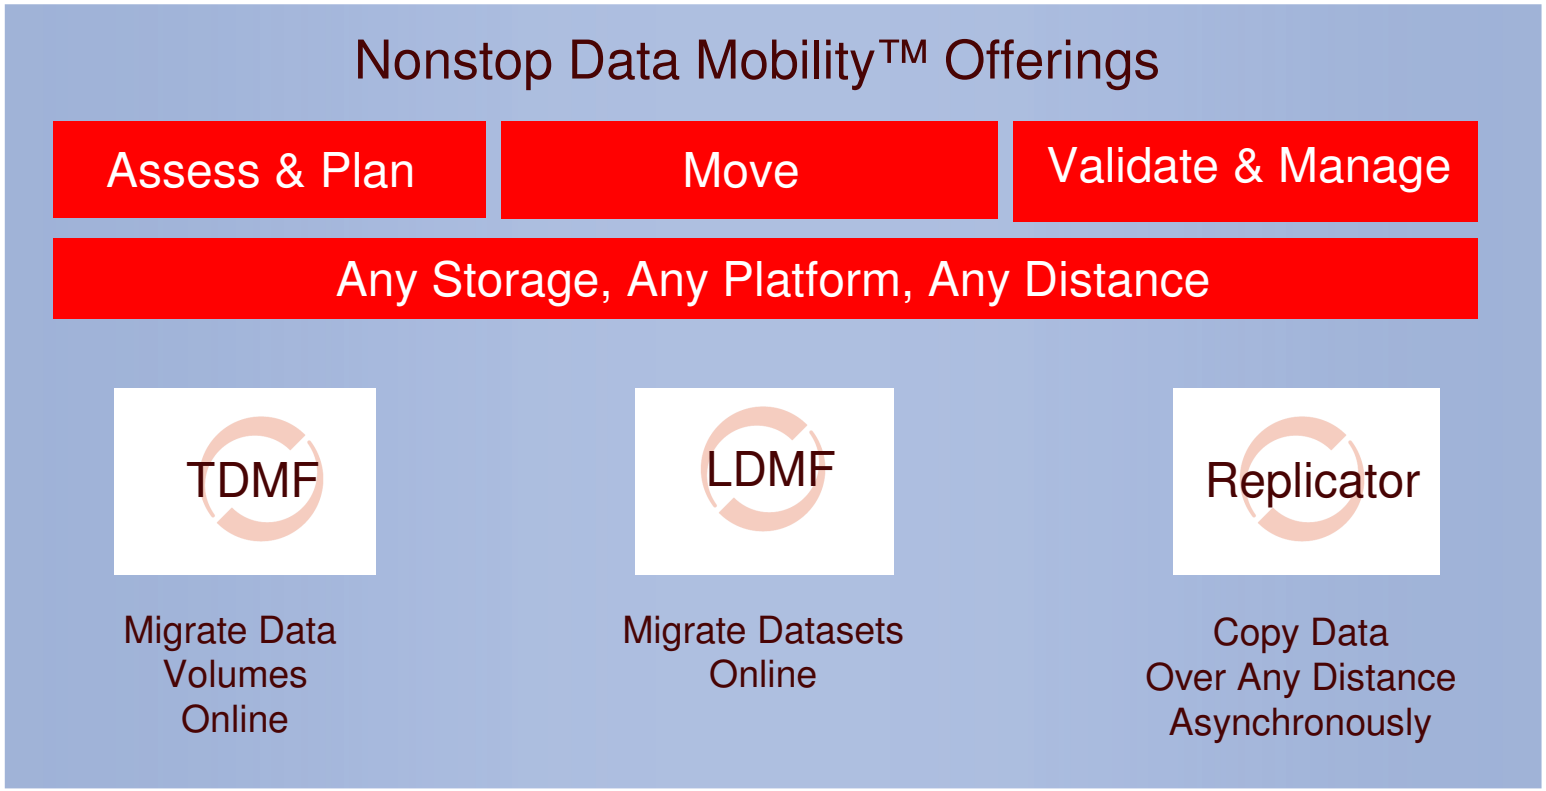

Note: All IBM/Softek products are host-based software solutions.

IBM and Softek – Ultimate Choice for Data Mobility

Any platform, any vendor, any distance as an asset or pre-integrated services offering

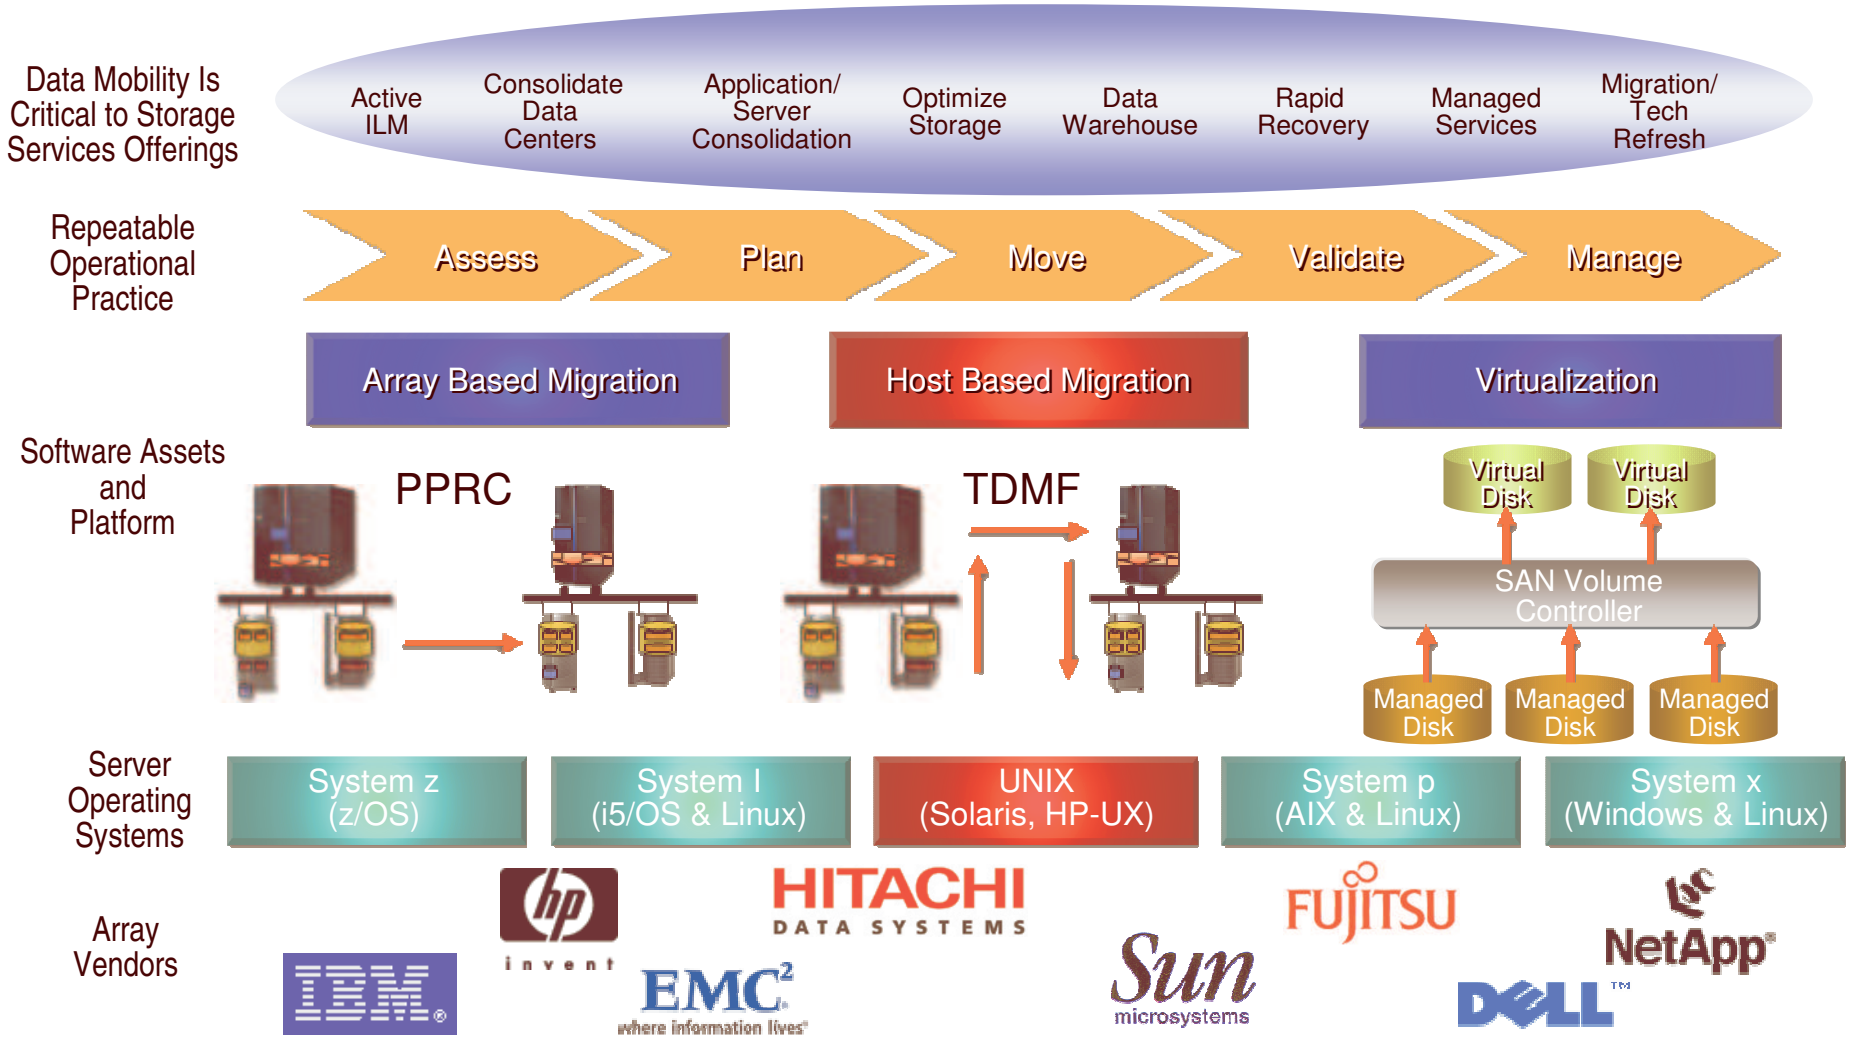

TDMF Product Family

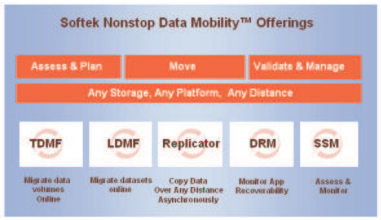

TDMF

Migrate Data
Volumes Online

Any Storage, Any Platform, Any Distance

z/OS	Unix / Linux	Windows
TDMF z/OS	TDMF Unix TDMF Unix/Linux IP	TDMF Windows IP

TDMF z/OS – Gold Standard for Mainframe Digital Migration

- Patented, non-disruptive “dynamic swap” migration
  - Transparent
  - Divert I/O to new storage
- Local or global migration Solutions
- Fast, easy, non-disruptive, host-based install
- Dynamic configuration
- Dynamic pacing
- Replication features:
  - PIT, PPIT, TCP/IP, OVA

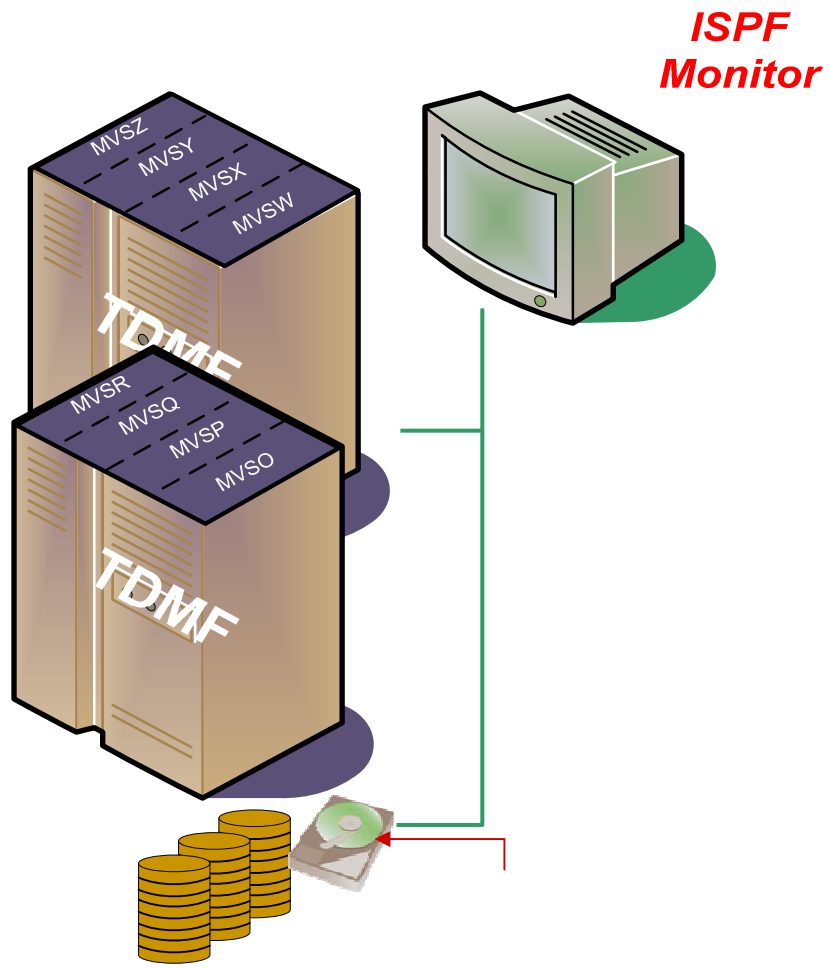

TDMF - the Phases

- Initialization
- Activation
- Copy
- Refresh
- Quiesce
- Synchronization
- Redirect
- Back Out
- Resume
- Terminate

Migration phases

1. Initialization

- Master and agent systems perform error checks for every volume defined in the session
- Volume and/or group confirmation, selection and initialization occur
- Required real storage frames are page fixed

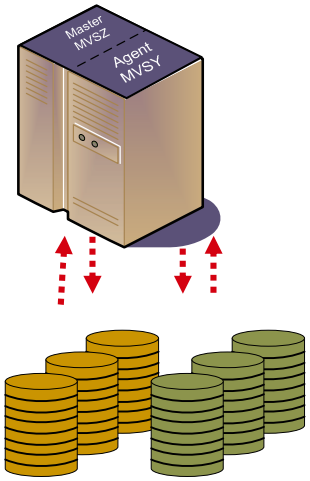

2. Activation

- The Softek TDMF I/O monitor is enabled, and the copy task is started for each volume pair

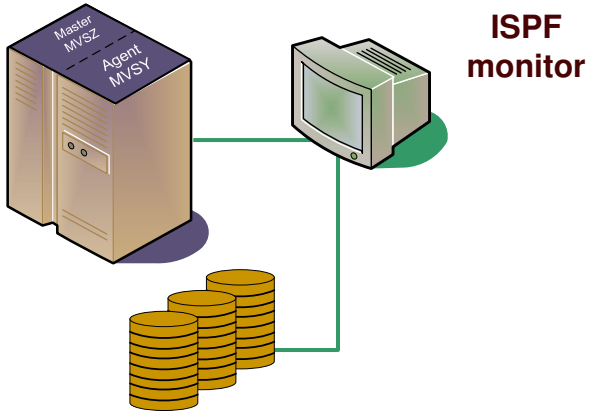

Migration phases (continued)

3. Copy

- Copies source data to target volume asynchronously
- Monitors updates to source volume from all involved systems
- Allows volume pacing, if selected, to occur during this phase

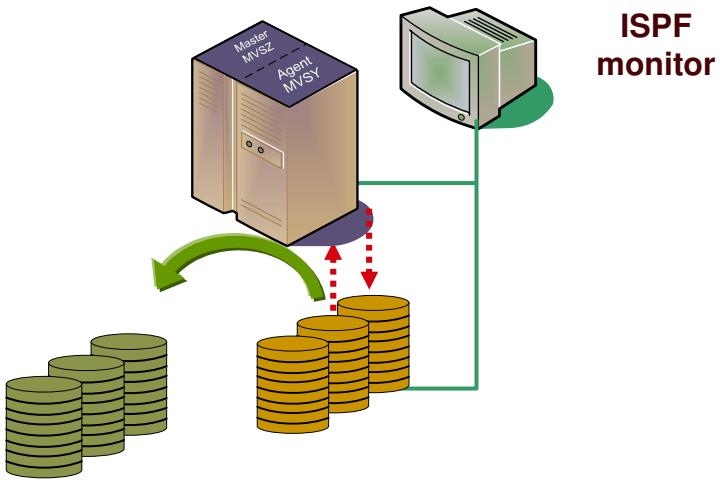

4. Refresh

- Updates made to the source volume during copy phase are asynchronously made to the target volume until the synchronization goal can be met
- Monitors updates to source volume from all involved systems
- Allows volume pacing, if selected, to occur during this phase

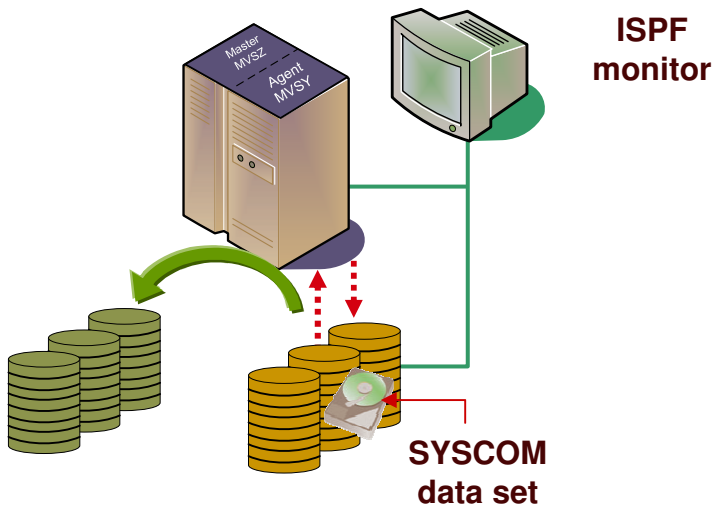

Migration phases (continued)

5. Quiesce

– Inhibits read/write from all systems to the source volume

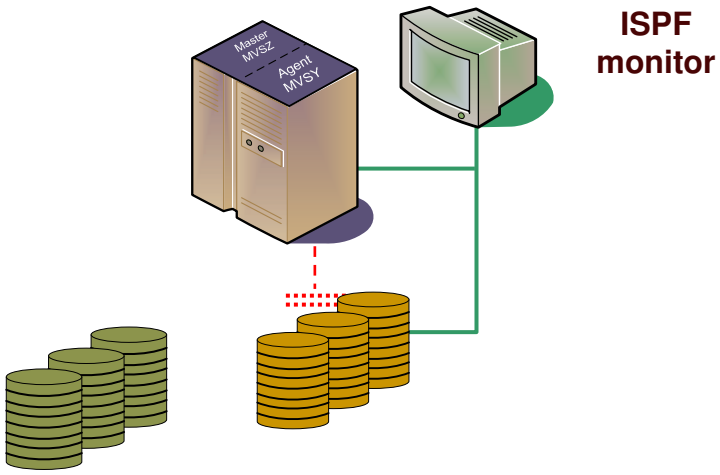

6. Synchronize

– Collects the final set of updates from the source volume and applies it to the target volume

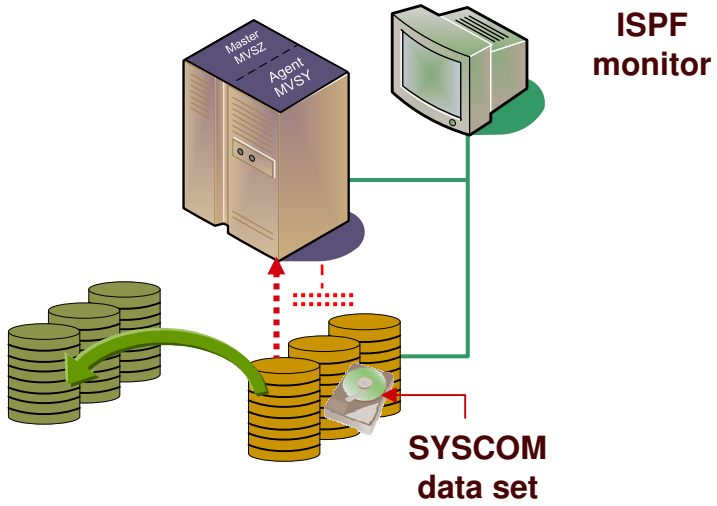

Migration phases (continued)

7. Redirect

– During a swap, the I/O is redirected to the new source volume (original target)

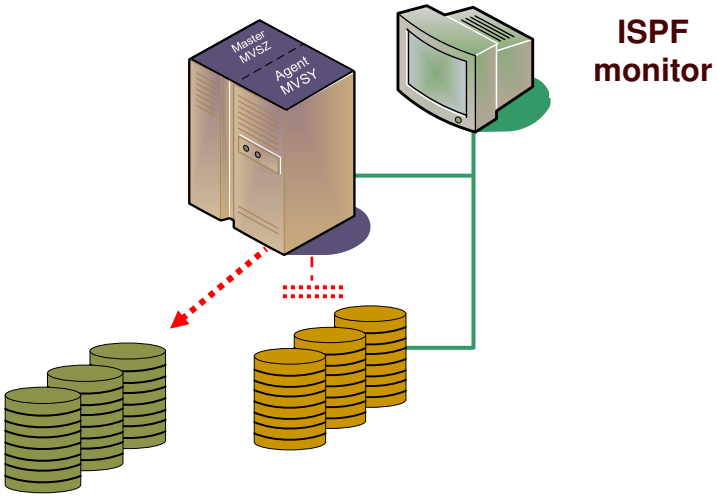

8. Back out

– If a failure occurs during the redirect phase, any volumes that have been redirected must be backed out (i.e., reversed back to the original way)

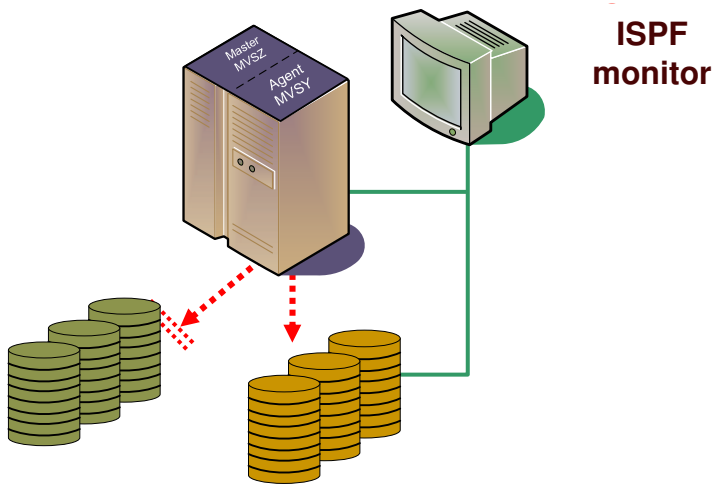

Migration phases (continued)

9. Resume

- Resume volume I/O to the new source volume for all systems
- The Softek TDMF I/O monitor is disabled and the appropriate target volume is marked offline

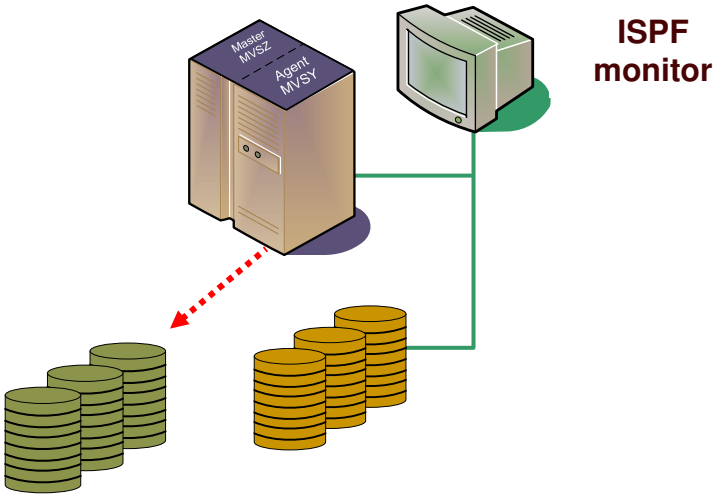

10. Terminate

- Final cleanup takes place and all fixed storage frames are freed

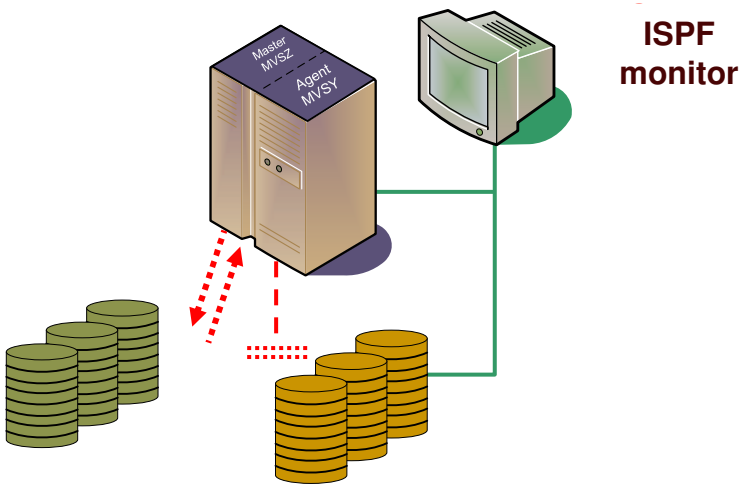

How TDMF z/OS Works with Channel Extenders

Architecture Overview - with Channel Extenders

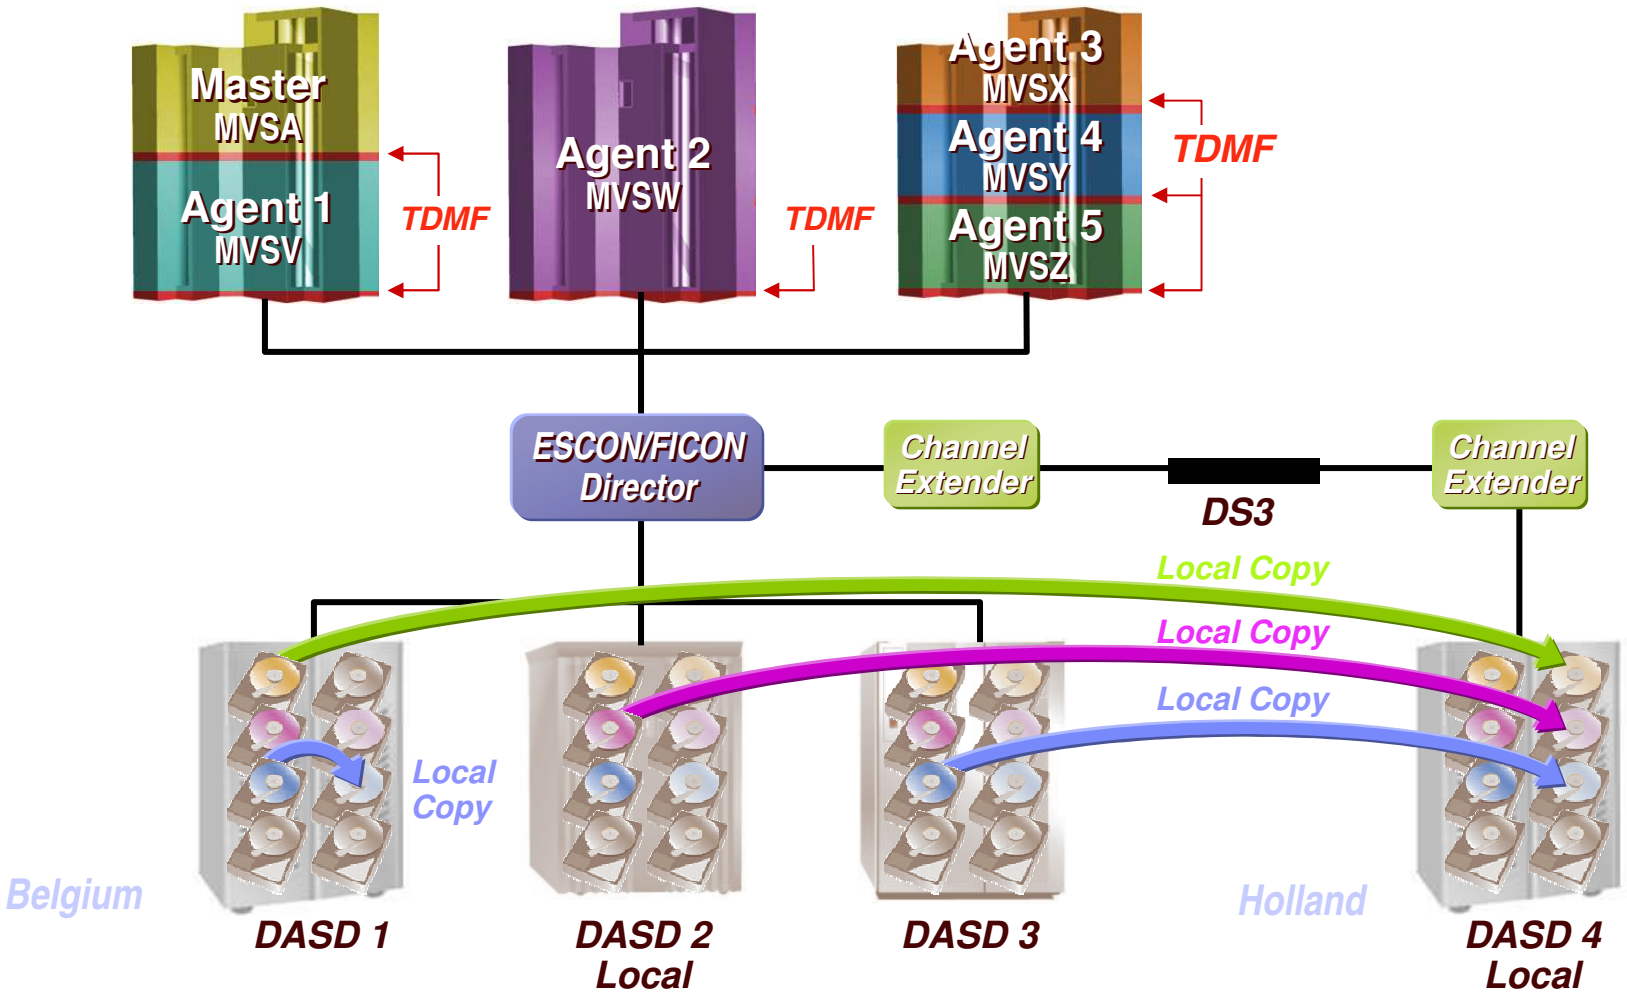

How TDMF z/OS Works Over TCP/IP

Architecture Overview - via TCP/IP

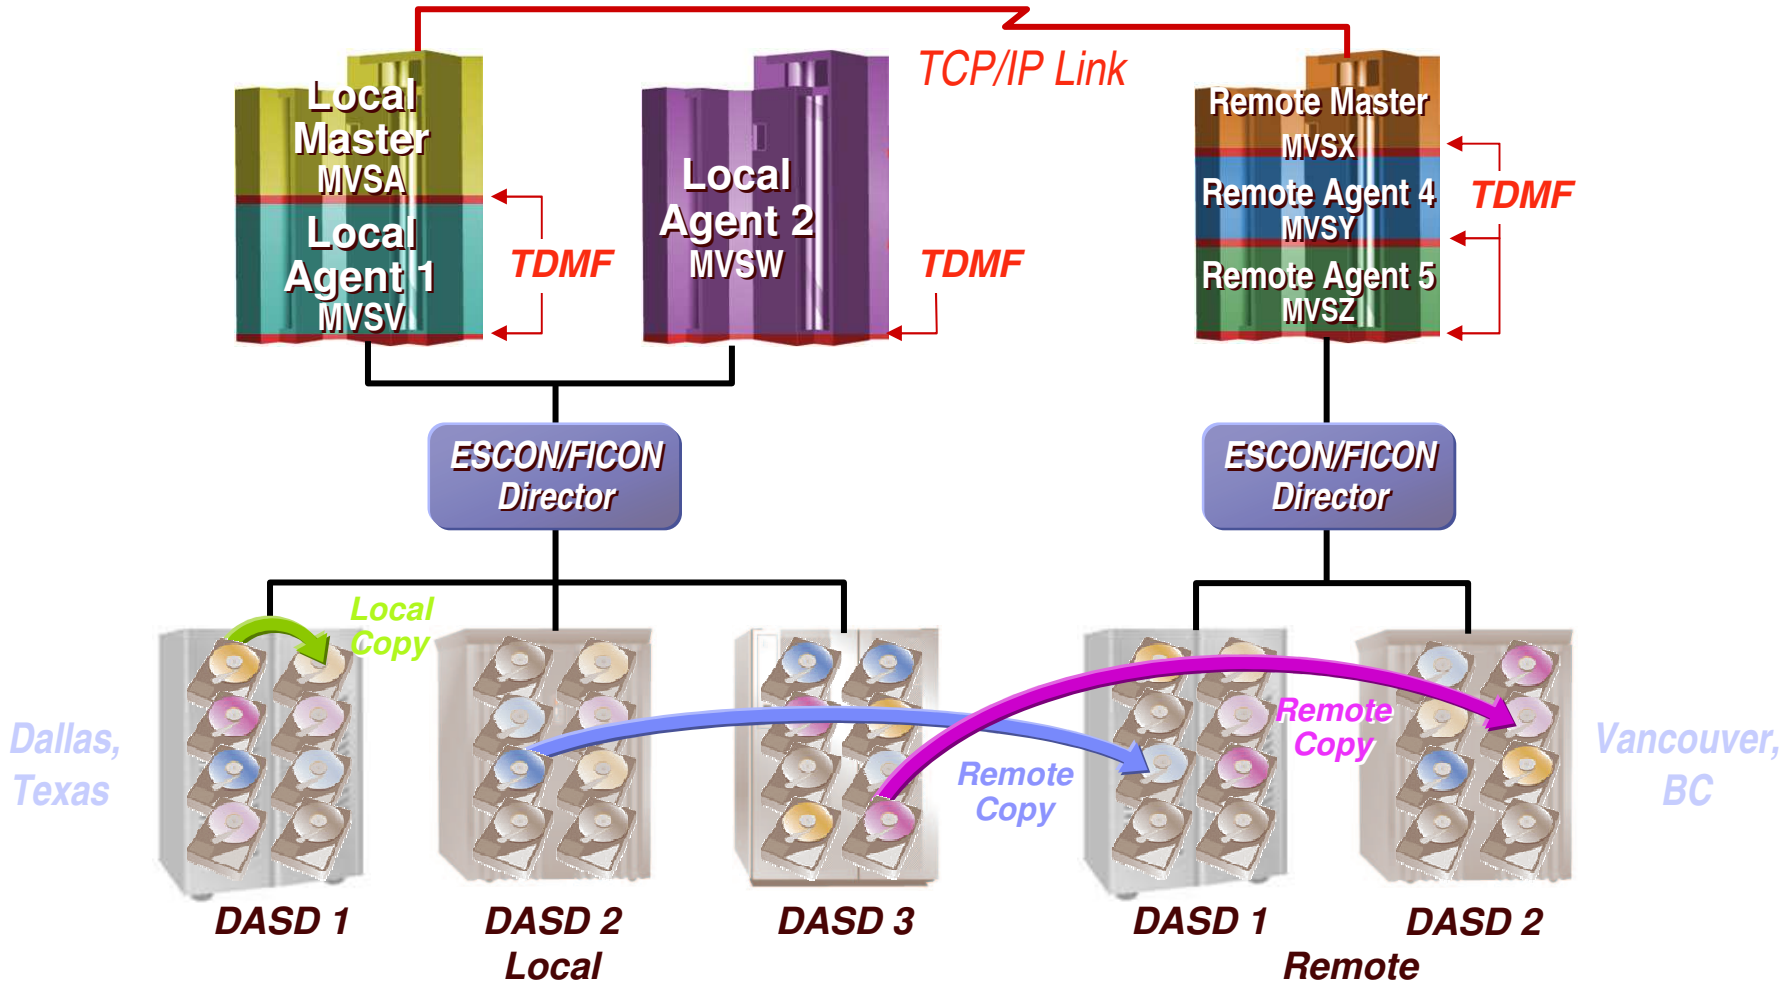

LDMF – Gold Standard for Dataset-Level Migrations

- Non-disruptive dataset migration in z/OS
- Storage-vendor independent
- Device-model independent
- Suspend/resume migration to provide maximum migration flexibility for planning or validation purposes
- Dataset Grouping for easy management of large migrations
- ISPF screens for easy configuration, monitoring and operation
- Improves application performance with dataset relocation
- Load-balance I/O activity across volumes
- Extends system resources
- Reclaim UCBs

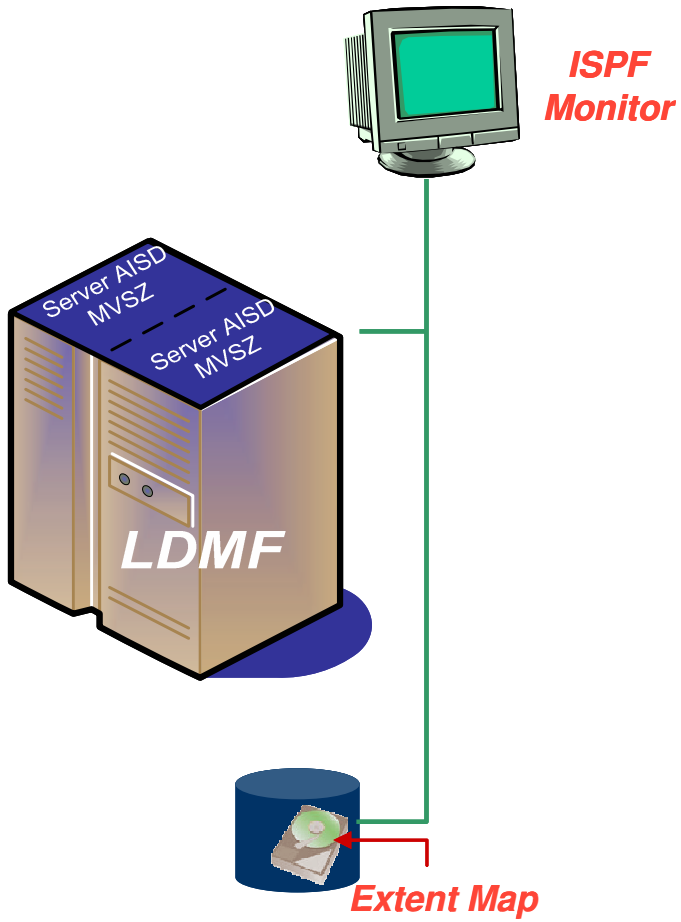

Benefits of LDMF

- Eliminate the need for extended application down time
- Move datasets to high capacity volumes
- Reclaim UCB's
- Improve application performance with dataset relocation
- Balance workloads by moving datasets from poorly performing volumes to lesser used volumes
- Timing

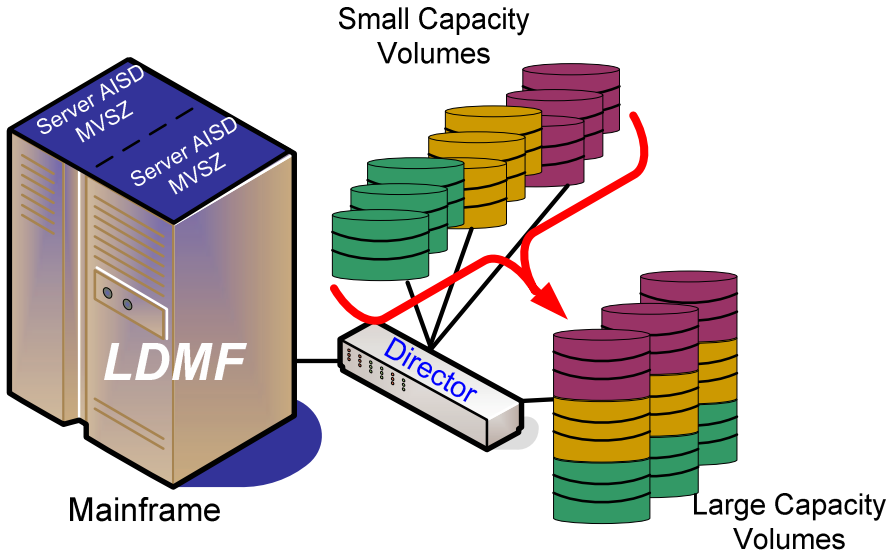

How LDMF Works

Migration Phases

1. Users define migration group with dataset information
2. Allocates datasets on target volumes
3. Copies source datasets to target datasets
4. Synchronizes source and target
5. I/Os to the source Extent are mirrored to its corresponding target Extent
6. Dataset information is swapped and I/O is diverted to new source (source devices still allocated by appl.!)

Post-Migration Phase

1. Completion

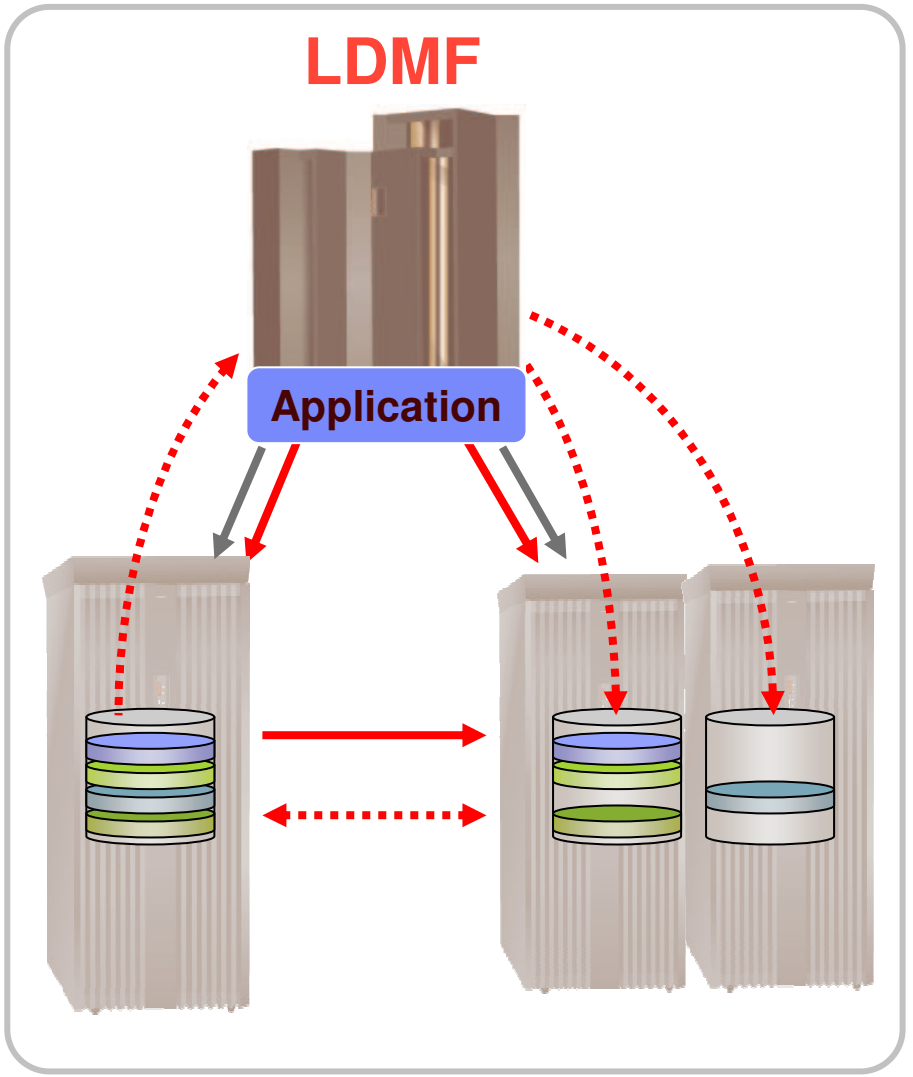

Restrictions

- Datasets not supported
  - VSAM with IMBED, KEYRANGE and REPLICATE
  - Catalogs
  - ISAM
  - Individual PDS members
  - HFS / zFS datasets
  - Page datasets
  - Unmovable datasets (DSORG=U|PSU)
  - VVDS
  - VTOC
  - VTOCIX
  - SYSRES resident data sets
  - Sysplex couple data sets
- 'TO' control unit must be equal or higher type
  - Example: no 2105→3990
  - CCW command set differences
- Hardware reserves must be converted
- Currently, other "UCB swap" software must be disabled (e.g. GDPS Hyperswap)
- SMS managed data sets "only"

LDMF is ONE of the tools for data set mobility

- Consider the whole spectrum of ”data set mobility solutions”
  - DFDSS copy/delete can do it (if you can allocate the data set)
  - DFHSM migrate/recall can do it (if you can allocate the data set)
  - Daily/weekly/monthly batch with delete/define will do fine
  - Database reorganization
  - .....
- Use LDMF for remaining supported data sets

TDMF UNIX

Simple Local Migration for UNIX

- Non-Disruptive “Switchover” Migration
  - Transparent
  - Divert reads to new storage
- Works with any File System
- Works with or without a Volume Manager
- Scalable to 1000 concurrent volumes of any size
- Fast, easy host-based install
- Dynamic Configuration
- Script-friendly
- Throttled Copy
- Supports HA environments
- Fully recoverable
- Similar functionality planned on Windows

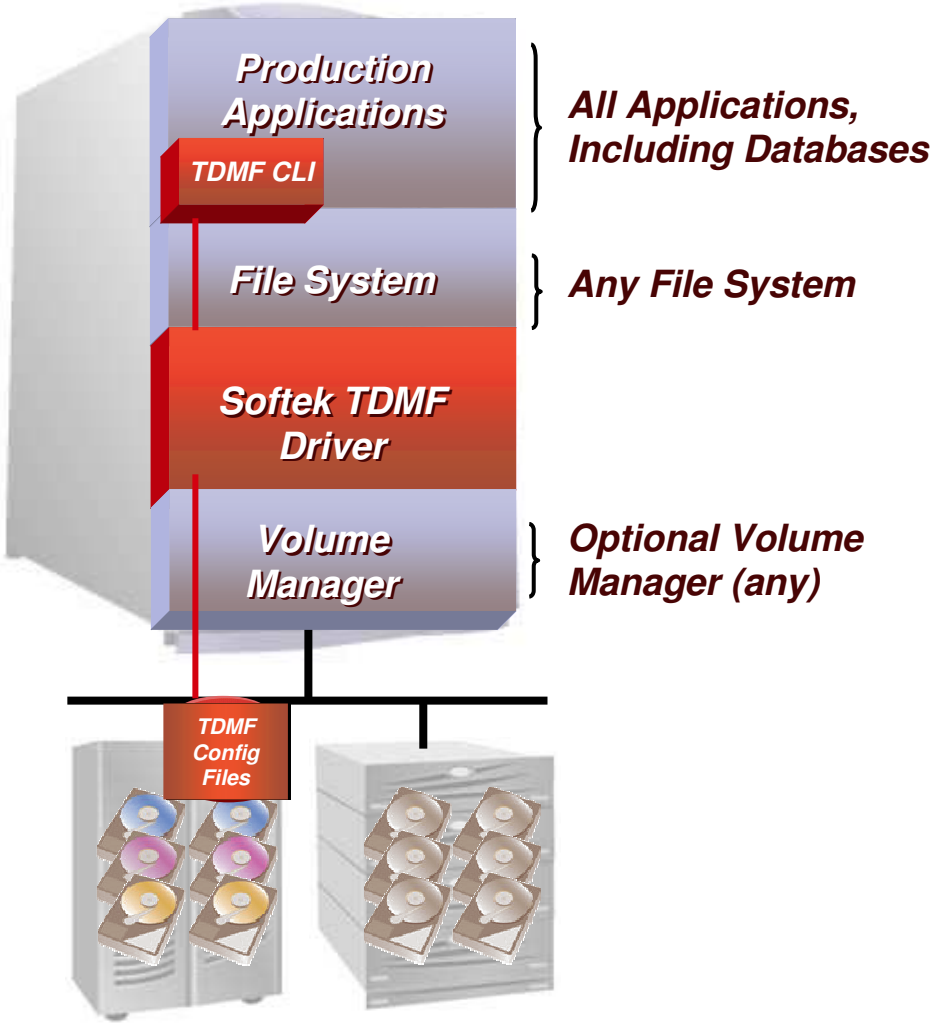

Softek TDMF UNIX Typical Migration Process

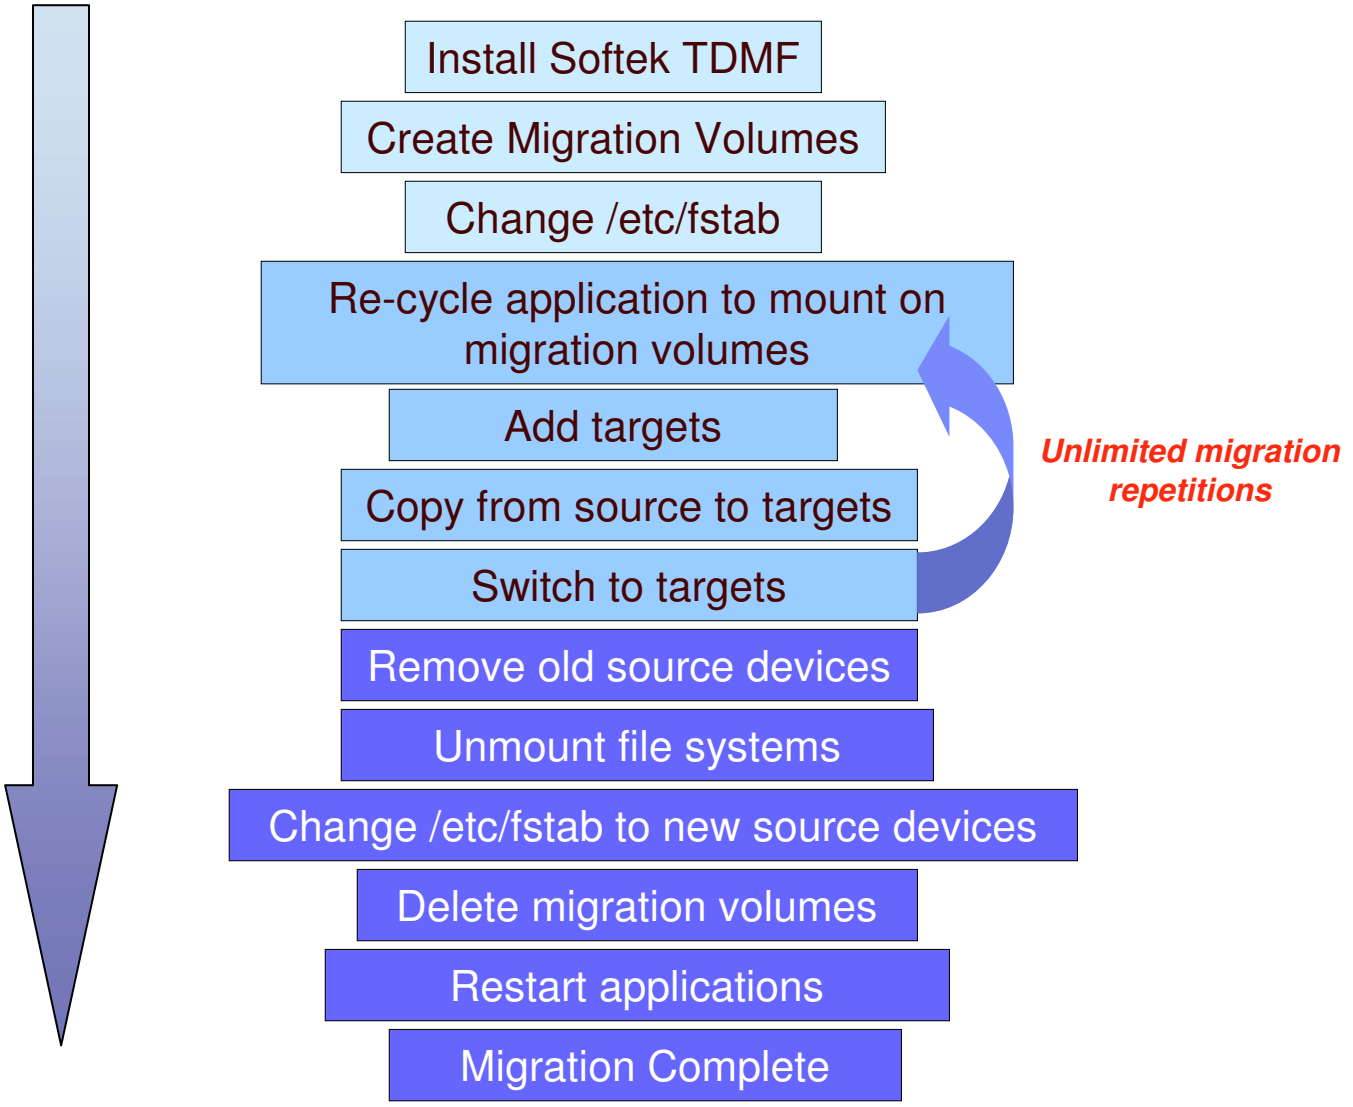

Working with Clusters

- TDMF UNIX only fully supports failover clusters
- Define one group per Cluster Service/resource group
- Database must be on shared storage otherwise copy must be restarted following failover
- New storage must be created as cluster resource to support failover
- Set dependencies to control start/stop order:
  - TDMF depends on underlying LV/disk
  - Application/mount point depend on TDMF
- Install/configure on one node only whilst offline
  - Provides fallback option if issues occur with TDMF on node
  - Typically will not want overhead of migration if issue occurs causing failover
- Only one node can be active on Oracle RAC whilst copy in progress
  - Write Serialization issue

SVC or TDMF? Different tools but TDMF is about Migration NOT Virtualization

- SVC is a powerful, open systems, virtualization appliance.
- A feature of SVC is Data Migration for volumes that have been placed under SVC control.
- SVC does like for like data migration (no consolidation or volume size changes).
- TDMF/Replicator is not a virtualization tool.
- TDMF/Replicator is a host-based, heterogeneous platform data migration tool.
- TDMF/Replicator can help migrate/consolidate volumes to SVC control non-disruptively.

TDMF (IP) - Global and Server Migrations for UNIX, Linux and Windows

- Simple installation
- Online, dynamic configuration
- Supports all native file systems
- Works with any Volume Manager or no Volume Manager
- Logical volume migration
- Throttling and compression to minimize network bandwidth impacts
- Ensures data consistency
- Script friendly
- No prerequisite software, hardware or firmware
- Continues to track updates if network or remote server fails
- No limit on number of files
- Unified Common Console across platforms

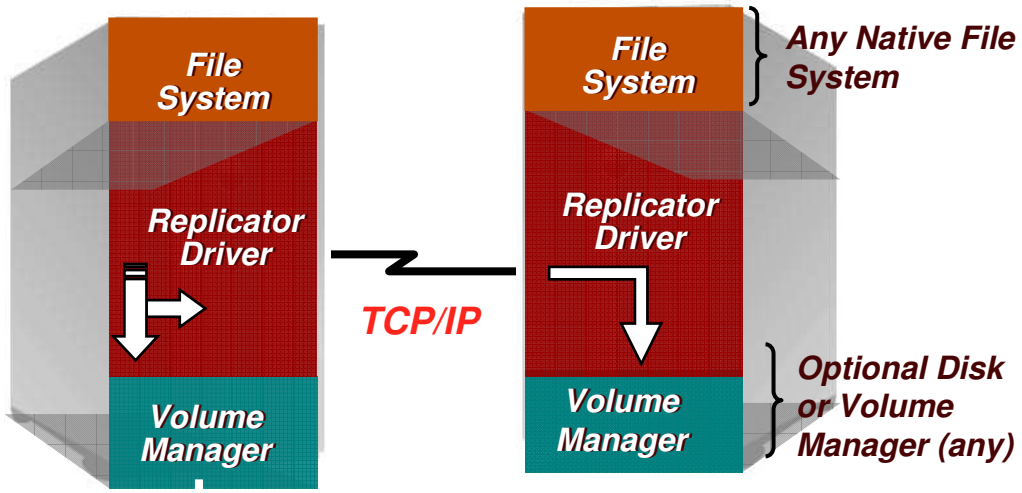

Mounted Source Volumes

Unmounted Target Volumes

Introduction to TDMF IP UNIX

Two Hosts Configuration

- Data migration to a remote site
- Ensures that target is always kept in a consistent/usable state
  - Memory buffer
  - Pstore
  - Journal
- Continues to track updates if network or remote server fails
- Facilitates migration back to the primary site

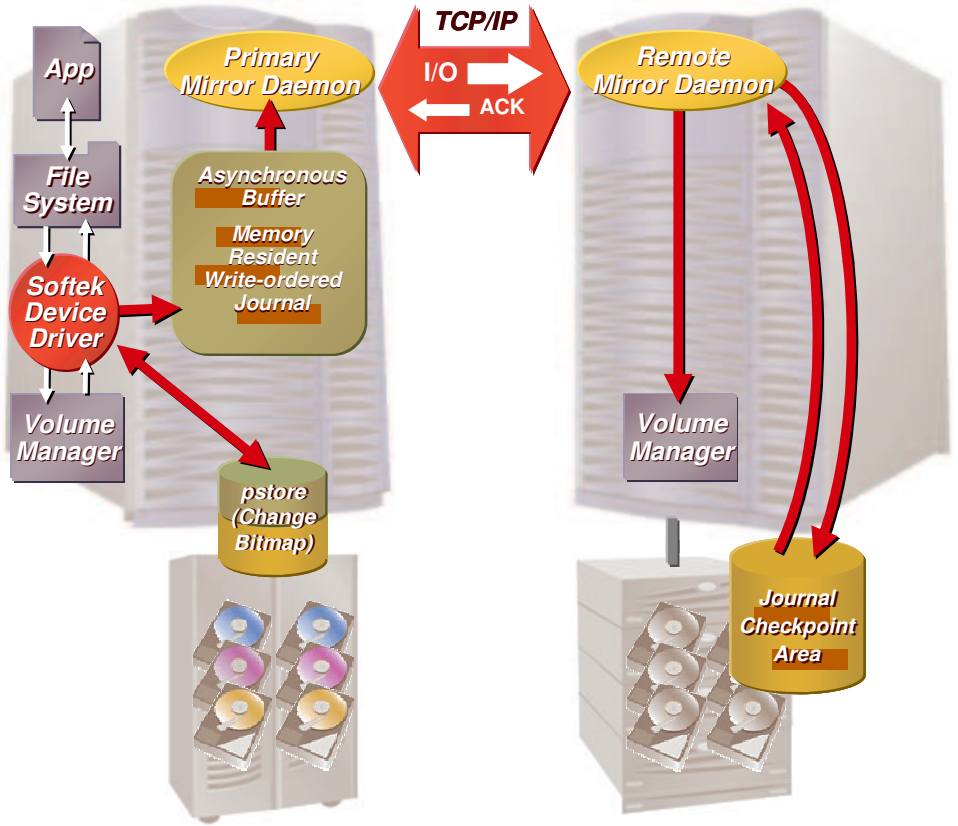

Operating Modes

Refresh

Initial synchronization

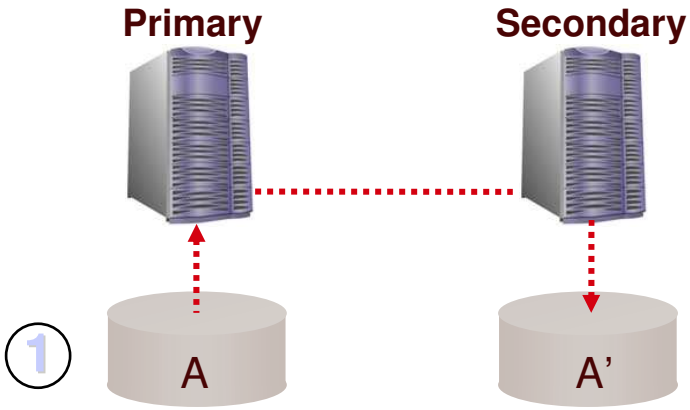

Normal Mode

Block-level changes are migrated using “copy-on-write” mechanism

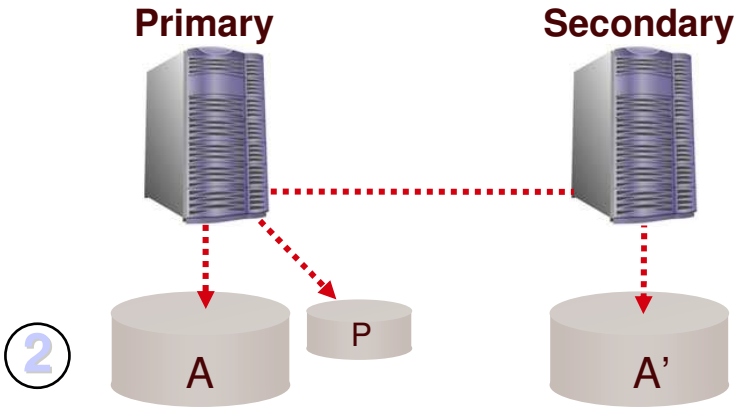

Tracking

A. Contact with secondary lost
B. Changes tracked in PStore

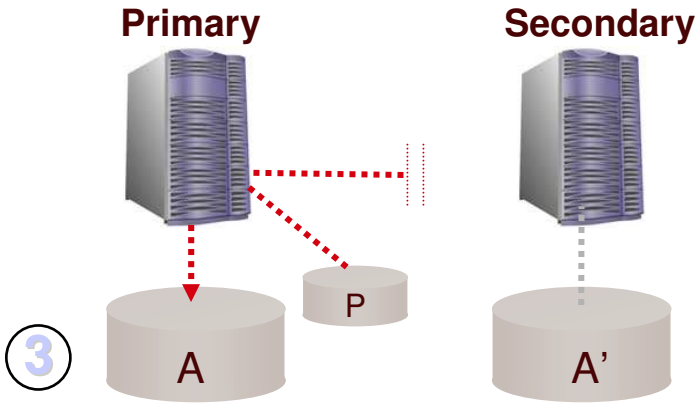

Smart Refresh

A. Bitmap used to send changes
B. Resume migration

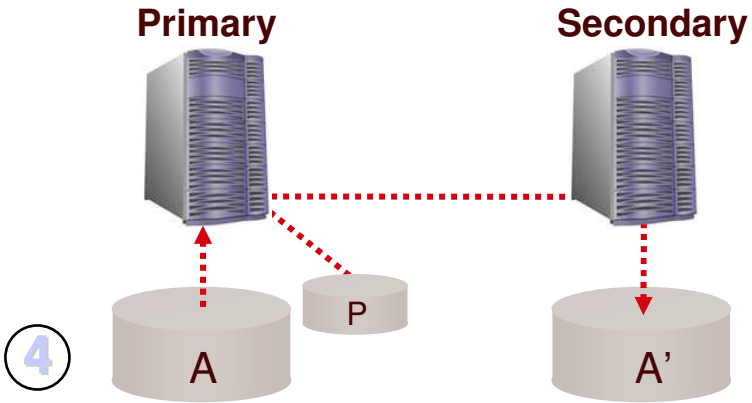

Taking a Checkpoint (PIT)

Normal mode

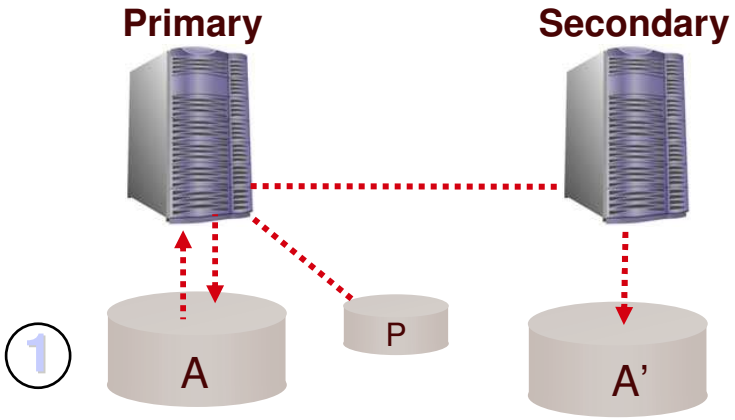

Checkpoint

Checkpoint unlock secondary volumes for read access

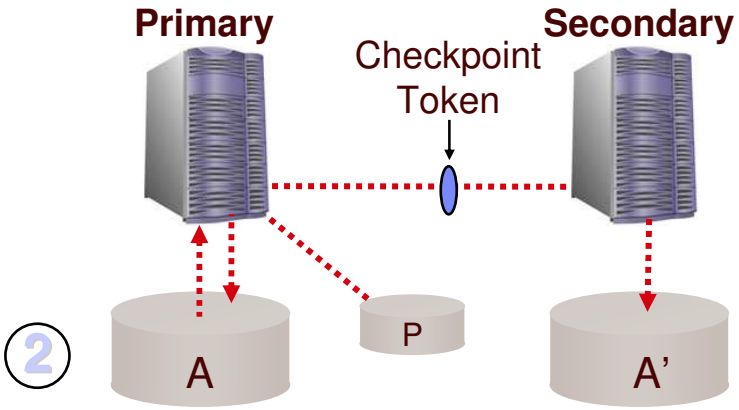

Checkpoint

PIT volumes are available in read only mode to validate. Updates tracked in pstore.

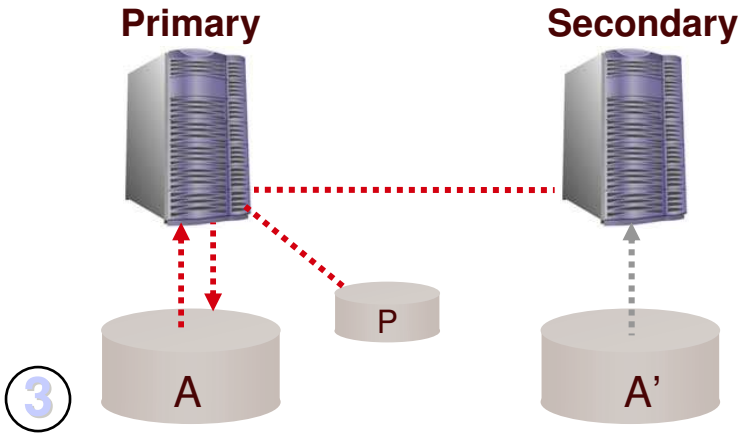

Resume

PIT volumes released, volumes re-synchronized from source using pstore.

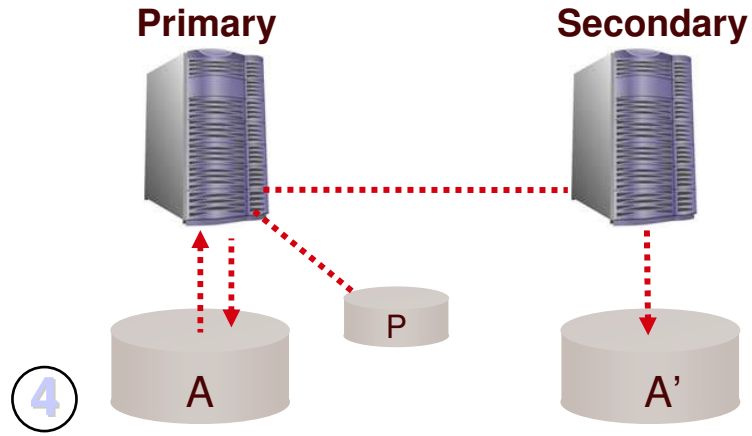

Replicator Key Features

- Asynchronous Distance Replication
- Maintains write order consistency
- Creates PIT (Point-In-Time) Copies
- Integrates seamlessly into existing production environment
- Managed by Common Console
- Recovery in the event of failure
- Persistent Store (pstore) to track all changed data
- Journal File System ensures data integrity
- Supports Virtualized Environments (VMWare)

Softek Replicator

- Replicator is independent of application
- Works with all systems
- Replicates both logical volumes and physical volumes (UNIX)
- No prerequisite software
- Works with or without any volume managers (UNIX) and with all disk managers (MS Windows)

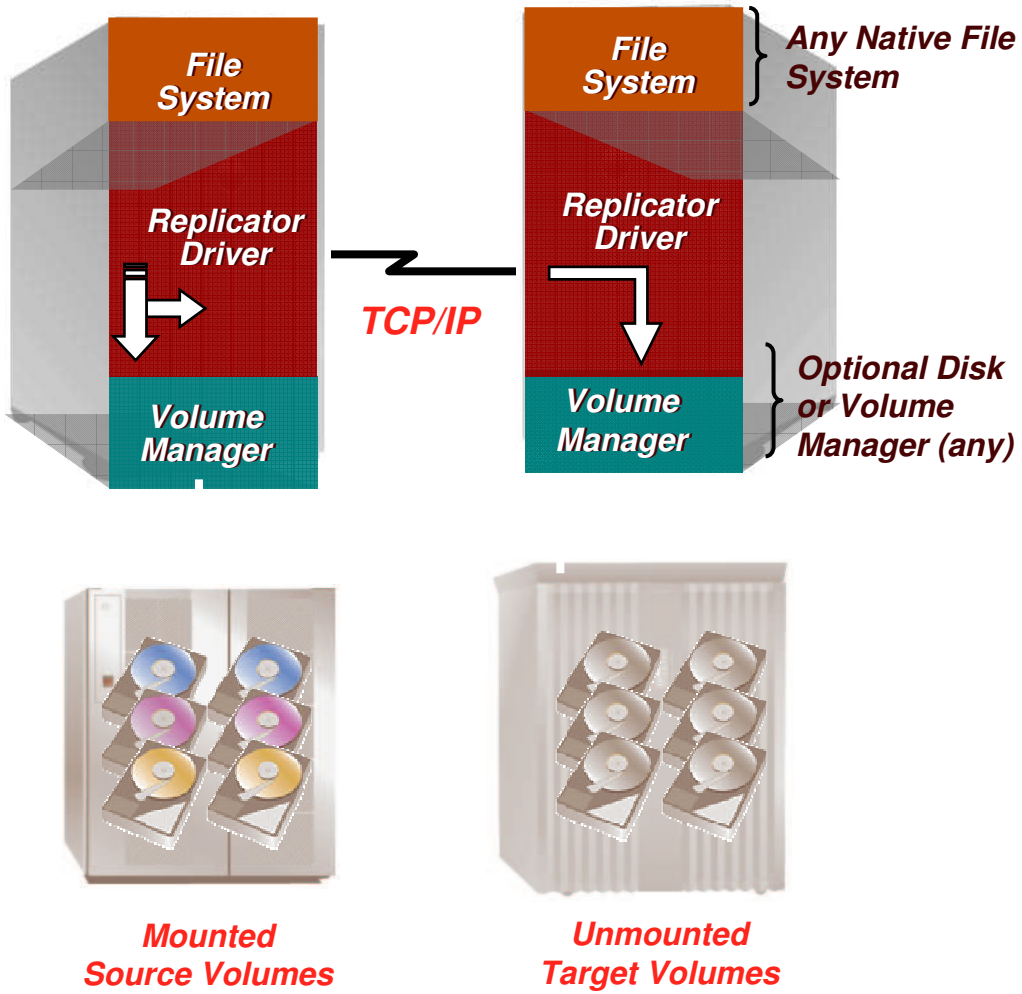

TS7700 Tape Virtualization Engine update

The new IBM Tape Virtualization Engine for System z

- The TS7700 Virtualization Engine is the latest in the tape virtualization solutions for the IBM System z environment
  - Builds on almost 10 years of enterprise tape virtualization experience
  - Designed to provide improved performance and capacity to help lower the total cost of ownership for tape processing
  - Introduces a modular, scalable high-performance architecture
- Supports attachment to multiple operating systems
  - z/OS, z/VM, z/VSE & z/TPF
  - Same level of host software support as 3494 - B20
- Expanded cache and performance
- Standalone and business continuation configurations

Began Customer Shipment 9/26/06

TS7700 Virtualization Engine Components

- TS7740 Virtualization Engine (3957 Model V06)
  - Power5+ architecture server based
  - Two dual-core, 64-bit, 1.9-GHz processors
  - Runs the V and H nodes
- TS7740 Cache Drawer (3956 Model CX6)
  - RAID array expansion
  - 16 15K 146GB FC HDDs
  - 1.5 TB usable capacity (after RAID and spares)
- TS7740 Cache Controller (3956 Model CC6)
  - Disk RAID array controller
  - 16 15K 146GB FC HDDs
  - 1.5 TB usable capacity (after RAID and spares)
- 3952 Model F05 Frame
  - Houses major components & support components
  - Dual Power

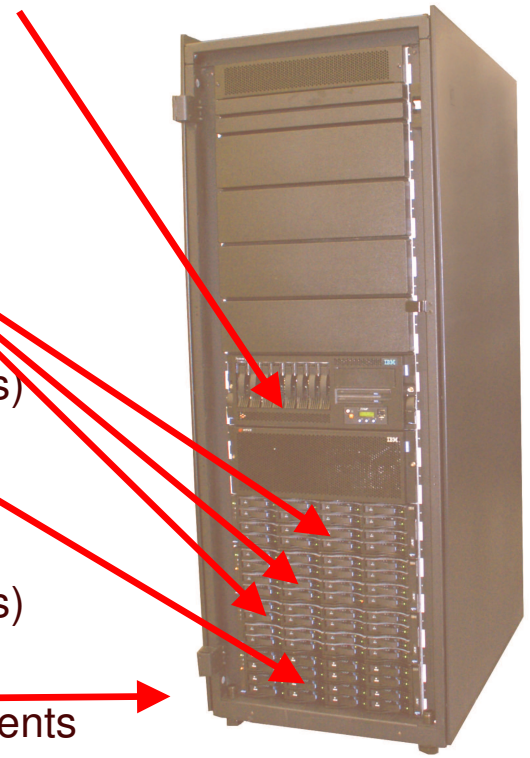

The combination of the Virtualization Engine components is called a TS7700 Cluster

Cache Configuration

- Two drawer (3 TB usable capacity)
  - One TS7740 Cache Controller (3956 Model CC6)
  - One TS7740 Cache Drawer (3956 Model CX6)
- Four drawer (6 TB usable capacity)
  - One TS7740 Cache Controller (3956 Model CC6)
  - Three TS7740 Cache Drawer (3956 Model CX6)
- Cache enablement features
  - Each feature enables the use of 1TB of cache space
  - 2-6 features can be ordered (2-6TB of cache)

With the two physical cache capacities and enablement features, the cache size can be tailored to a customer's needs and provides for future growth

TS7700 Solution Elements

3494 library support - 5/2007

Series p server base (power5)

RAID Disk Cache

Virtualization engine software with integral HSM and peering

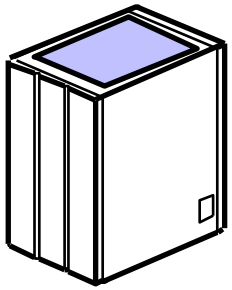

TS7700

TS7700 Cluster

Drives/Library

Industrial Intel-based server

FCP switches for tape drives

Library Manager software

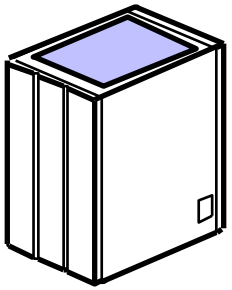

3953 (Only for TS3500)

3494/TS3500 - 3592/TS1120

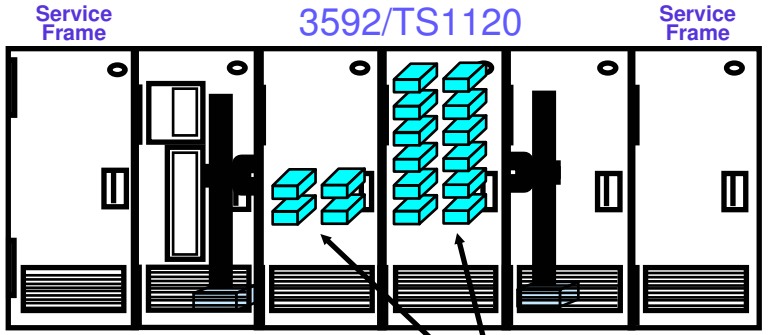

3592 J1A, 3592 E05/E JA, JB and JJ media (up to 700GB native Capacity)

TS7700 Virtualization Engine Solutions

- TS7700 Virtualization Engine
  - Up to 6TB usable cache (18TB @ 3:1 C/R)
  - Four 4Gbps FICON interfaces (512 logical paths)
  - 256 virtual tape devices
  - 500,000 logical volumes
  - Logical volume sizes up to 4,000 MB
  - Advanced policy management
    - Volume pooling
    - Cache management
    - Logical volume size
    - Reclamation
  - Up to 16 4Gbps fiber TS1120 tape drives in a TS3500 library, up to 12 drives in a 3494
    - Native model E05 mode stores up to 500GB on JA media and 700 on JB media
    - Encryption of data on physical volume pools

The combination of the Virtualization Engine components is called a TS7700 Cluster

TS7700 Support for TS1120 Encryption Capable Drives

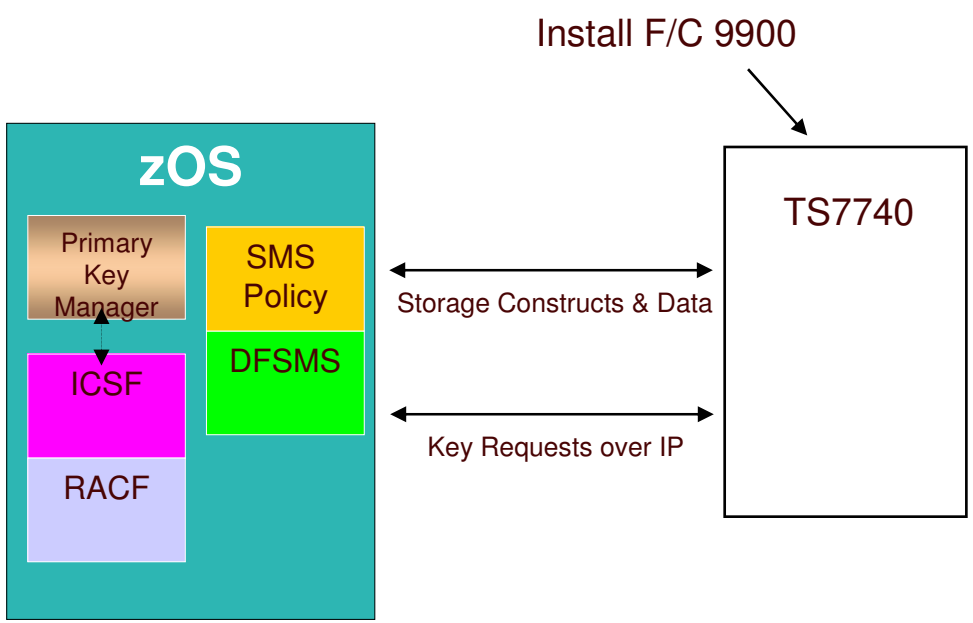

- Logical volumes assigned Storage Group names
- Storage Group name points to pool
- Each pool can be set up to encrypt or not
  - Two - 64 Character Key Labels
  - Label or Hash
- Data is encrypted when moved to a pool set up for encryption
  - Keys obtained through out of band method

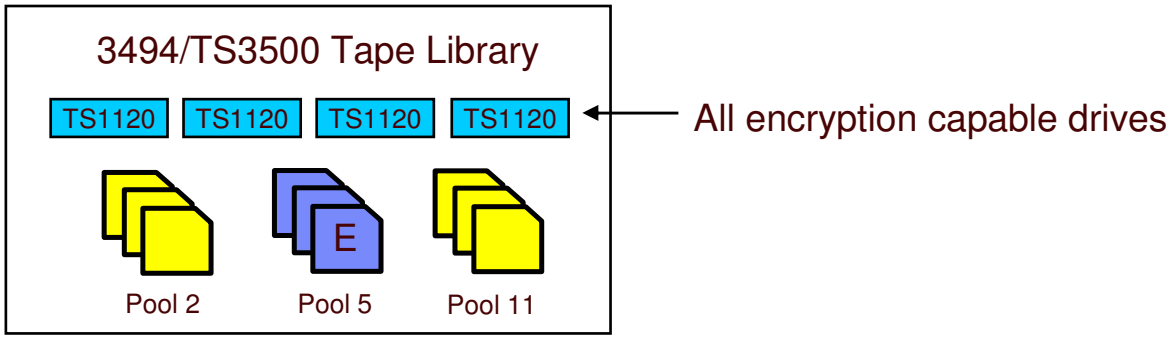

Available 3/09/2007

TS7700 Virtualization Engine Solutions

Supports RTO/RPO measured in seconds

- TS7700 Grid Configuration
  - Couples two TS7700 Clusters together to form a Grid configuration
  - VTCs have been eliminated
  - Hosts attach directly to the TS7700s
  - Any volume accessible through either TS7700 cluster in the Grid configuration
  - I/P based replication
    - Two 1 Gbps Ethernet links
    - RJ45 Copper (Cat 6) or SW Fiber Optic
    - Standard TCP/IP
  - Policy-based replication management
- Can be configured for disaster recovery or higher availability environments

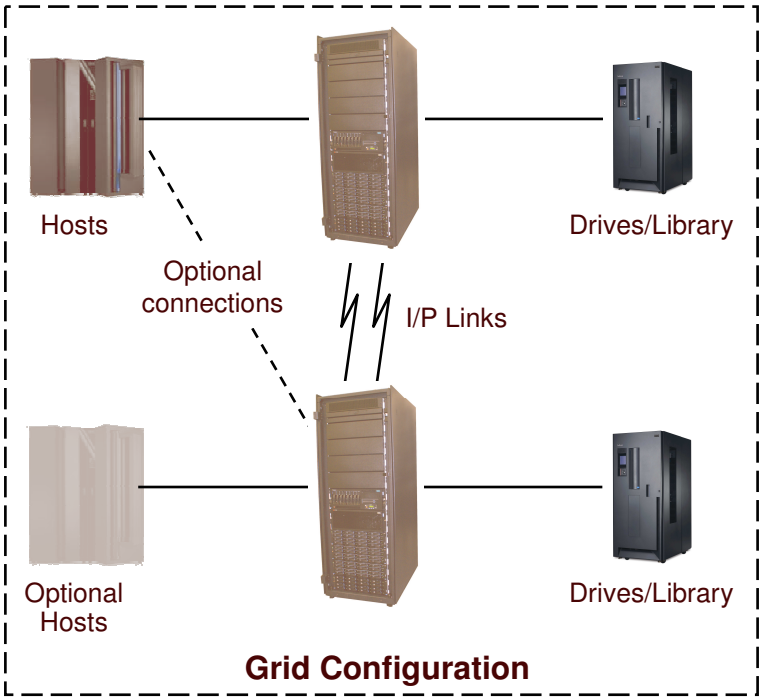

I/P replication may greatly simplify the infrastructure and management needed for a disaster recovery solution as compared to IBM's existing PTP VTS

TS7700 Virtualization Engine Solutions

Supports RTO/RPO measured in seconds

- Dual Site HA+DR Grid Configuration
  - Both site hosts connected to both TS7700s
  - All data accessible through either TS7700
  - Copies use RUN copy consistency point (at unload time)
  - Operations continue with the failure/service/upgrade of either TS7700
    - Automatic with service/upgrade
    - Automatic for most failures
    - One action for all other failures
  - Operations continue with site failure
    - Automatic for most failures
    - One action for all other failures

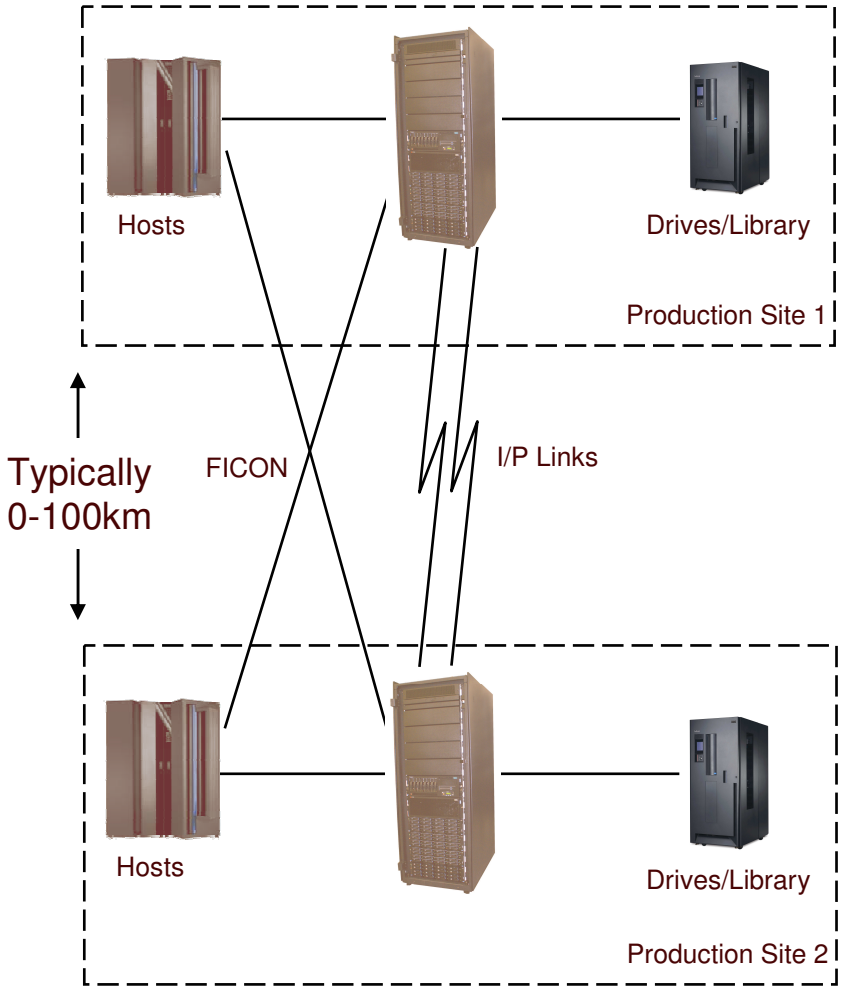

Grid Configuration

TS7700 Virtualization Engine Solutions

Supports RTO/RPO measured in seconds

- HA+DR Grid Configuration
  - Main site hosts connected to both TS7700s
  - All data accessible through either TS7700
  - Copies use RUN copy consistency point (at unload time)
  - Operations continue with the failure/service/upgrade of either TS7700
    - Automatic with service/upgrade
    - Automatic for most failures
    - One action for all other failures

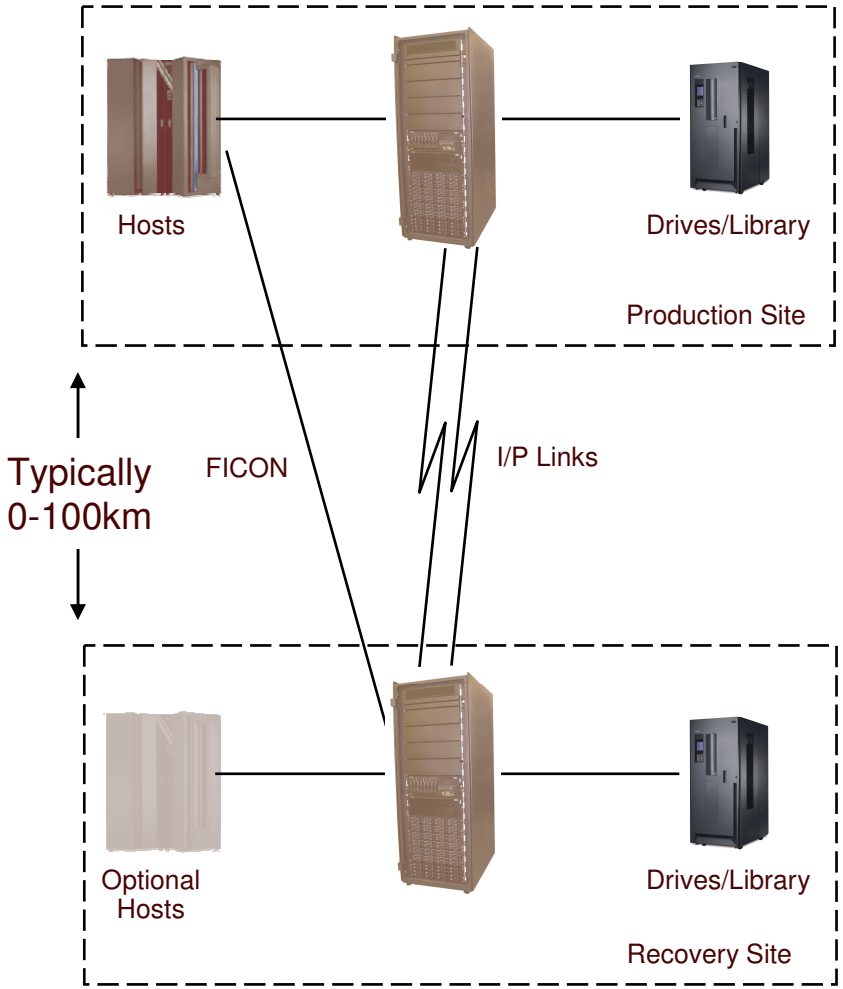

Grid Configuration

TS7700 Virtualization Engine Solutions

Supports RTO measured in seconds/RPO measured in minutes

- DR Grid Configuration
  - Main site hosts connected to one of the TS7700s
  - Typical separation is 100s of kms
  - Deferred copy consistency point
  - Operations continue after a disaster
    - Place recovery site TS7700 in one of the takeover modes

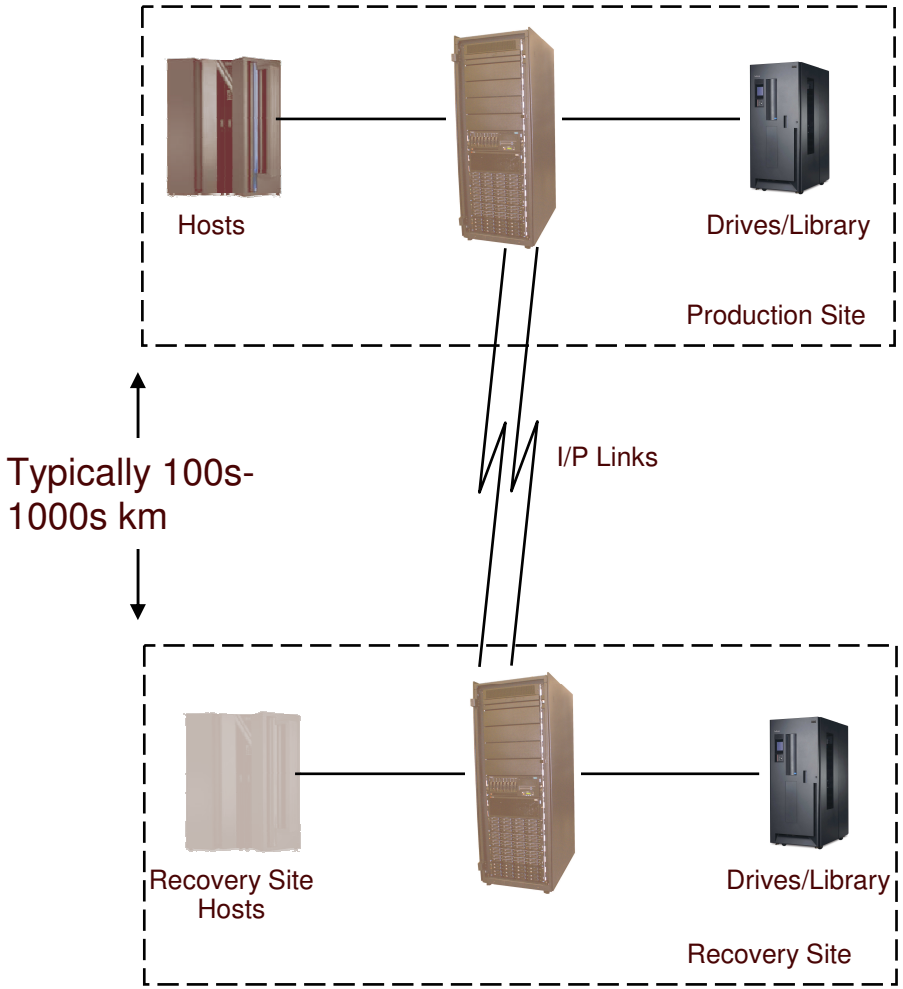

Grid Configuration

TS7700 Key Attributes - Grid Configuration

Virtual Drives	512 (combination of both TS7740s)
Virtual Volumes	500,000
Host Interface Types	4Gb FICON
Peak Host Bandwidth	900MB/Sec write, 900MB/Sec read @ 3:1 c/r
Cache Size (max)	6, 9, 12 TB (combination of both TS7740s)
Physical Drives	Up to 32 (combination of both TS7740s)
Interconnect	Dual 1Gb Ethernet, TCP/IP
Data Movement	Full duplex, dual production sites
Distance	LAN and WAN distances
Replication	Immediate, Deferred, None
Replication Control	DFSMS Policy Based, Defaults
Resynchronization	Automatic
Continued Data Accessibility	Planned/Unplanned Outages, Upgrades
Testing Support	User set special modes

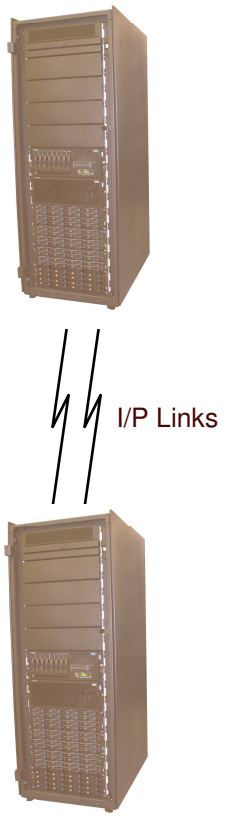

Industry Leadership

TS7700 Enhancements during 2007

R1.1 1/26/07

- Autonomic Ownership Takeover Manager
- 3592 E05 native, new only, no intermix with J1A
- JB Media
- Disk Cache two drawer configuration
- Cache increments 2 – 6 (no MESs) No removal
- Fiber connection for Grid – Shortwave only, no intermix with copper
- B10 or B20/3584/3592 data migration (standalone box swap only)
- 256 virtual devices

R1.2 03/09/07

- Broadband call home
- Out of Band Encryption by Pool
- 3494/3592 Support (5/18)

Outboard Data Migrations

- B10 or B20 3584/3592 PTP to Grid Data Migration (5/04)
- Migrations with 3494 (5/18)
- Library Move Migrations (6/1)
- Merge Migrations (3Q)

R1.3 3Q/07

- 3 Cluster Grid
  - Grid MESs
- Host Console Query
- Copy Export for standalone
- 1 million logical volumes
- Remote Access Read/Write Pipelining
- Automated Read-Only Recovery
- Secure Data Erase w/encryption
- Performance Increments 1-6
- Field upgrade features

R1.4 4Q/07

- Copy Export for Grid
- Single Cache Drawer
- 1TB Cache Enablement

List of potential enhancements is an estimate of current intent and is subject to change and is not a commitment to deliver.

MES & Migration Rollout Plan

- R1.1/R1.2
  - B10/B20 3494 to B10/B20 3584
  - B10/B20 Standalone 3584 to TS7740 3584
  - B10/B20 PtP 3584 to TS7740 Grid 3584 – 05/04/07
  - B10/B20 Standalone or PtP 3494 to TS7740 Standalone or Grid 3494 – 05/11/07
  - B10/B20 Standalone or PtP 3494 to TS7740 Standalone or Grid 3584 – 06/01/07
- R1.3
  - 2 Standalone VTS B10/B20 attached to 3584 to Standalone TS7740 attached to same 3584
  - 2 Standalone VTS B10/B20 attached to 3494 to Standalone TS7740 attached to diff 3584
- R1.4
  - Standalone VTS attached to a 3584 to Standalone TS7740 attached to diff 3584
  - 2 Standalone VTS attached 3584 to Single Cluster TS7740 diff 3584
  - 2 Standalone VTS B10/B20 attached to 3494 to Standalone TS7740 attached to same 3494

TS7700 – Three Cluster Grid - H/A & Remote

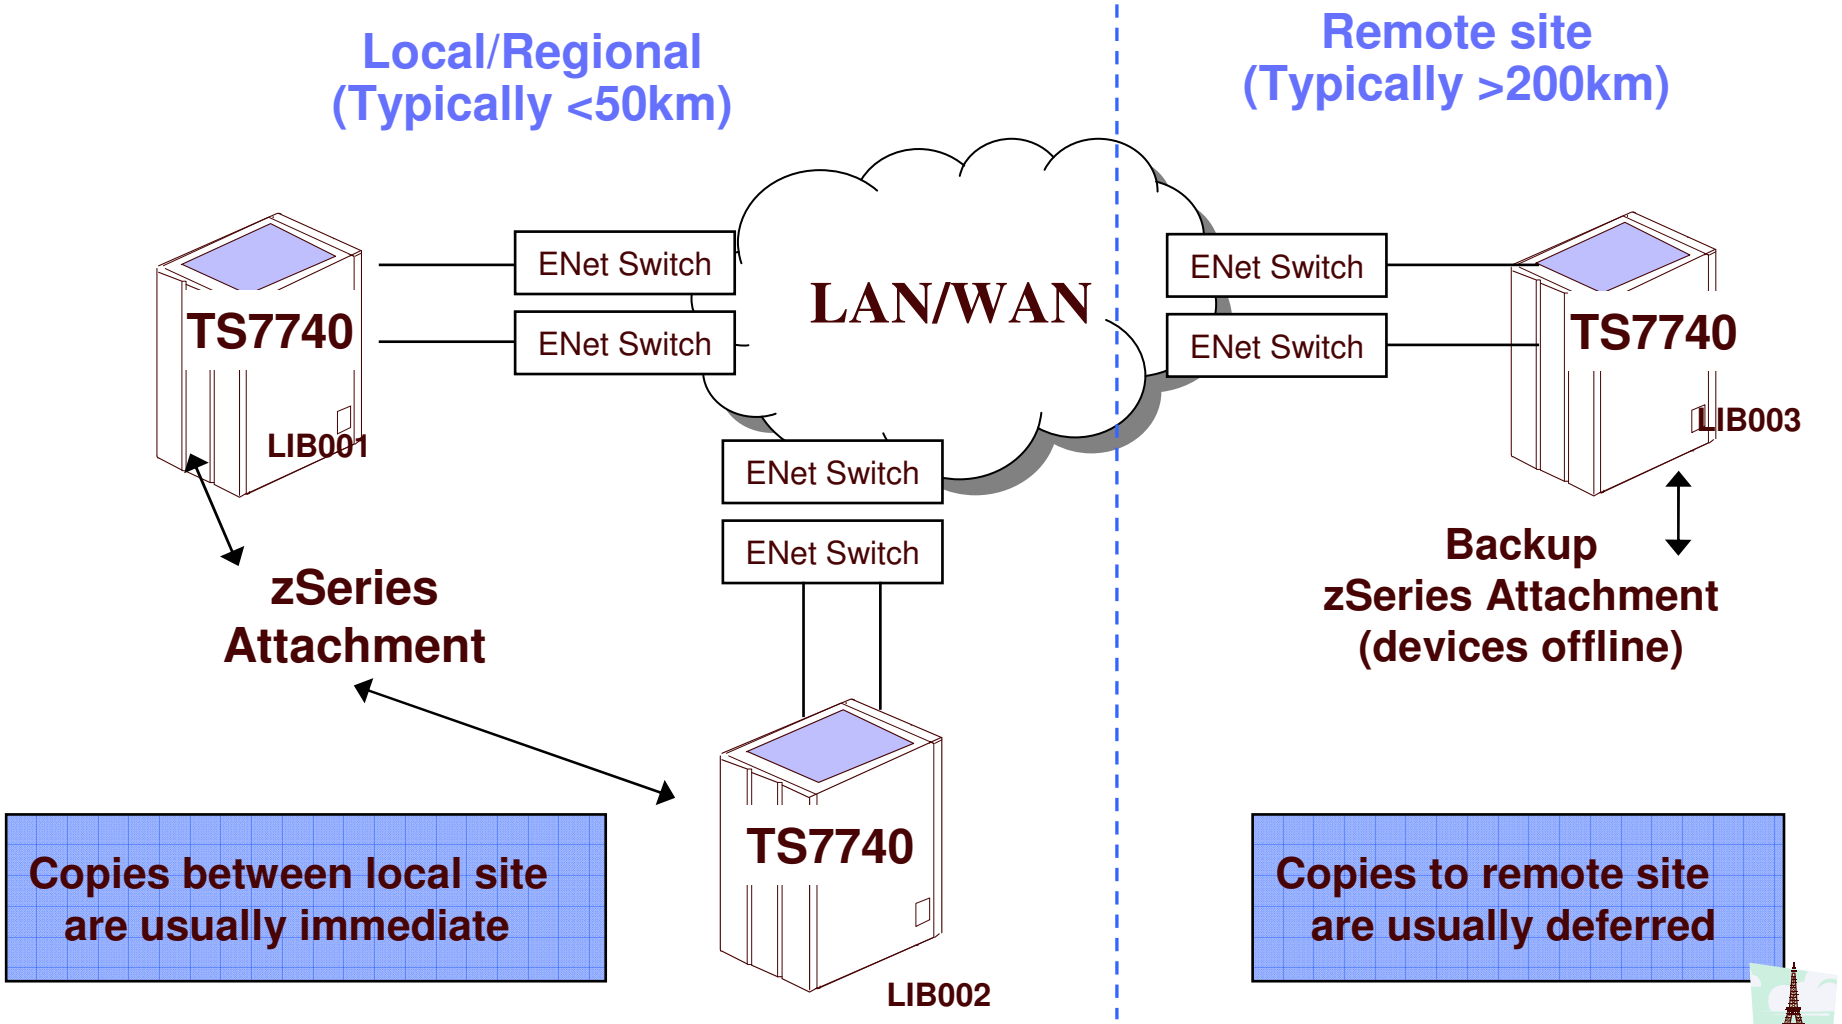

[1]List of potential enhancements is an estimate of current intent and is subject to change and is not a commitment to deliver.

TS7700 – Three Cluster Grid - Two Centers & Remote

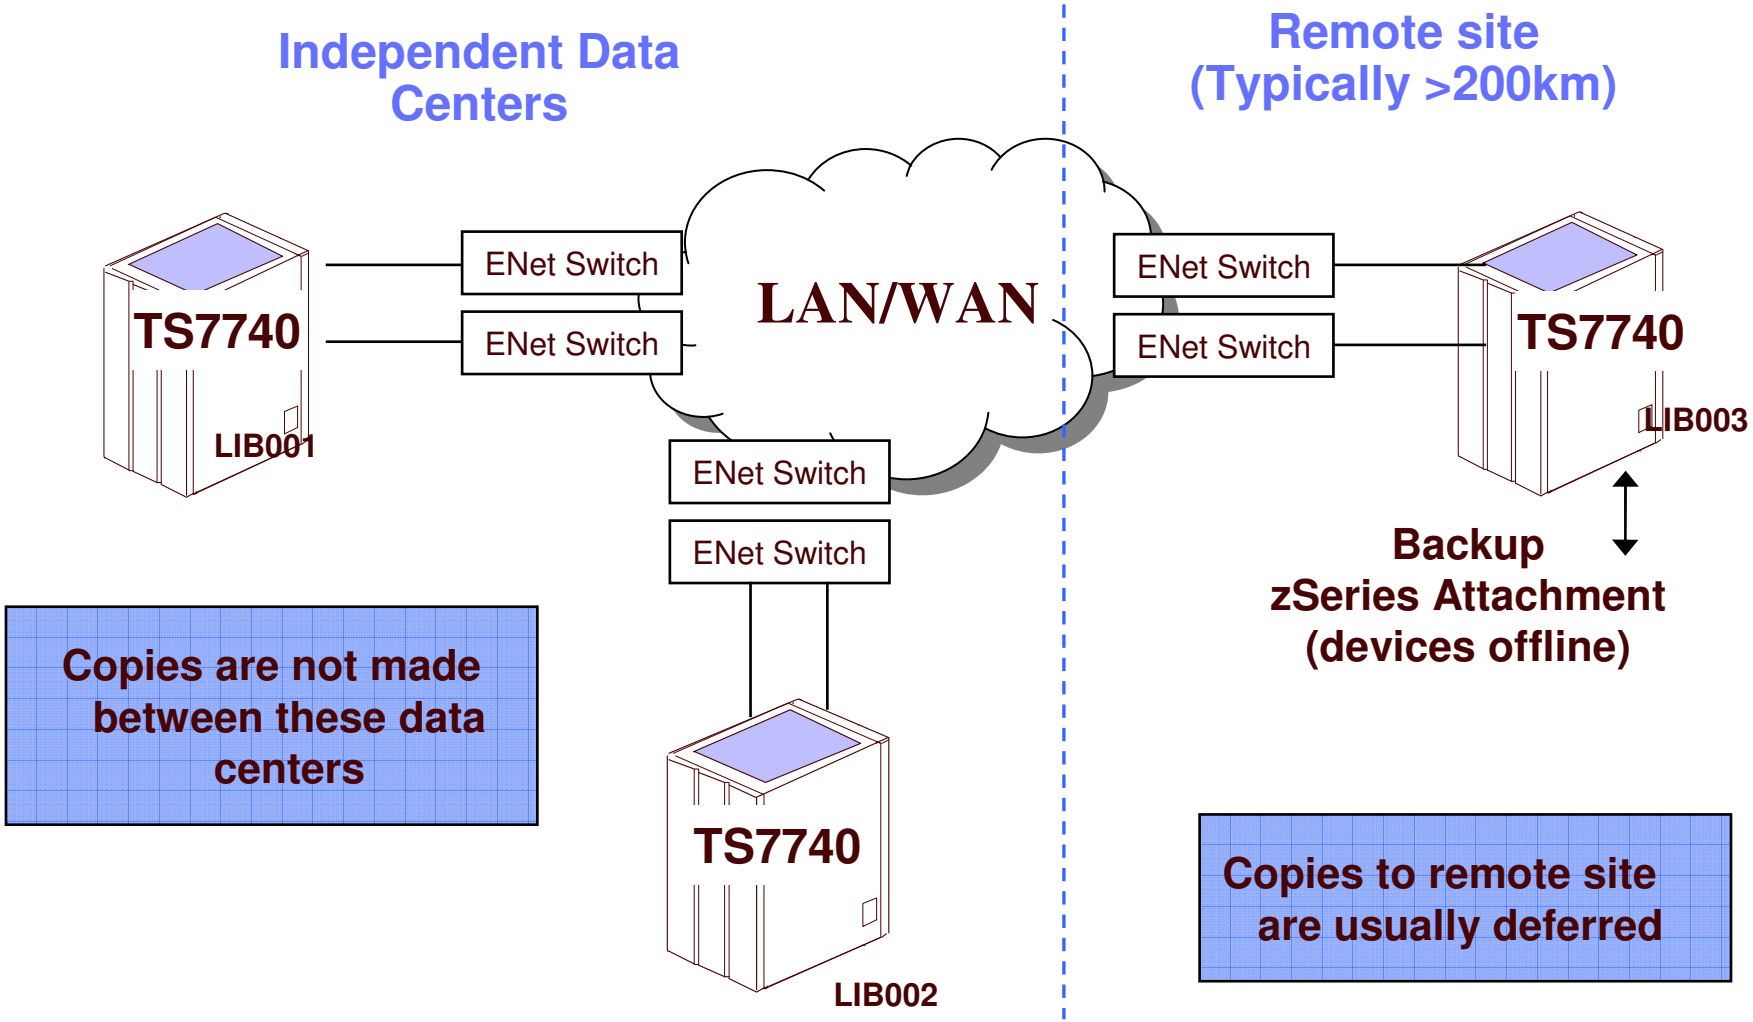

[1]List of potential enhancements is an estimate of current intent and is subject to change and is not a commitment to deliver.

TS7700 - Multi-Site Policy Based Copy Management

- Management Class Storage Construct
  - Where copies are to reside
    - By distributed library
  - When copies are to be consistent with host that created the data
    - At volume close time (Rewind/Unload: RUN)
    - After volume close time (Deferred)
    - No copy

Examples

Management Class: PROD01

LIB001	LIB002	LIB003
RUN	RUN	Deferred

Management Class: TEST01

LIB001	LIB002	LIB003
RUN	No Copy	No Copy

Copy Export for Disaster Recovery - Overview

- Export function to support transfer of data for offsite disaster recovery
- Exports a copy of selected data, leaving the primary copy in the TS7700
  - For selected data, customer specifies a management class that defines a secondary storage pool
    - Data with different locations or expiration cycles may be directed to different storage groups/secondary pools.
  - As logical volumes are copied from cache, they are written to their assigned primary and secondary storage pools.
  - During a copy export operation, all physical volumes of a specified secondary storage pool that contain active logical volumes are removed from the TS7700
- Customer performed recovery process
  - All copy exported data from a source TS7700 is recovered on an empty TS7700
    - May include multiple secondary pools
  - Recovery options for test vs an actual disaster recovery
- Copy exported physical volume continue to be managed by the source TS7700
  - Volume records maintained in database - support status queries
  - Active logical volume content updated as volumes change/expire
  - Manages offsite volume reclamation
  - Information about the exported physical volumes available through host console and BVIR request
- Operation in a Grid Configuration (supported in R1.4)
  - Executed on a specific TS7700
  - Logical volumes must have been copied to the TS7700 to be exported
  - Recovery is to a standalone TS7700

Copy Export Operation - Data Creation

Host creates application data

DSN = PROD.ACCOUNT.MAIN
DSN = PROD.ACCOUNT.MERGE
DSN = PROD.ACCOUNT.OUTPUT.Gu.Vu
. . .

ACS { SG = PROD01
MC = DRDATA
SC = CACHE
DC = TAPE

Storage Group PROD01 -> Pool 01

Management Class DRDATA -> Pool 09

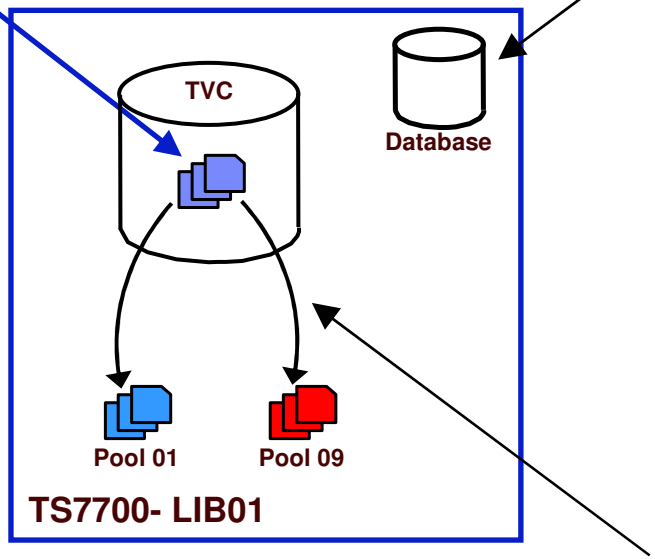

Management Class DRDATA is used to identify data to be taken offsite for disaster recovery; outboard definition at library manager assigns physical pool 09

Logical volume data is copied from cache to physical tapes in pool 09 as well as to pool 01 because of the management class assignment.

Copy Export Operation - Export Operation

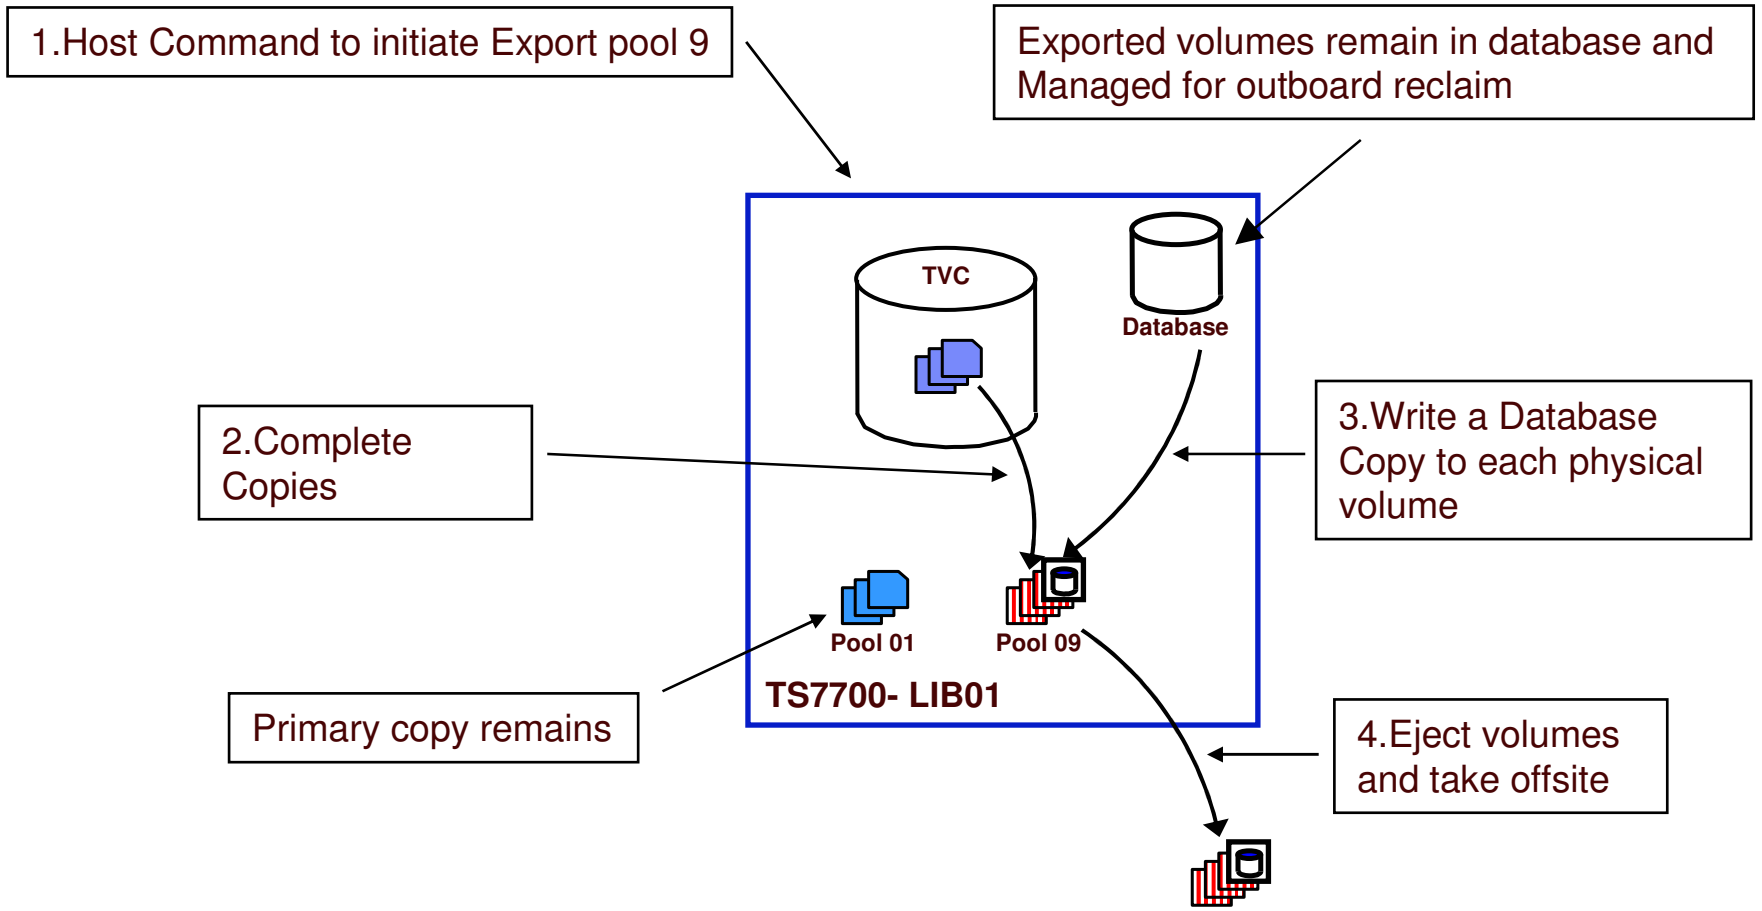

Copy Export Recovery

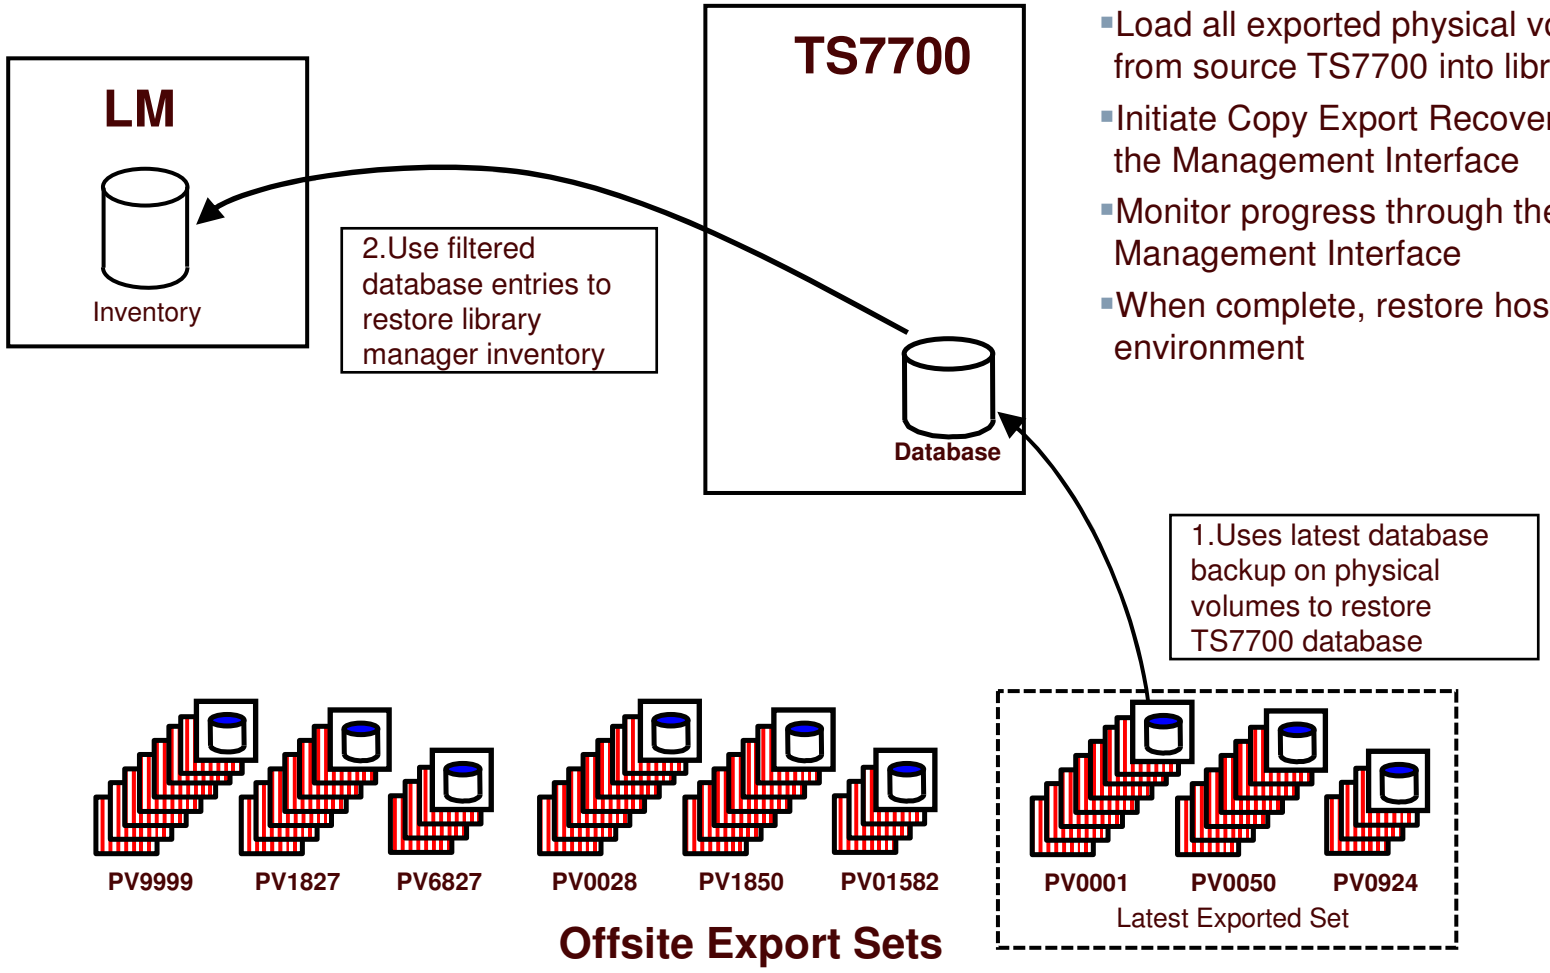

- Load all exported physical volumes from source TS7700 into library
- Initiate Copy Export Recovery from the Management Interface
- Monitor progress through the Management Interface
- When complete, restore host environment

z/OS Host Console Request Facility

- Address key customer requirements for their operators to perform problem determination from the z/OS console
- Host console command to request TS7700 functions
  - Similar to existing library query commands
  - Customer provides keywords based on request needed
- Response is a set of formatted text lines
  - Text lines formatted by the TS7700
  - z/OS sends them to the console as a multi-line response
- Host Command and Response Mechanism easily expanded
  - z/OS does not understand any of the keywords in the request
  - z/OS does not interpret the text lines
- Suitable for Automation
  - Publish format of returned data so that customers can use the information for automating tasks

zOS Host Console Request Command

- LIBRARY REQUEST,*libname*, k*eyword1,keyword2,keyword3,keyword4*
  - *libname* is the name of the composite or distributed library
  - Provide for 4 separate keywords
  - Keyword is 1-8 characters (A-Z, 0-9, $*@#%), no blanks
- DFSMS generate a console message when the request is initiated
  - CBR1020I Processing Library command: request,keyword1,keyword2,keyword3,keyword4
- DFSMS generates a multi-line WTO to the console
  - CBR1280I
  - Up to 50, 70 character width lines are returned
  - Formatted by the TS7700

Host Console Request

```
LIBRARY REQUEST,BARR50A,PDRIVE
CBR1020I PROCESSING LIBRARY COMMAND: REQUEST,BARR50A,PDRIVE.
CBR1280I LIBRARY BARR50A REQUEST. 301
KEYWORDS: PDRIVE
------------------------------------------------------------------------
PHYSICAL DRIVES V1
   SERIAL NUM     TYPE  MODE  AVAIL  ROLE  POOL    PVOL    LVOL
 000007875183 3592E05E  E05E      Y  RCLS    21  Q00888  ZL8507
 000007874867 3592E05E   E05      Y  RECA    01  Q00873  ZM8791
 000007875179 3592E05E   E05      Y  IDLE    01  Q00860
 000007874935 3592E05E  E05E      Y  RCLT    21  Q00904  ZL8507
 000007874798 3592E05E   E05      Y  IDLE    01  Q00900
 000007875018 3592E05E  E05E      Y  IDLE    02  JA7618
IFU008I - SETUP JOB - TAB52202 - SUCCESSFUL. OUTPUT DELETED
$PJ07002
$HASP395 TAB52202 ENDED
IEC501A M 82B7,PRIVAT,SL,NOCOMP,TAB5220C,CREATE1,TAPE.TAB5220C
IEC501A M 82D8,PRIVAT,SL,NOCOMP,TAB52208,CREATE1,TAPE.TAB52208
```

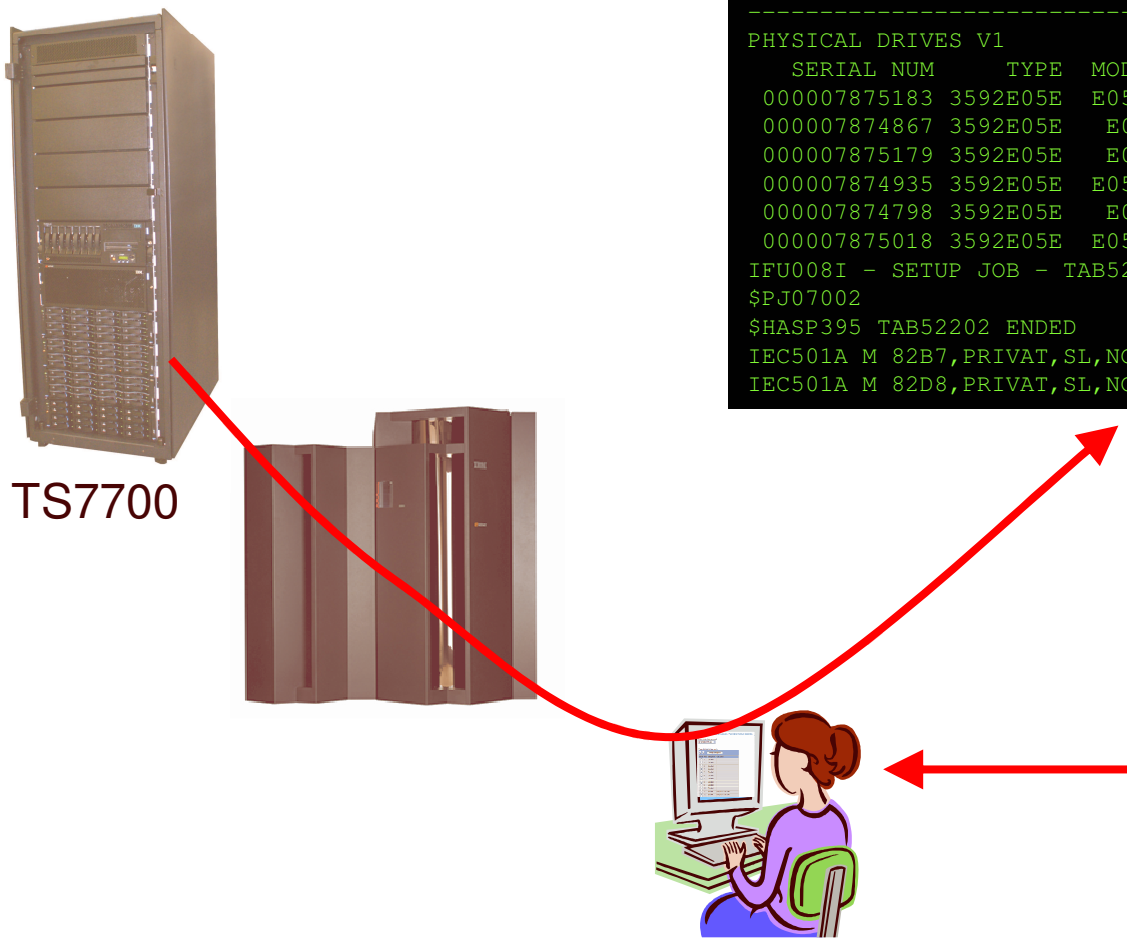

- Requests Supported
  - Cache Status
  - Logical Volume Information
  - Physical Volume Information
  - Physical Drive Activity
  - Physical Volume Counts
  - Recall Queue
  - Grid Status
  - Copy Export Volume Reclaim
  - Copy Export Volume Delete

z/OS Host Console Request Keywords

Keyword 1	Keyword 2	Keyword 3	Description	Comp	Dist
CACHE			**Requests information about the current state of the cache and the data managed within it associated with the specified distributed library.**	**N/A**	**Y**
COPYEXP	**zzzzzz**	**RECLAIM**	**Requests that the specified physical volume that has been exported previously in a copy export operation, be made eligible for priority reclaim.**	**N/A**	**Y**
COPYEXP	**Zzzzzz**	**DELETE**	**Requests that the specified physical volume that has been exported previously in a copy export operation, be deleted from the TS7700 database. The volume must be empty.**	**N/A**	**Y**
LVOL	**zzzzzz**		**Requests information about a specific logical volume.**	**Y**	**N/A**
PDRIVE			**Requests information about the physical drives and their current usage associated with the specified distributed library.**	**N/A**	**Y**
POOLCNT	**[0-32]**		**Requests information about the media types and counts, associated with a specified distributed library, for volume pools beginning with the value in keyword 2.**	**N/A**	**Y**
PVOL	**zzzzzz**		**Requests information about a specific physical volume.**	**N/A**	**Y**
RECALLQ	**[zzzzzz]**		**Requests the content of the recall queue starting with the specified logical volume.**	**N/A**	**Y**
RECALLQ	**zzzzzz**	**PROMOTE**	**Requests that the specified logical volume be promoted to the front of the recall queue, then returns the content of the recall queue.**	**N/A**	**Y**
STATUS	**GRID**		**Requests information about the copy, reconcile and ownership takeover status of the libraries in a Grid configuration**	**Y**	**N/A**

1 M Logical Volumes

- TS7700 will now support up to 1M logical volumes
- If two TS7700s are attached to a 3953, each TS7700 supports the full 1M
  - Total of 2M for the 3953
- Support on 3494 depends on the hardware level of the library manager
  - Must at least have the 1.2Mhz platform

TS7700 White Papers

- IBM Virtualization Engine TS7700 Series Copy Export Function User's Guide Version 1.0
- IBM Virtualization Engine TS7700 Series z/OS Host Command Line Request User's Guide Version 1.0
- IBM Virtualization Engine TS7700 Series Bulk Volume Information Retrieval Function User's Guide Version 1.3

- Available at G.A. on techdocs

Other tape news

New IBM Tape Systems Portfolio (Apr '07)

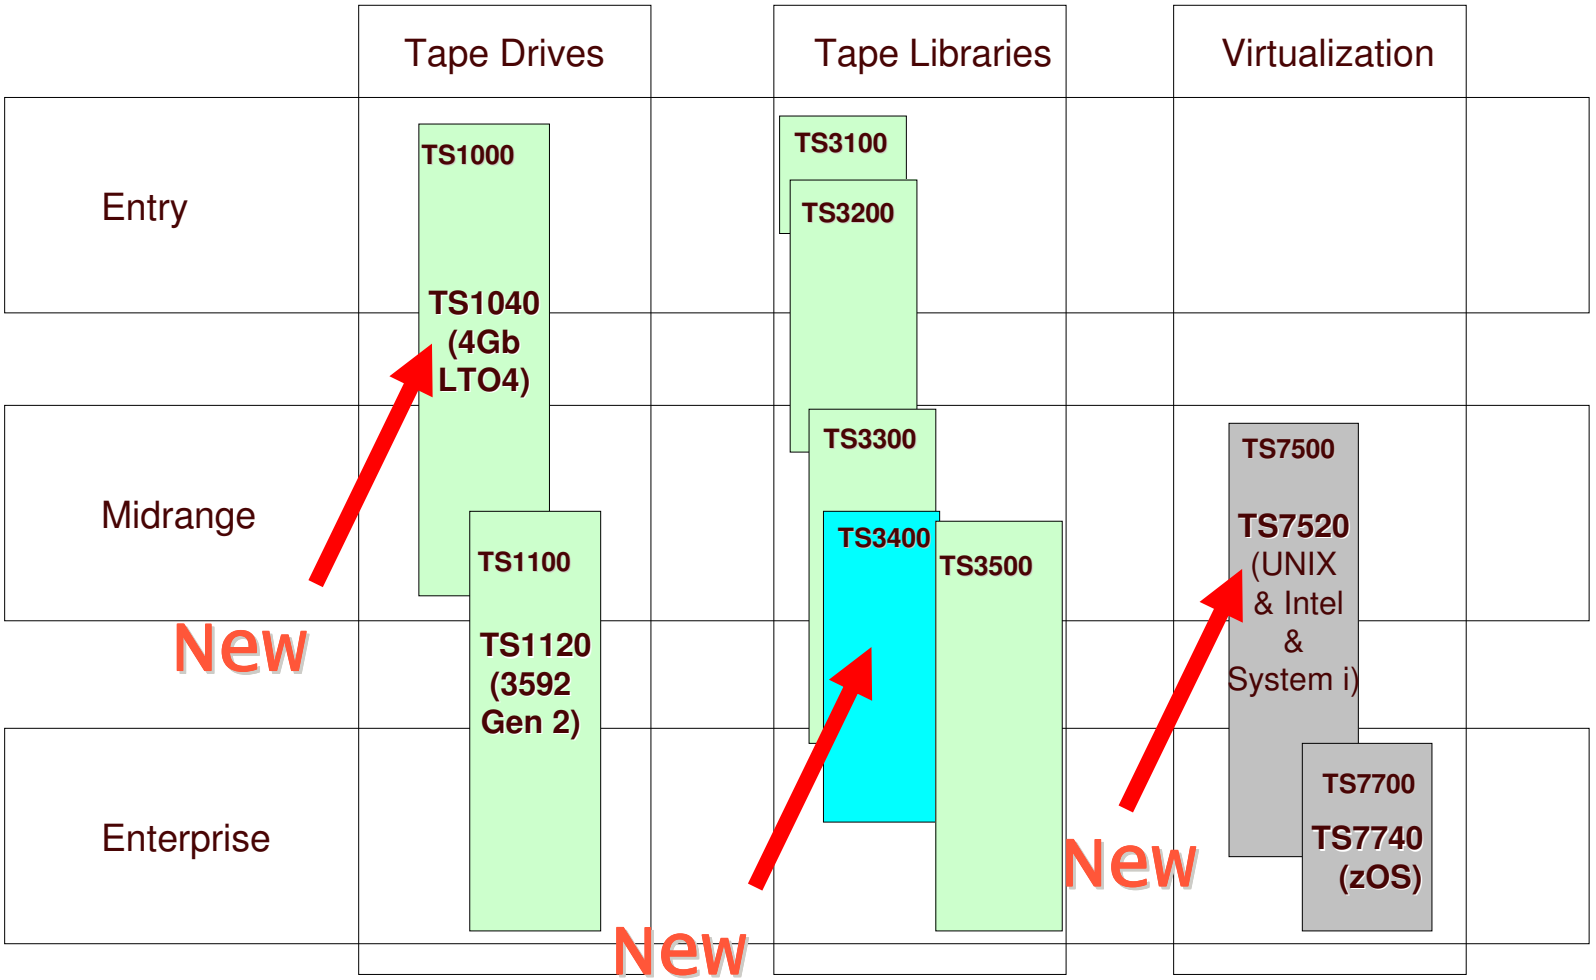

TS3400 Tape Library Overview

- 1 to 2 TS1120[1] tape drives
  - 4Gbps dual port Fibre Channel attachment
- 2 removable cartridge magazines
  - Each holds up to 9 cartridges
    - Front three slots of the lower magazine can be configured as I/O station slots
    - Two slots in the upper magazine, if installed, can be configured as cleaning slots
- Bar code reader standard
- Ability to partition the library into two logical libraries
  - Each logical library comprised of one drive and one magazine
  - Run any single-drive logical library in either sequential (autoloader) mode or random (library) mode
- Manageable by local operator panel or remote web GUI
- Storage capacity of up to 12.6TB (up to 37.8TB with 3:1 compression)
- Stand alone or rack mount configurations

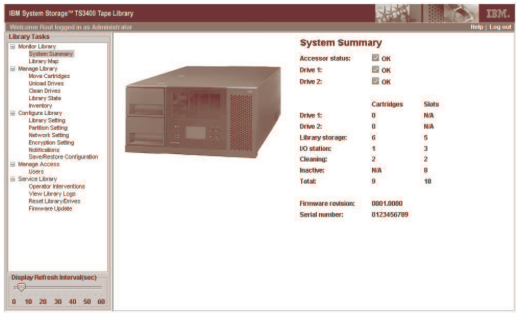

[1]3592 J1A drives are not supported

System z Support

- Requires TS1120 tape controller
  - –Other required features:

MT/Model	Feature Code	Description
3592 C06	9014	Attach to TS3400
3592 C06	4641	Rack mount kit
3592 C06	5247	Enhanced router
3577 L5U	9014	Attach to System z
3577 L5U	7004	Rack mount kit

- Two modes of operation for System z hosts
  - –System mode (z/OS only)
  - –Auto mode
  - –No random access of cartridges

Why virtualize the open system tape process?

- Virtual tape can be employed to help
  1. Improve the backup process
     - Address network limitations
     - Manage immense data growth
  2. Improve the recovery time objective
     - Eliminate physical tape movement
     - Reduce contention for resources
  3. Improve the recovery point objective
     - By creating incremental backups more frequently
     - By writing incremental backups to cache
  4. Augment the existing business continuance infrastructure
     - Reduce bottlenecks and utilize tape assets more efficiently
     - Optimize IT resources
  5. Optimize the IT business continuance infrastructure

Virtual Tape Concepts

- Virtual Tape Drives
  - Appear as multiple LTO or 3592 tape drives
  - May reduce the number of real tape drives required
  - Can be shared / partitioned like real tape drives
- Tape Volume Cache
  - Eliminates mechanical tape mount and thread operations
  - Allows creation of virtual volumes on RAID5 disk cache buffer
  - May help improve tape performance / reduce backup window
- Tape drive attachment
  - Allows volumes to be moved/copied to real tape drives by
    - The backup application (such as Tivoli Storage Manager)
    - The TS7520 Virtualization Engine (using cache management routines)
  - Supports Business Continuity and Information Lifecycle Management

Virtual Tape Concepts (continued)

Physical Topology

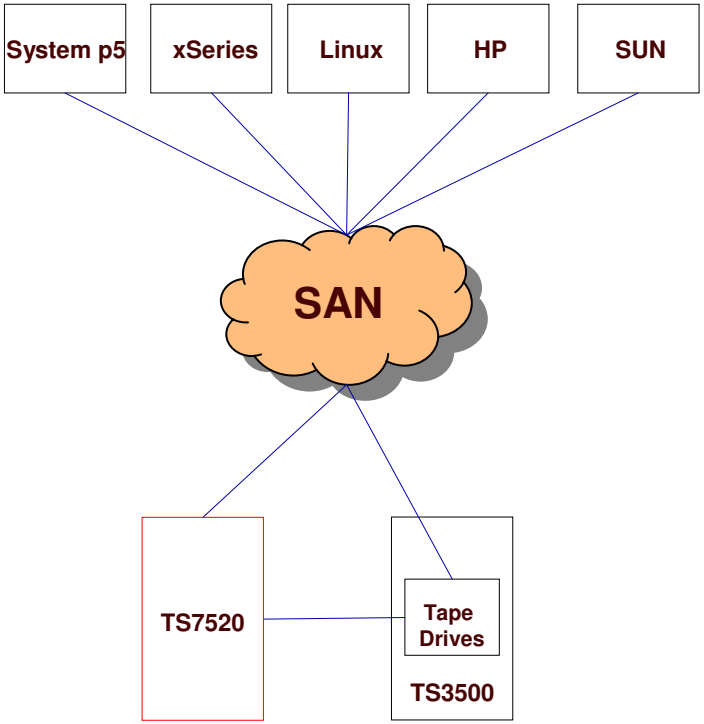

Logical Topology

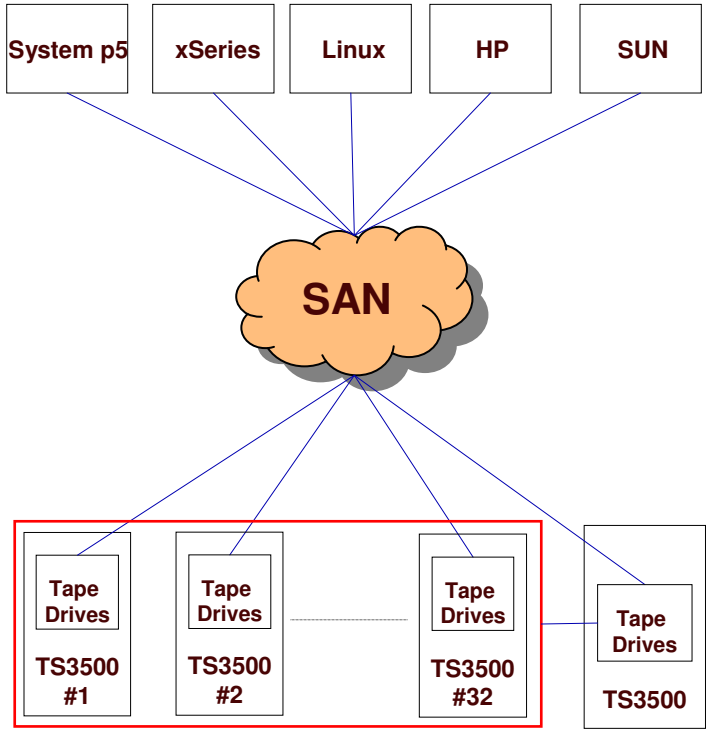

Improving the Backup Process

- Consider a simplified hypothetical backup scenario that includes
  - Eight backup tasks that can exploit native tape
  - Twenty backup tasks that cannot exploit native tape

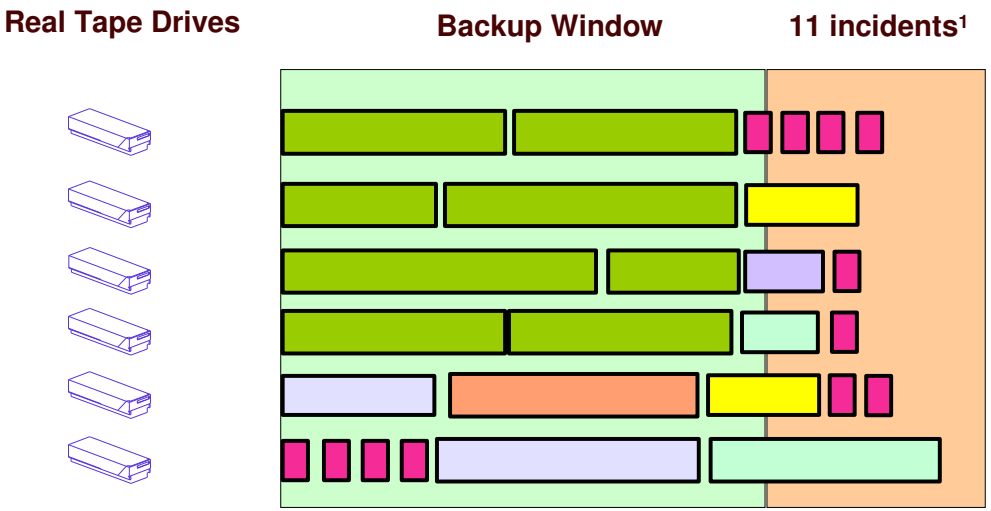

[1] Incident refers to a Service Level Incident and not a tape outage

Improving the Backup Process (continued)

- By adding more tape drives
  - The issues listed on the previous page may be resolved
  - But at a cost (Purchase, Maintenance, Management)

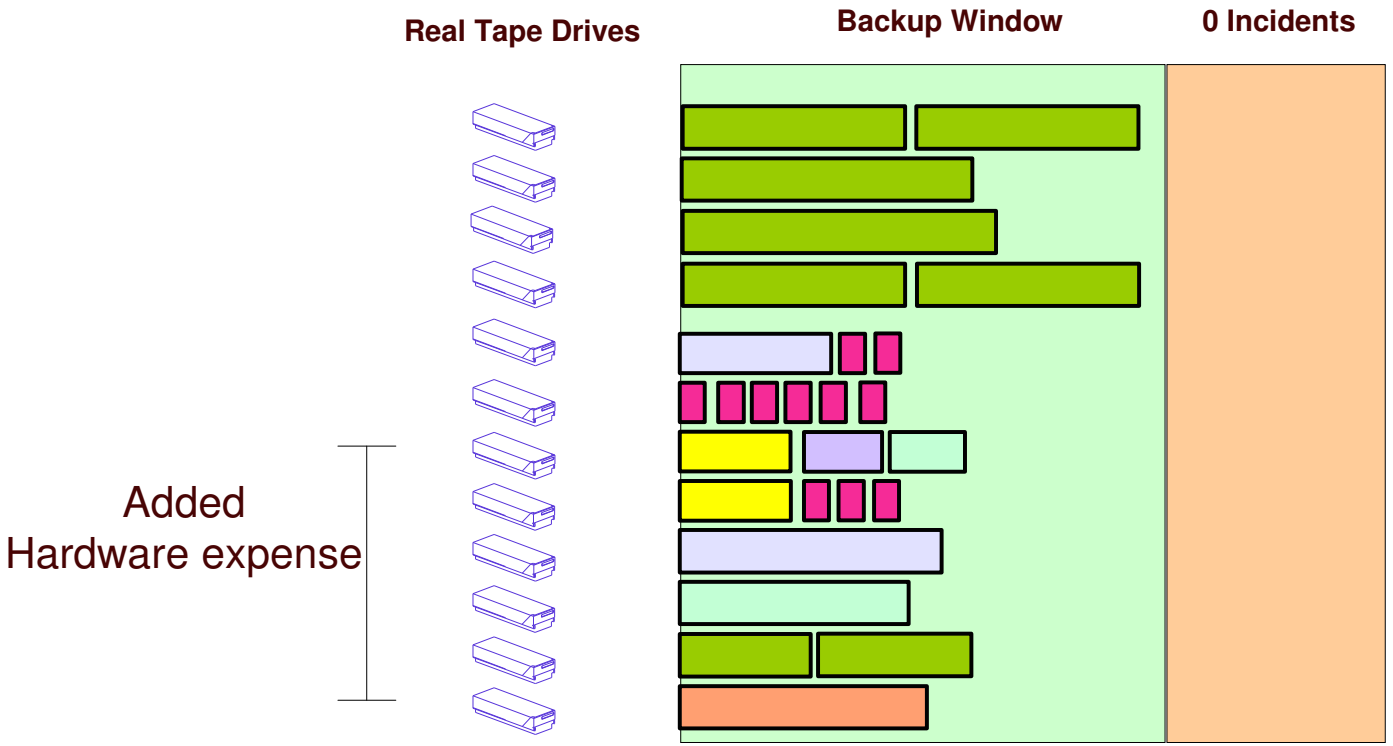

Improving the Backup Process (continued)

- By implementing a TS7520 Virtualization Engine
  - The issues listed on the previous pages may be resolved
  - Existing tape drives may be re-deployed support creation of offsite copies

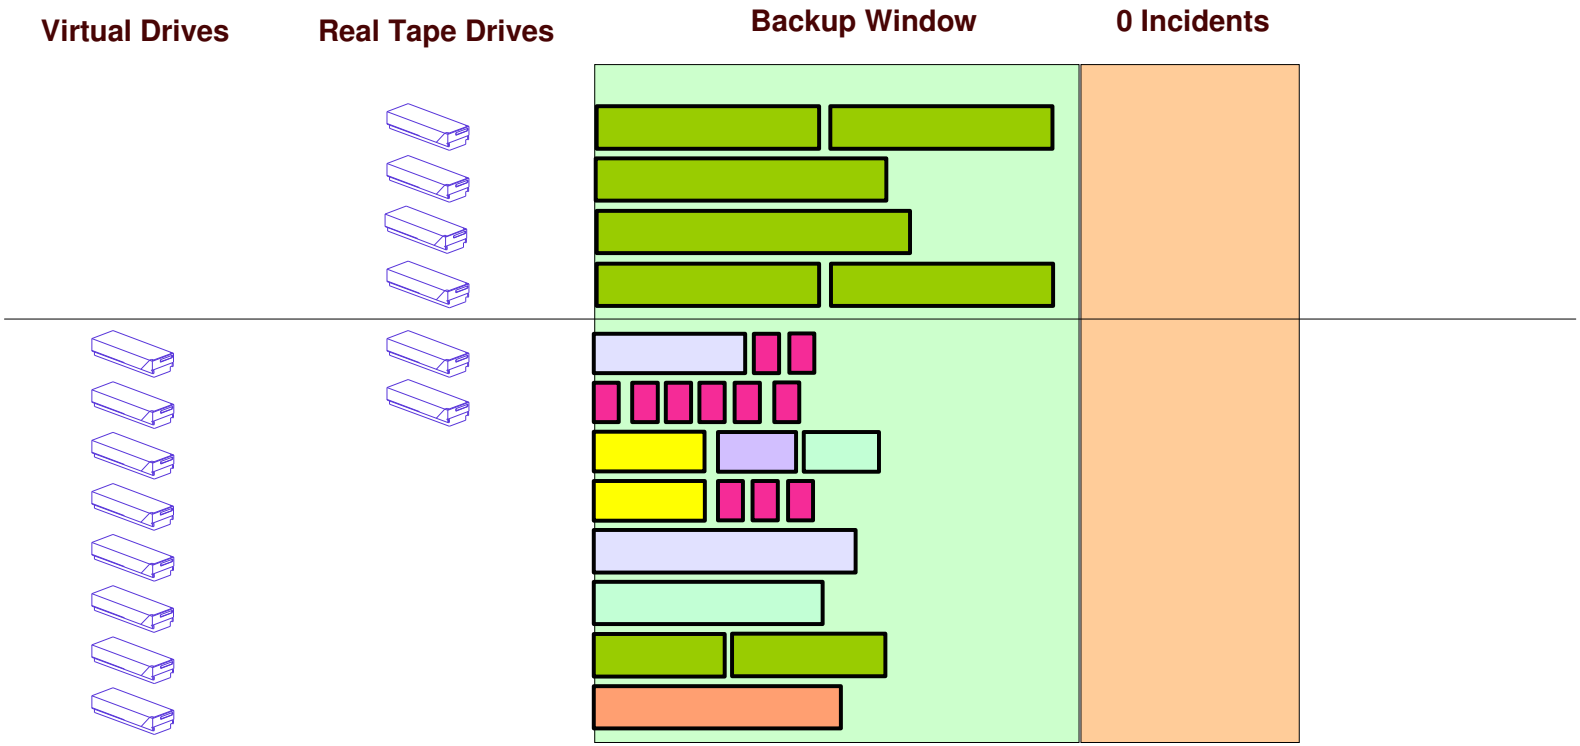

Traditional Backup

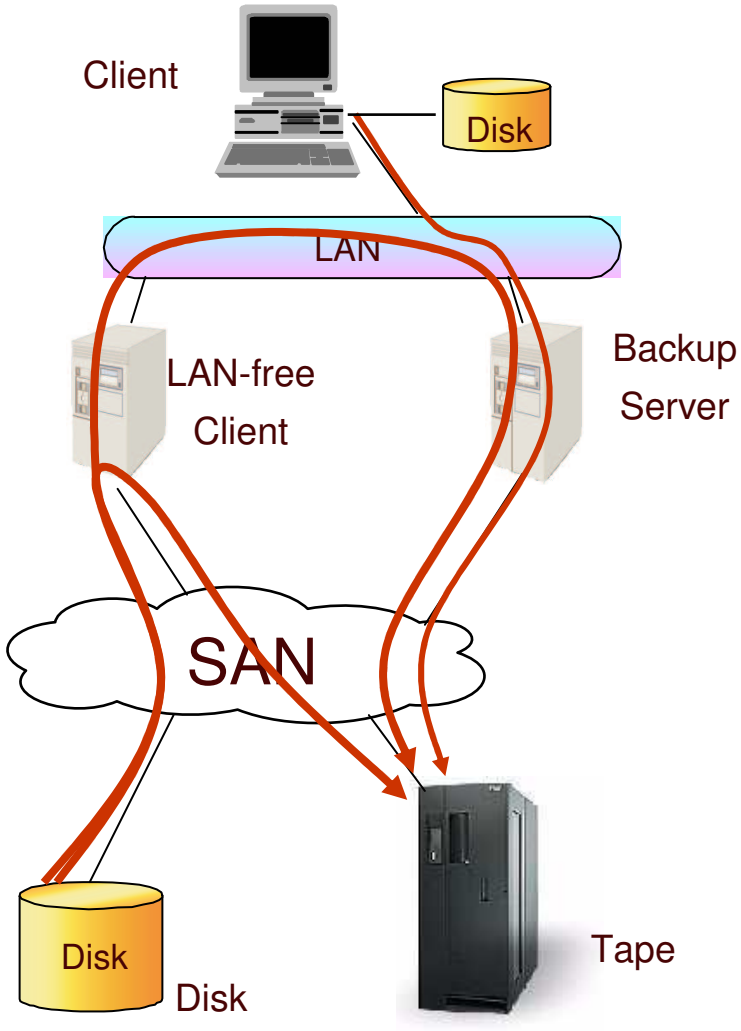

- Backup data written via LAN to Backup Server and tape
  - Issues
    - LAN utilization
    - Components throttling drive
    - LTO3 = 80 MB/sec / 3592 = 104 MB/sec
    - 100 Mbit Ethernet = 8 MB/sec
    - No advantage from fast tape drives
    - Migration window consideration
- LAN-free approach writes Backup data directly to tape
  - Issues
    - Dedicated tape drives required
    - Large numbers of drives required to meet backup window
    - Time to restore client data

The Solution for LAN-free (and other Backup S/W without a Disk Buffer)

Disk appears as tape

No application changes required

No additional software

Reduces real drive quantity

Addresses backup/restore issues

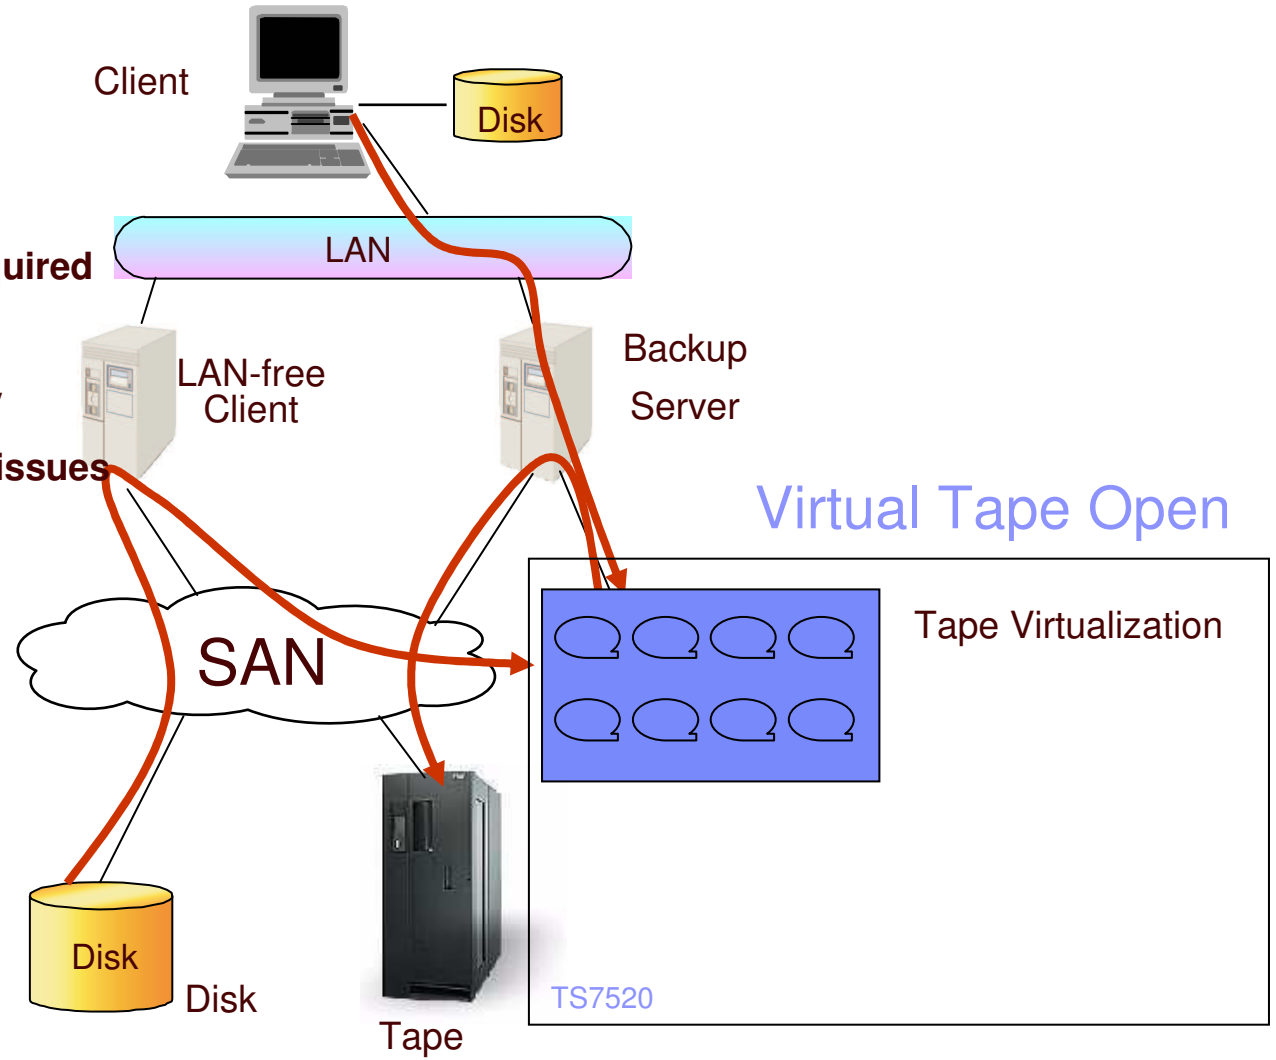

IBM TS7520 Virtualization Engine Overview

- **New Virtualization Engine for tape**
  - Up to 4.8GB (4,800 MB) per sec and 884TB native capacity
  - Up to 512 virtual tape libraries, 4,096 virtual drives and 256,000 virtual volumes
- **Supports**
  - Supports all current IBM tape library models
  - Supports all current IBM tape drive models[1]
  - Optional features include network replication, hardware compression, encryption for data in flight and on tape
- **Attaches to**
  - IBM System p, System z (Linux) and System x
  - Selected HP and Sun Microsystems servers
  - Servers running supported Microsoft Windows™ and Linux operating systems

[1] LTO Gen 4 tape drive support is planned for 4Q07. This represents a statement of IBM's current plans and direction, such plans and direction are subject to change without notice.

VTL Benefits in a TSM Environment

- Good for LAN-free Backup. If you have many drives (virtual or not) then you can do more in parallel and LAN-free to tape is more easily and commonly deployed than LAN-free to disk. Virtual Drives may be more cost-efffective than physical drives (depending on the numbers required).
- Migration: typically from disk to tape. VTL is similar to setting up a huge disk pool with Global Mirror (though may be a similar cost, while offering more benefits).
- VTL remote replication may also offer benefits.
- Reclamation will still have to happen. Logical tapes get gaps where data expires just like physical tapes.
- Collocation: a VTL may not have the wait time of physical robotics, but restores will still wait if another process is using the virtual tape volume that they need so the benefits of collocation still stand.

Integrated Removable Media Manager for the Enterprise (iRMM)

Today‘s tape and media management challenges

- **Customers increasingly have a need for consolidated tape and media management solutions**
  - Hardly any tools for open systems tape media management beyond what’s built in to backup/archive applications
    - For example, cross platform serialization of cartridges and drives
    - And enterprise view of tape assets
  - z/OS® and open systems attached tapes are managed separately - two systems for doing the exact same thing
  - Duplicate administration effort can severely complicate management of steadily growing tape environments.
- **Hardly any dynamic resource sharing**
  - Not among open systems and not among z/OS and open systems
  - Dedicated hardware resources have to be managed separately
  - Inefficient utilization of tape resources

IRMM in a Nutshell

- **Integrated Removable Media Manager (IRMM) is:**
  - A new robust systems management product for Linux® on IBM System z™ that manages open system media in heterogeneous distributed environments and virtualizes physical tape libraries, thus combining the capacity of multiple heterogeneous libraries into a single reservoir of tape storage that can be managed from a central point
- **IRMM is designed to provide:**
  - Centralized media and device management
  - Dynamic resource sharing
- **IRMM extends IBM's virtualization strategy to tape library resources (drives and cartridge pools)**
- **IRMM complements Linux on System z consolidation efforts**

Tape Library Virtualization and Enhanced Tape Resource Management

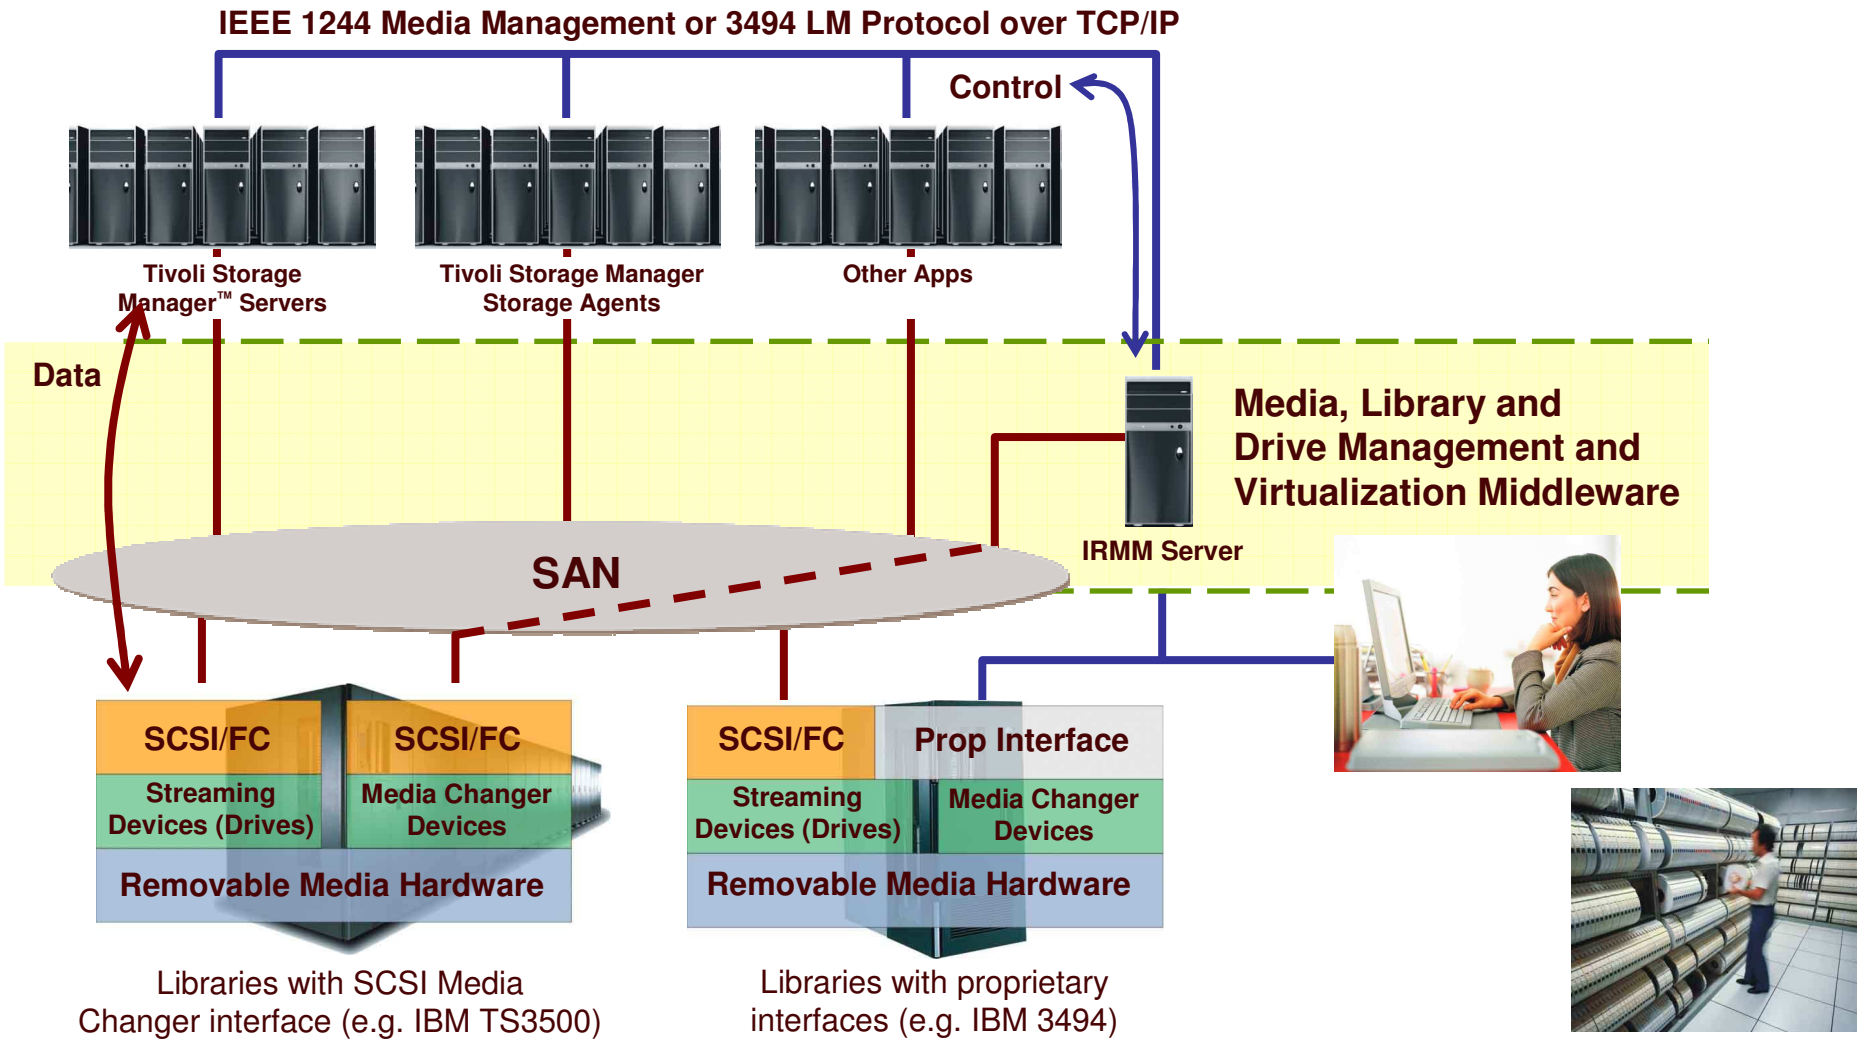

Tape Library Virtualization and Enhanced Tape Resource Management

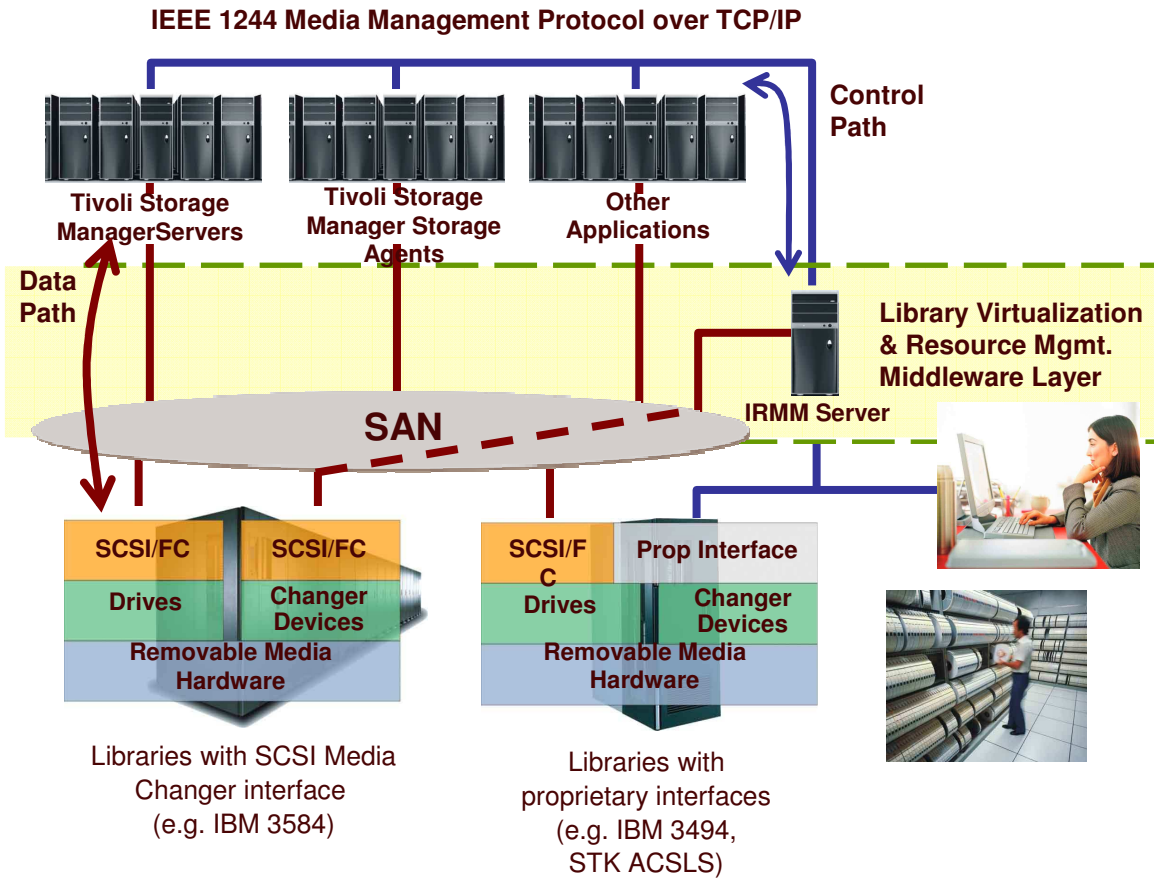

IRMM systems management for open systems is designed to provide:

- **Tape Library Virtualization**
  - Similar to what SAN Volume Controller provides for disk
  - Which combines the capacity from multiple heterogeneous libraries into a single reservoir of tape storage which can be managed from a central point
  - Which may increase utilization, availability, administrator productivity and efficiency
  - Which may reduce downtime for planned and unplanned outages and maintenance
  - Which may reduce administration of Tivoli® Storage Manager (and other BA products) through “on demand” provisioning of devices/libraries
- **Tape Storage/Resource Management**
  - Includes mainframe-class management of removable media for Open Systems
  - Enables single point of control for all removable media management tasks through integration with z/OS DFSMSrmm™
  - Advanced tape hardware management, monitoring and reporting
  - Policy-based administration tasks

IRMM Features

- **Consolidated, mainframe-class media management services**
- **Centralized repository, access control and administration**
- **Centralized reporting and monitoring**
- **Management beyond physical library boundaries**
  - Access a cluster of multiple physical libraries as one large logical library
- **Dynamic sharing of resources across heterogeneous application boundaries and heterogeneous operating systems**
  - Helps sharing of tape resources much easier
- **Security features to permit or prevent application access to tapes**
  - Helps to enable common scratch pool and private pools for every application
  - Ensures secure usage and visibility
- **“Virtualizes” the media changer Interface**
  - Can help to allow selection of best-of-breed software and hardware
- **“Decouples” volumes and cartridge**
  - Prepared to support other removable media
  - Supports application transparent movement/copying of volumes to other cartridges

IRMM Features (continued)

- **Policy-based drive and cartridge allocation**
- **Policy-based media-lifecycle management**
- **DFSMSrmm integration**
  - Consolidated management of tape resources (mainframe and open systems)
- **z/OS integration**
  - IRMM command line interface available for z/OS
- **Tivoli Storage Manager enhancements**
  - Dismounts all tapes currently mounted when Tivoli Storage Manager stops
  - Releases tapes to scratch if open is not successful

Integrating IRMM and DFSMSrmm

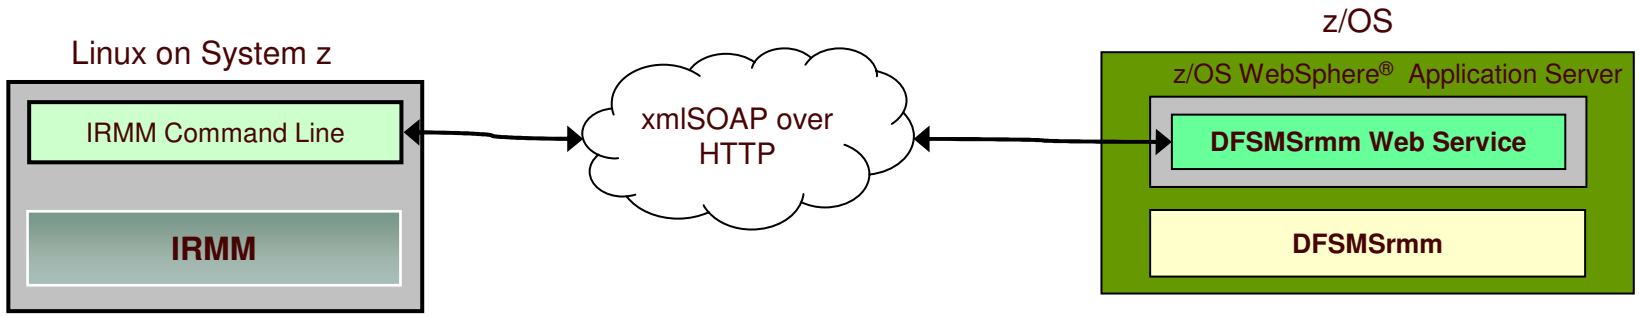


```
ermmtool> lsrmmvol -s 9.11.214.154
Volume   Type     Owner    State    HLOC     Voltype
=========================================================
A00001   ETC      CHKEIL   MASTER   SHELF    PHYSICAL
A00003   ETC      ERMM     USER     SHELF    PHYSICAL
483AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
485AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
486AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
488AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
490AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
491AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
495AFQ   ETC      ERMM     USER     SHELF    PHYSICAL
```

```
                DFSMSrmm Volumes (Page 1 of 2)                Row 1 to 9 of 9
Command ===> _____________________________________________ Scroll ===> PAGE

Enter HELP or PF1 for the list of available line commands
Use the RIGHT command to view other data columns
   Volume          Assigned   Expiration S                    Dest-    Tr- Data
S  serial Owner    date       date       R Status  Location  ination  ans sets
-- ------ -------- ---------- ---------- - ------- --------  -------- --- -----
__ A00001 CHKEIL   2006/116   2006/121     MASTER  SHELF               N   0
__ A00003 ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 483AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 485AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 486AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 488AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 490AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 491AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
__ 495AFQ ERMM     2006/116   2006/121     USER    SHELF               N   0
******************************* Bottom of data ********************************
```

- **Consolidated view on tape resources**
- **Open systems interface to DFSMSrmm**

- **Consolidated view on tape resources**
- **Centralized retention, movement and vaulting capabilities**
- **Policy-based management of all cartridges (z/OS and open)**

Integrating IRMM and DFSMSrmm (continued)

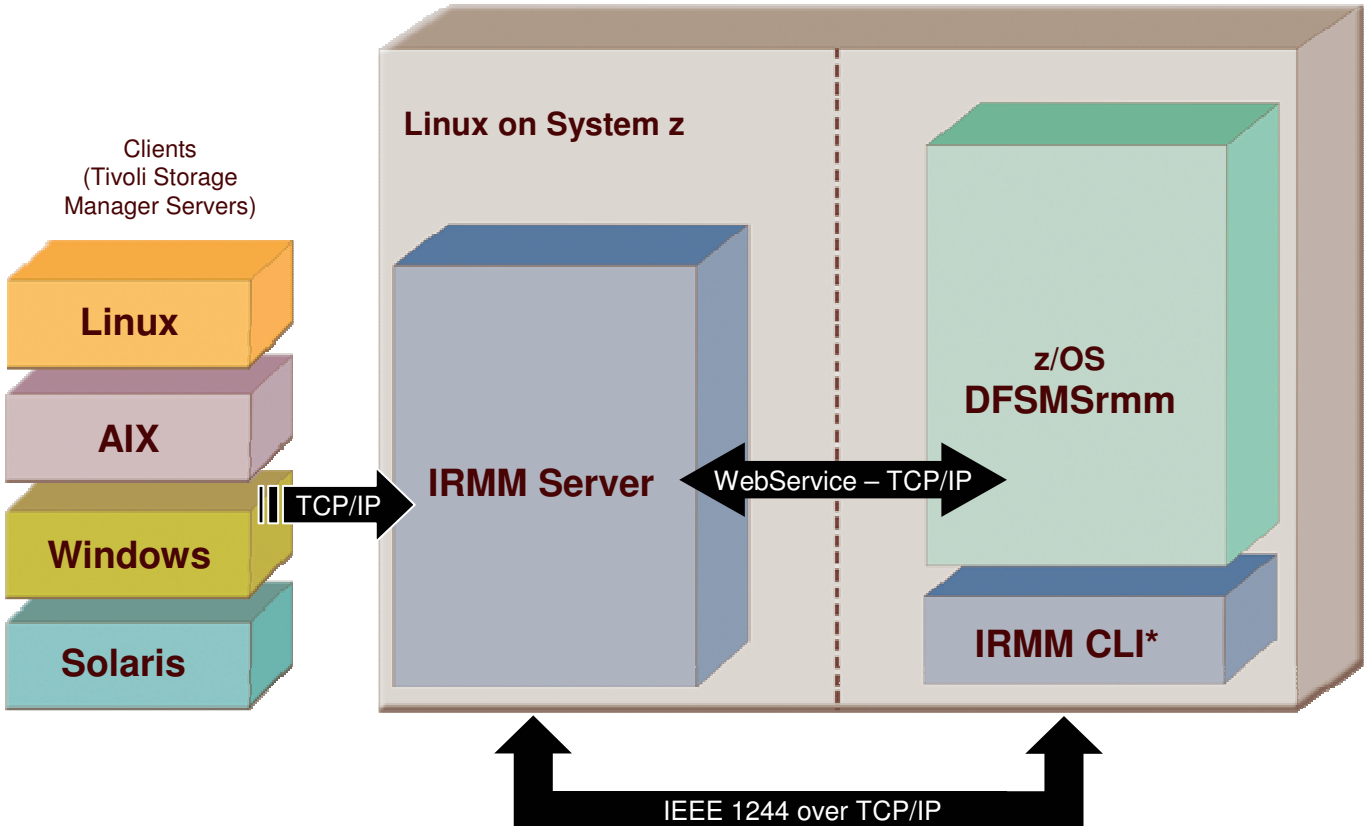

* Can run on Linux on System z as well

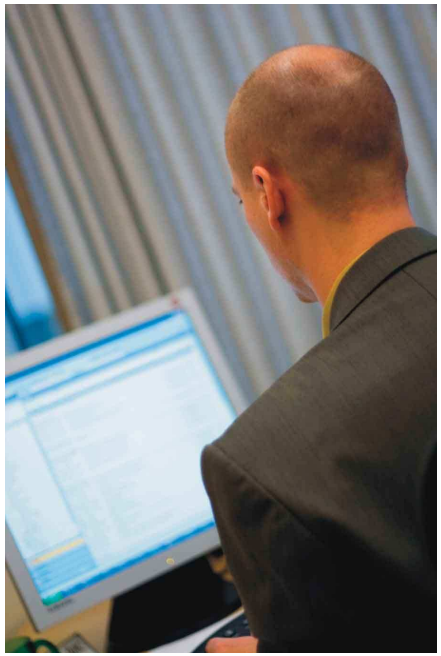

Operator can use z/OS to manage both z/OS and distributed environment!

Tape Library Virtualization – One Pool of Tape Storage

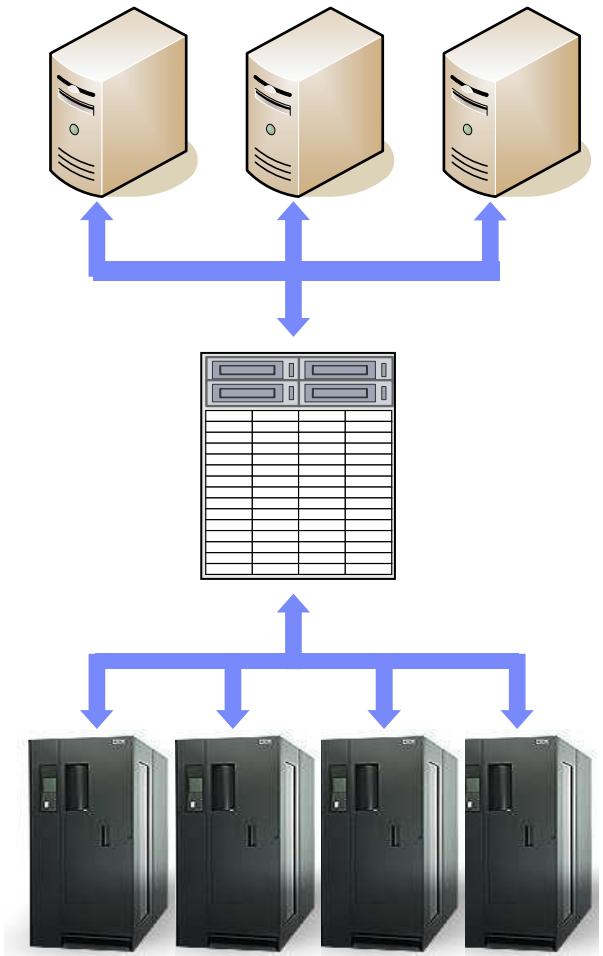

- Logical pool of tape storage beyond physical library boundaries
- Applications only have to deal with one library
- Physical storage capacity can be added without changes to the application
- Scratch mounts can still be possible as long as one physical library is online
- IRMM can provide workload balancing by distributing scratch mount requests
- Preferably done with homogeneous physical tape drives / media

That’s it, folks!