Simplify and Improve DB2 Administration by Leveraging Your Storage System

Ron Haupert
Rocket Software, Inc.

August 5, 2010
Session Number 7961

Session Agenda

- Database and Storage Integration Overview
- DB2 System-Level Backup Methodologies and Storage Integration Considerations
- Cloning DB2 Systems Using Storage-Based Fast Replication
- Refreshing DB2 Table and Index Spaces by Leveraging Your Storage Facilities
- Storage-Aware DB2 Product Examples
- Implementation Planning Considerations
- Session Summarization

Database and Storage Administration
Trends and Directions

- Large DB2 systems require high availability
  - Fast and non-intrusive backup and cloning facilities are required
  - Fast recovery capabilities are required to minimize downtime and promote high availability
  - Most backup, recovery and cloning solutions do not leverage storage-based fast-replication facilities
- Storage-based fast-replication facilities are under-utilized
  - Tend to be used by storage organizations
  - Tend not to be used by database administrators (DBAs)
- Storage aware database products allow DBAs to use fast-replication in a safe and transparent manner
  - Provides fast and non-intrusive backup and cloning operations
  - Simplifies recovery operations and reduces recovery time
  - Simplifies disaster recovery procedures

Database and Storage Integration

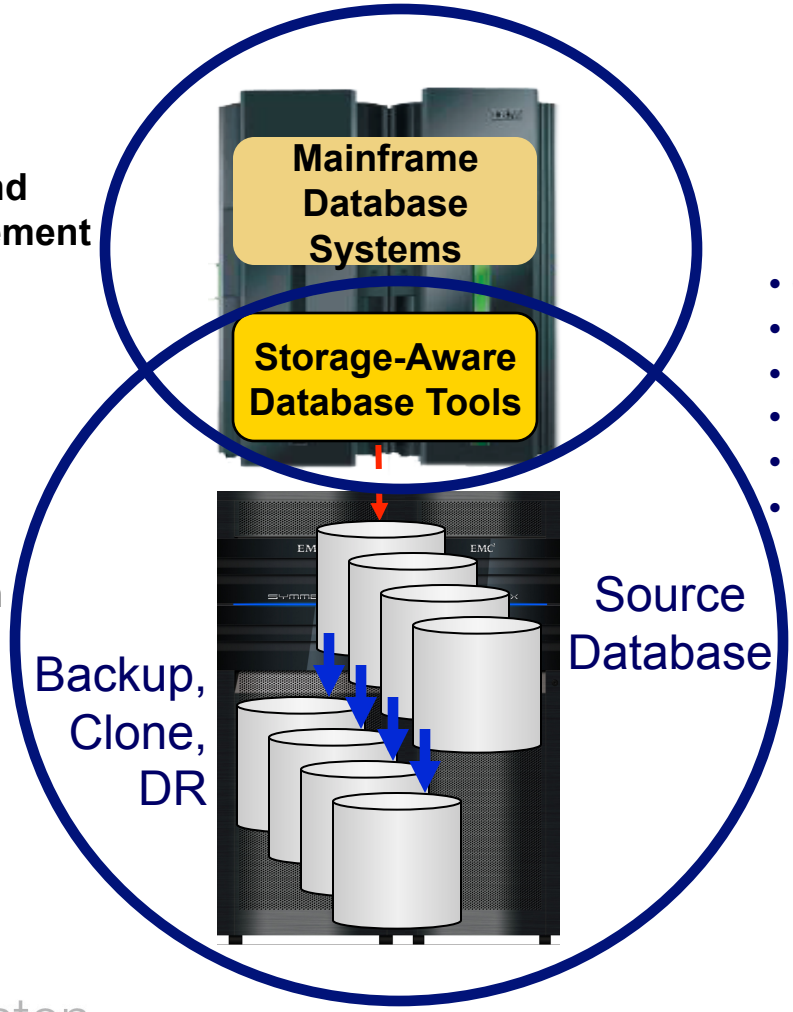

- Organizational Integration
- New Backup Methods
- New Recovery Strategies
- Business Recovery Monitoring
- Cloning Automation
- Disaster Restart Solutions

Database and Storage Integration
Operational Advantages

- Reduce backup, recovery, and cloning administration costs
- Reduce host CPU and I/O resource utilization
- Perform backups and create clone copies instantly
- Fast restore and parallel recovery reduces recovery time
- Simplify disaster recovery operations and procedures
- DBMS and storage-based fast-replication integration
  - Leverage storage processors and fast-replication investments
    - IBM, EMC, HDS, STK
  - Expose fast-replication capabilities to the DBAs ***safely and transparently*** using "***storage-aware***" database utilities
- Provide a sophisticated infrastructure and metadata to manage the DBMS and storage processor coordination

Database and Storage Integration
New Solutions for DBAs to Consider

- DBAs use traditional database backup and recovery tools
  - Difficult to integrate new backup and recovery methodologies
  - Uncomfortable with new backup and recovery solutions
- Lack of database and storage administration coordination
  - Storage processor fast-replication facilities are not well understood by application and DBA personnel
    - *DBAs don't trust storage technologies*
  - Database applicability of fast-replication not well understood by storage administrators
    - *Storage groups don't trust DBAs*
- Storage-aware database utilities resolve these issues

System Level Backup Methodologies

- Backup complete database systems as a unit without affecting running applications
  - Backup components include:
    - Active and archive logs
    - Recovery metadata
    - All database data sets
    - Appropriate libraries, and system data sets
    - All associated ICF User catalogs
  - Backups performed instantly using storage-based fast replication
- System-level backups are the foundation for federated backup and recovery solutions
- System backup and cloning methodologies are difficult to implement without sophisticated automation
    - “Split mirror” backup methodologies pioneered in late 1990s
    - Valuable concept - but hard to implement

Storage-aware DB2 Backup
Functional Requirements

- Integrate DB2 backup, restore, and recovery process with storage-based fast replication
- Provide easy and fast backup and restore of DB2 systems and applications
- Support common storage systems
  - IBM – FlashCopy (FC)
  - EMC – TimeFinder/Mirror/Clone/Snap, FC
  - HDS – Shadow Image, FC
- Feature requirements include:
  - DB2 system discovery and configuration management
  - DB2 system backup and recovery operations
  - System backup validation
  - Application and object data set backup
  - Image copy creation
  - Object and application recovery
  - Active metadata repository
  - Encrypted tape offload support
  - DR preparation and management

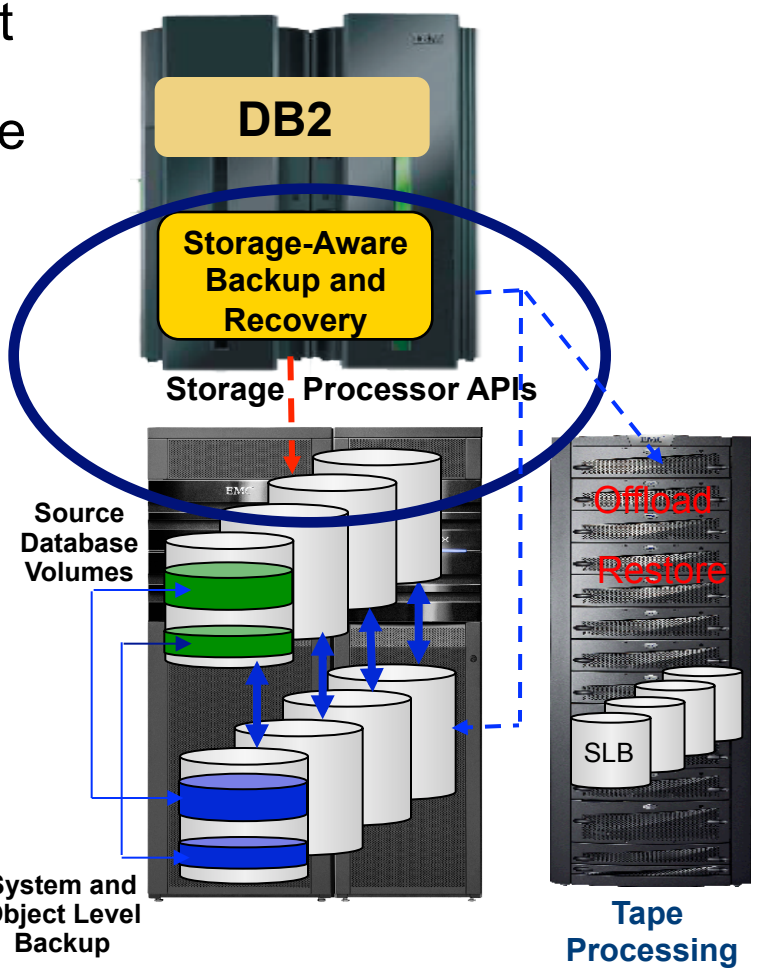

DB2 System Level Backup

- Storage-based backup reduces processing and administration costs
- Fast replication is used to perform database backup and restore functions
  - Full system backups complete in seconds
  - Backup performed without host CPU or I/O
- Back up large groups of databases with no application affect or down time
  - Backup windows are reduced or eliminated
  - Extend online or batch processing windows
- Data consistency ensured
  - Database suspend process
  - Storage-based consistency functions
  - DB2 BACKUP SYSTEM
- Automated backup offload management

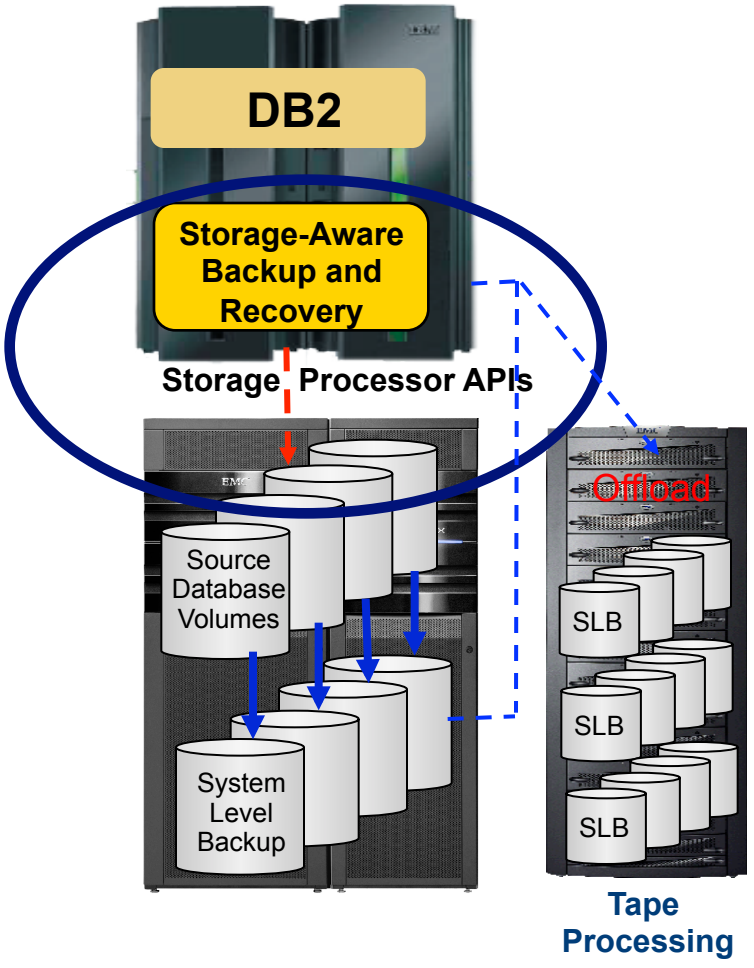

DB2 Application and Object Backup Using Data Set Based Fast-Replication

- DB2 backups performed at application or object level
- Supports share levels reference and change
- Backups performed using data set fast replication facilities
- Backups can be registered in repository and used for fast restore and parallel recovery
- DB2 image copies can be created and registered

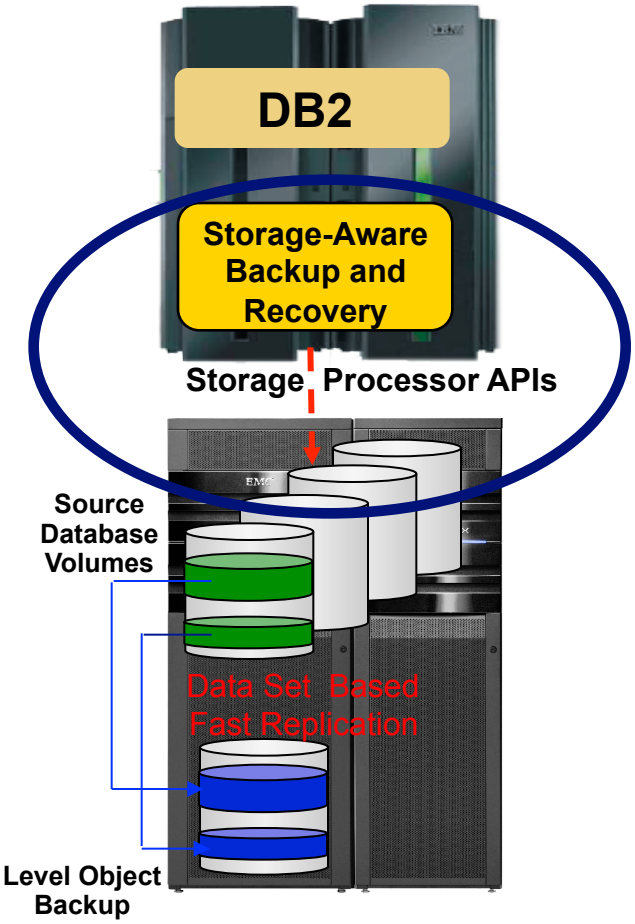

DB2 Image Copy Creation

- Image copies created from a system level backup
  - Image copies can be created and registered in DB2 syscopy
  - Eliminates I/O contention to maintain production performance
  - All image copies created at the same point in time
  - Reduces recovery time
- Image copy created from a data set fast-replication
  - Can be share level change or reference
  - Share level reference performs tablespace quiesce before data set fast-replication operations
  - Copies can be deleted after image copy creation
  - Fast-replication backups can be persistent, registered in repository, and used for restore and recovery operations

DB2 System Level Backup
System and Application Recovery

- Recover DB2 systems or application objects from disk or tape automatically
- Intelligent Recovery Manager invoked to optimize recovery plans
- Faster recovery
  - Instantaneous system-restore process
  - Coordinated and parallel restore and DBMS recovery operations minimize system downtime
- DB2 system backup can be used for object or application recovery
  - Data sets snapped to restore data
  - Parallel log apply reduces recovery time
- One system backup used for system, application, and disaster recovery

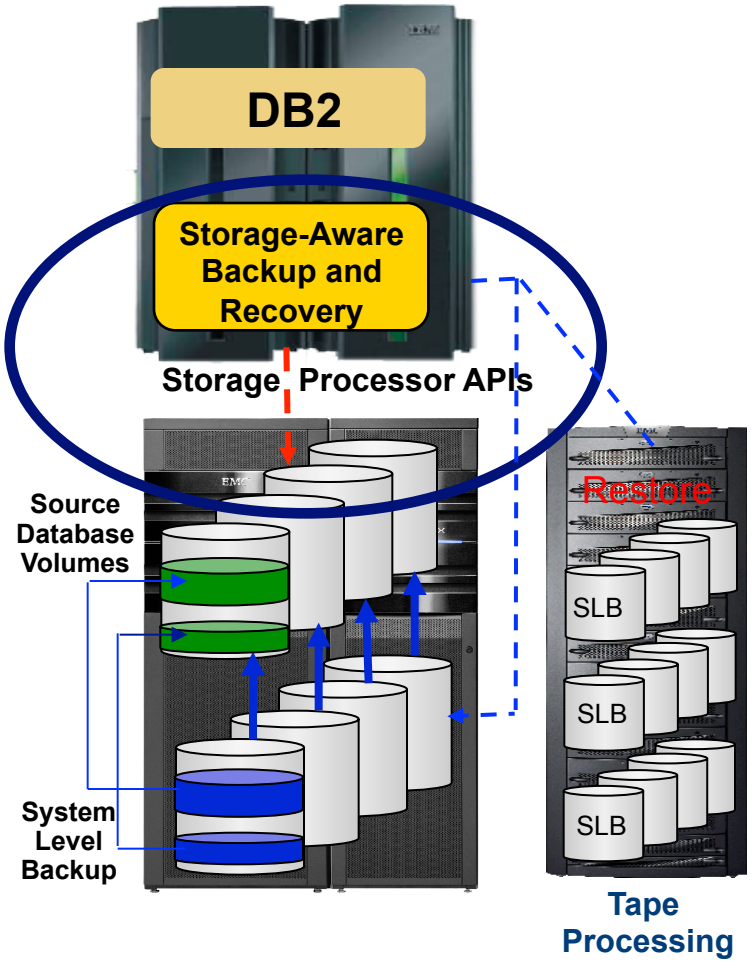

DB2 System Level Backup
Disaster Recovery Benefits

- Simplifies disaster recovery operations
  - System level backup for restart
  - System level backup and roll forward
- System backup is “restartable”
  - Restore volumes containing the last SLB
  - Performs recovery during normal database initialization process
  - Disaster recovery is as simple as restarting from a power failure
- Intelligent Disaster Recovery Manager
  - Prepares recovery assets and manages remote restore and recovery operations
- Reduced recovery time at a DR site
- Transform disaster recovery procedures into a tape-based disaster restart process
  - Similar benefits as storage-based remote replication solutions

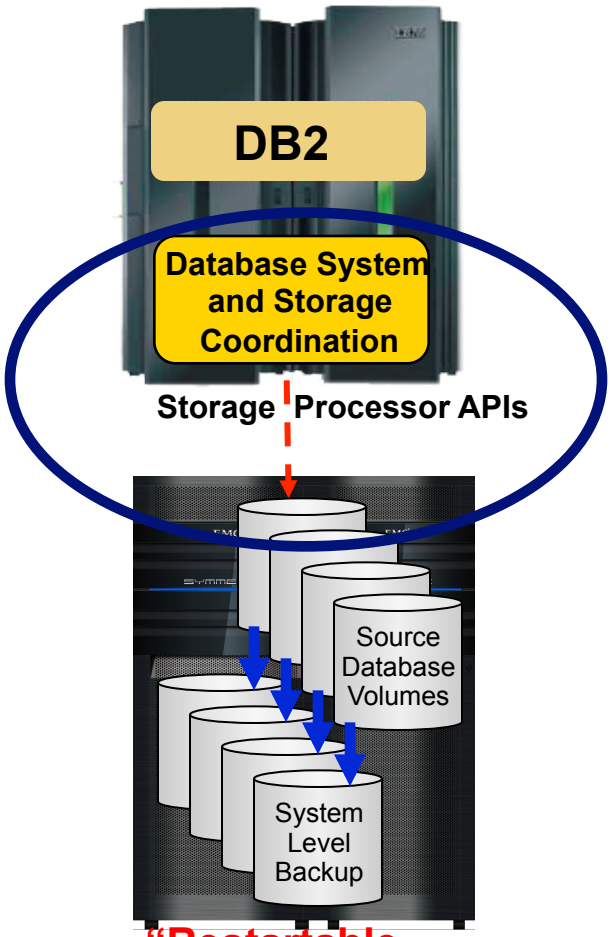

DB2 System Level Backup
Storage Benefits

- A system backup used for multiple functions
  - Saves storage and processing resources
- Leverages storage-processor and fast-replication software investments
- Expose fast copy capabilities to the DBAs ***safely and transparently*** using "storage-aware" database utilities
- Provides a sophisticated infrastructure and metadata to manage DB2 and storage processor coordination
- Multiple storage vendor support
  - IBM - FlashCopy
  - EMC - TimeFinder/Mirror/Clone/Snap, FlashCopy
  - Hitachi – ShadowImage, FlashCopy
  - IBM RAMAC Virtual Array, STK - SnapShot

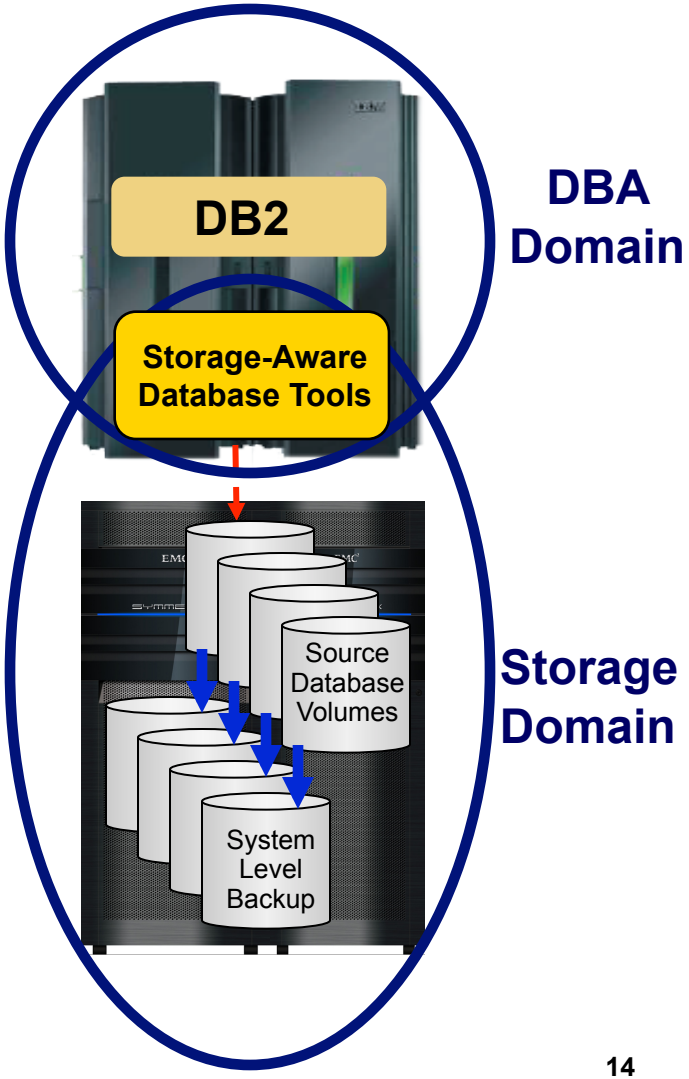

Cloning DB2 Systems

- Performs DB2 cloning automation
  - Simplifies database system cloning processes
  - Reduces cloning time and administration costs
- Leverages fast-replication facilities to clone data
  - Data can be cloned while online or offline
- Performs rapid volume reconditioning and data set renaming on cloned database volumes
  - Critical component of the database system cloning process
- Adjusts target database system to accommodate and accept the cloned data
  - DB2 catalog, directory, BSDS, active / archive log, etc.

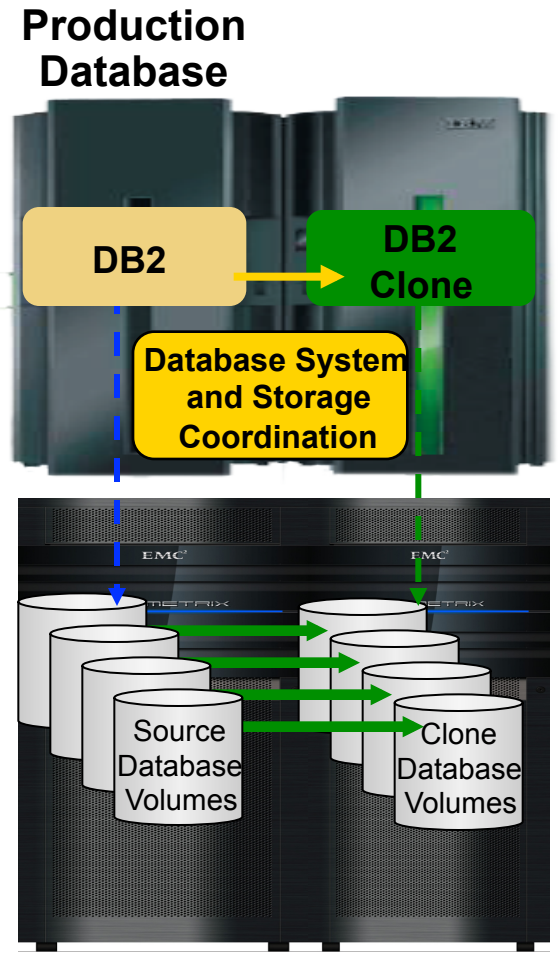

Refreshing DB2 Objects

- Performs automated DB2 table and index space refresh operations
  - Fast refresh of database objects
  - DB2 RI relationships, LOBS, and Identity columns
- Verifies source and target database compatibility
- Objects copied using storage-based data set fast replication
  - Target takes up the same amount of space as the source
- Performs object ID translations and target system metadata management

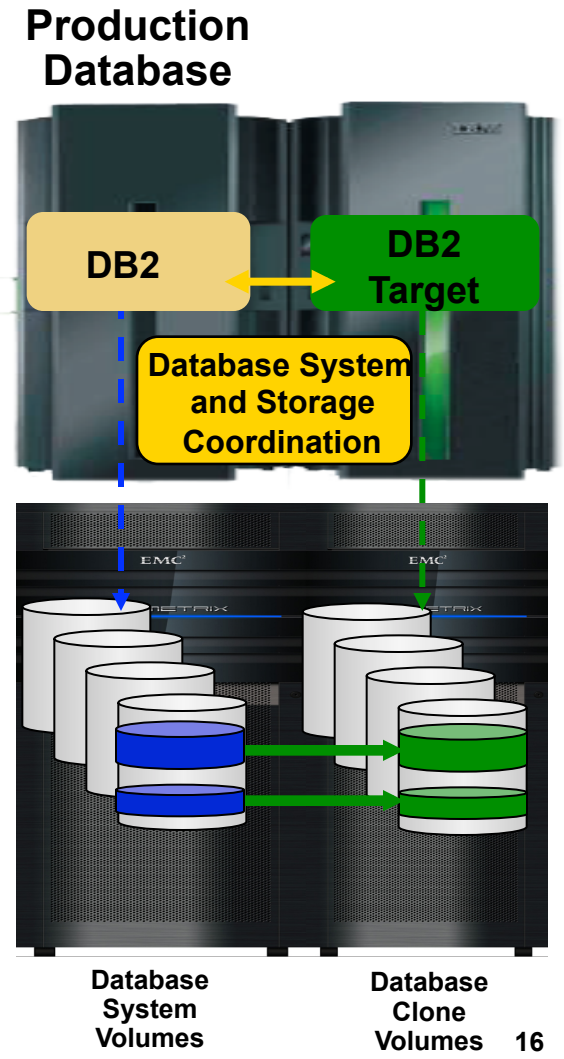

DB2 Storage-aware Products

- **IBM DB2 Recovery Expert for z/OS**
  - DB2 for z/OS backup and recovery
- **IBM DB2 Cloning Tool for z/OS**
  - DB2 system cloning and tablespace refresh
- **Mainstar Database Backup and Recovery for DB2 on z/OS**
  - DBR for DB2 – DB2 for z/OS backup and recovery
- **Mainstar Volume Clone and Rename**
  - VCR – DB2 system cloning automation
- **Mainstar Fast Tablespace Refresh**
  - FTR – DB2 tablespace refresh automation
- **EMC - Rocket Backup and Recovery for DB2 on z/OS**
  - RBR – DB2 for z/OS backup and recovery
  - EMC Select product

Implementation Planning Considerations
Examples based on DBR for DB2, VCR, EMC and IBM Storage

- System level backup usage
  - Determine how SLB(s) will be used
- SLB type
  - Determine full, data-only, or partial SLB requirements
- Backup frequency and space utilization
  - Determine backup frequency, performance, and space efficient fast-replication requirements
- Disaster restart considerations
  - Determine offsite disaster restart resources and preferences (RTO, RPO) to define appropriate disaster recovery profiles
- Copy blade selection
  - Determine storage processor capabilities, available facilities and fast-replication preferences

DB2 System Level Backup Usage and Data Set Layout Considerations

- DB2 SLB used for local system recovery
  - DB2 data and recovery structure isolation required
  - DB2 system isolation may be required
    - Non-database data sets will get restored when DB2 system is restored
    - User catalogs will get restored
- DB2 SLB used for application or DB2 object recovery
  - Data and recovery structure isolation is not required
- DB2 SLB used for remote disaster restart operations
  - Recovery structure isolation is not required
  - Database system isolation may be required
    - Non-database data sets will get restored when database system is restored
    - User catalogs will get restored

DB2 System Level Backup Usage
Data Set Layout for Full Backup / System Recovery

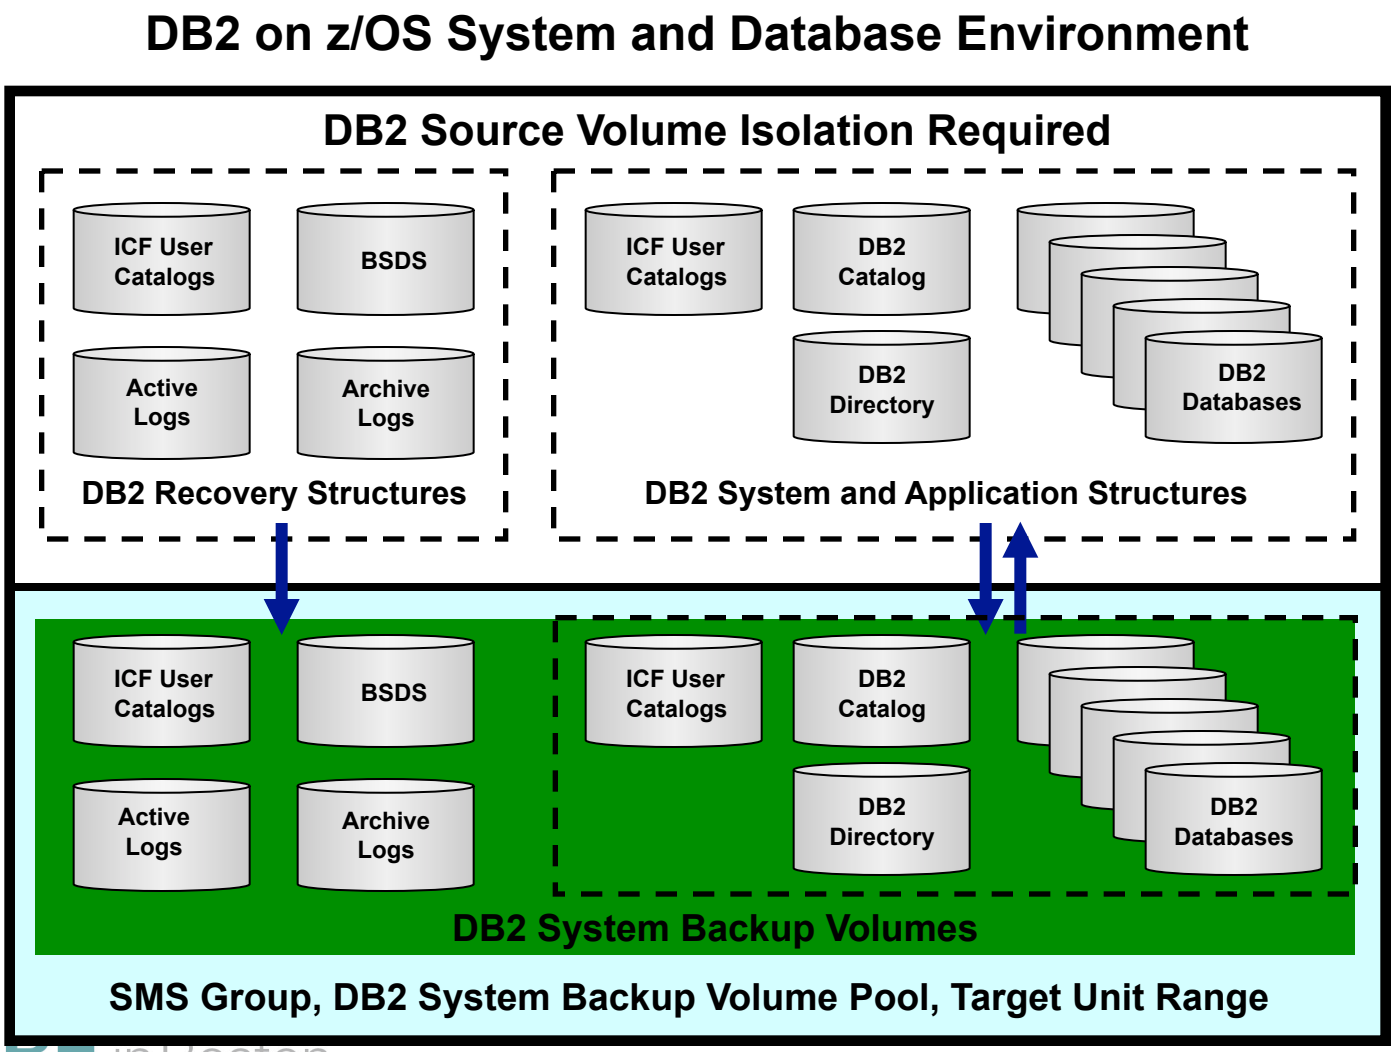

DB2 System Level Backup Usage
Data Set Layout for Data Only / Application Recovery

DB2 on z/OS System and Database Environment

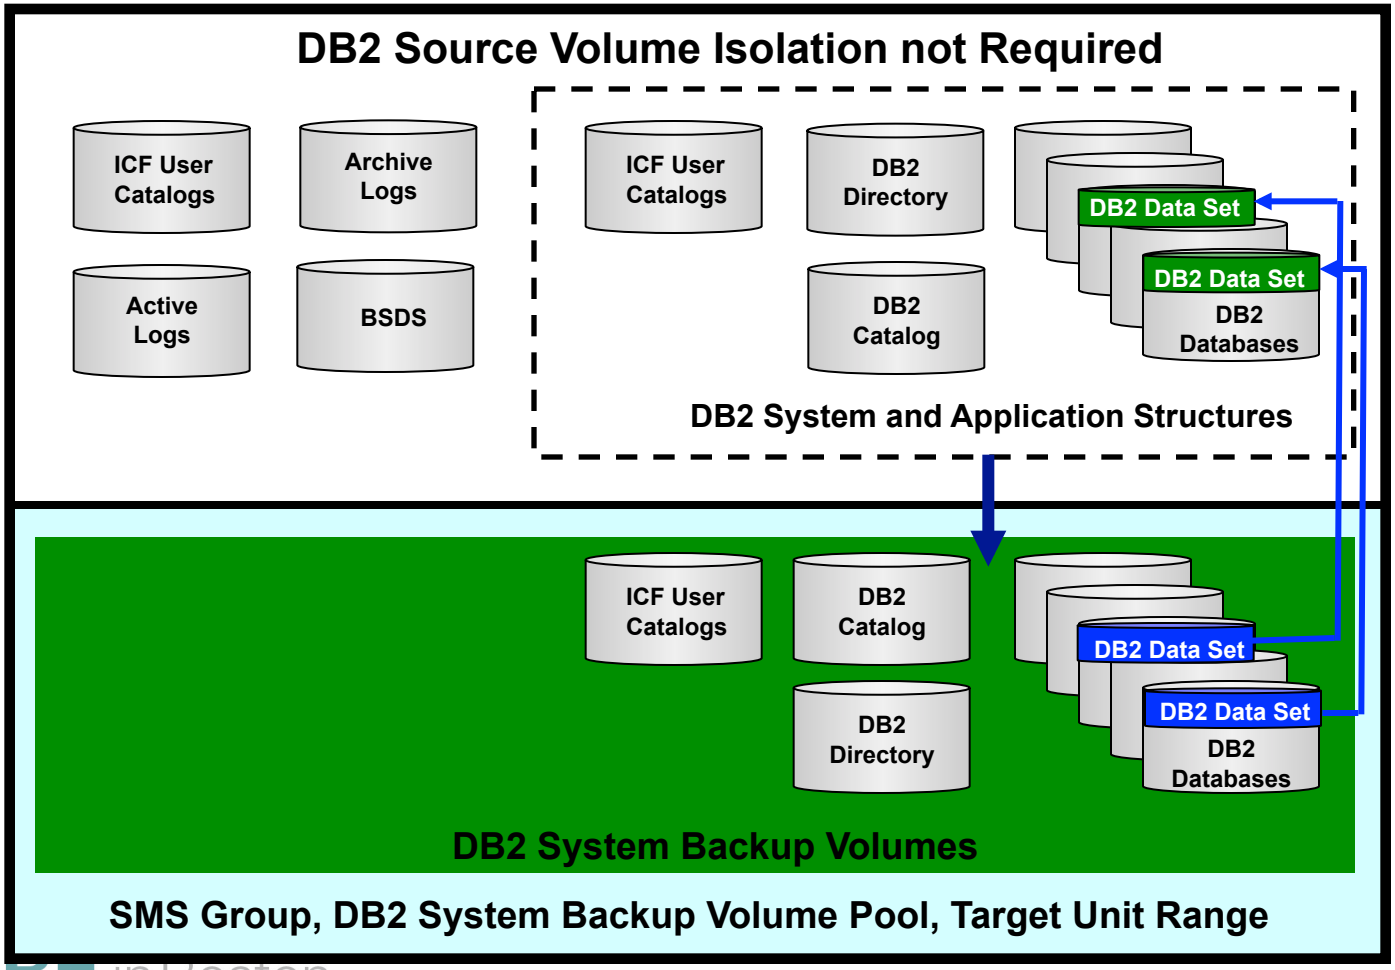

DB2 Partial System Level Backup
Data Set Layout for Application Recovery

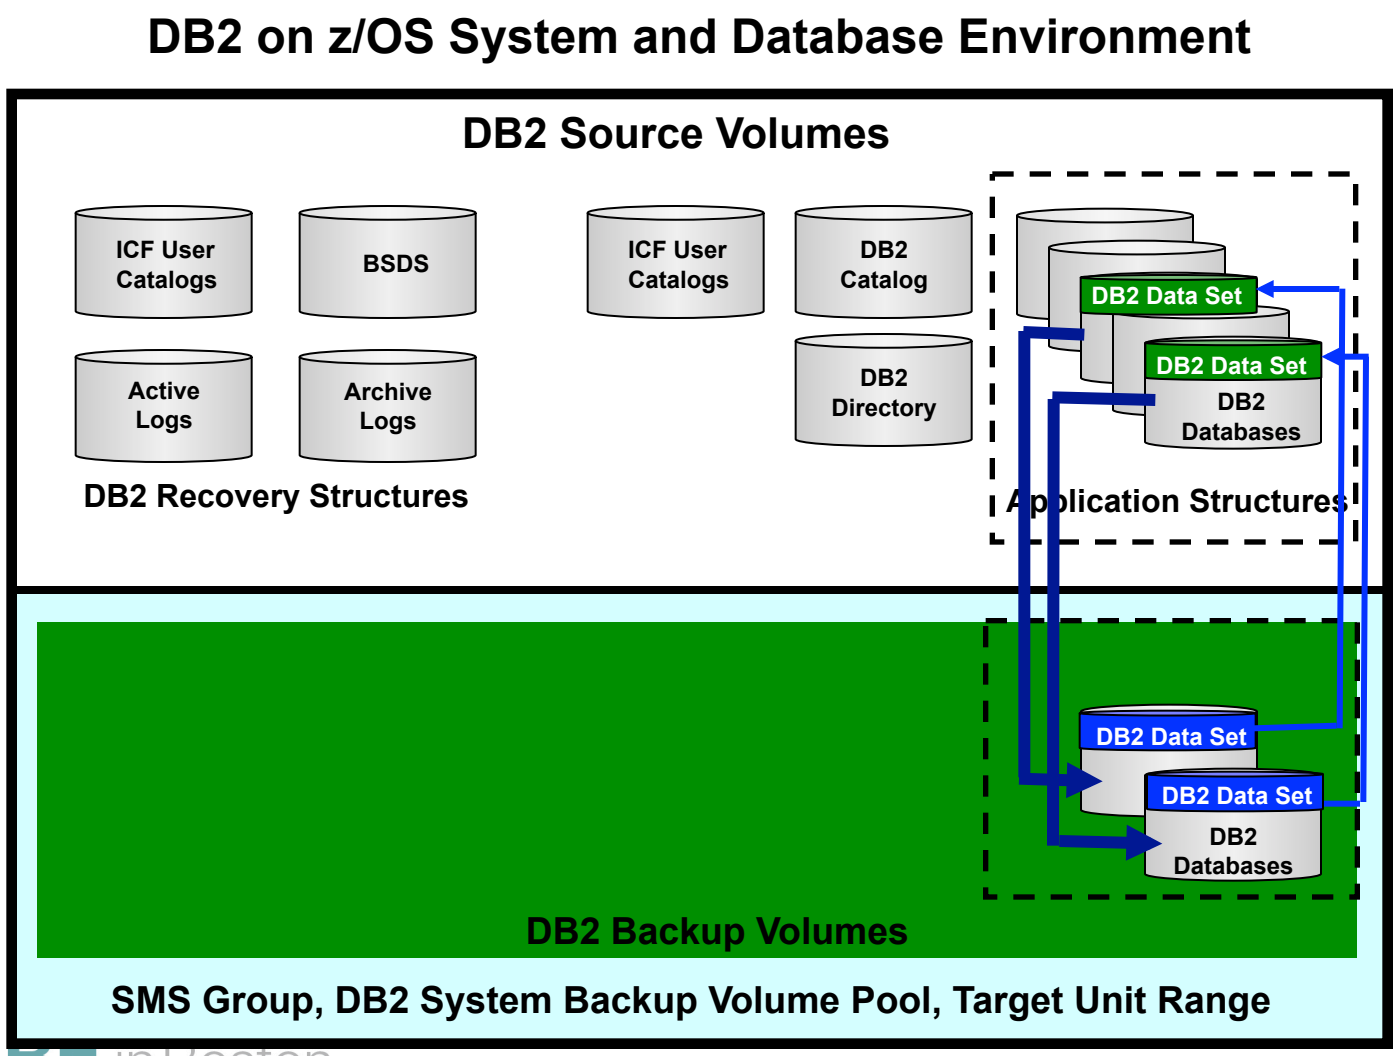

DB2 Partial System Level Backup

- Partial system level backup (PSLB)
  - Backup volumes representing a subset of the DB2 system
  - PSLB's used for database or application recovery only
  - Data set fast replication used to restore data
  - Log and data isolation not required
  - Desired application database data should be grouped on volumes as a best practice
- PSLB cannot be used for system recovery
  - System recovery requires all volumes in SLB
- PSLB usage
  - Large databases or applications having unique backup requirements
  - Creating image copies from a PSLB
  - Reduce disk utilization
  - Support more backup generations

Implementation Planning
Backup Frequency, Space, and Resource Usage

- SLB type: full, data-only, or partial – shown in previous slides
- Determine optimal backup frequency
- Determine number of backups to keep online (on disk)
  - Establish online backup duration requirements
    - SLB or PSLB used for IC creation may be deleted after ICs complete
- Determine offline (tape) backup requirements
- Consider incremental fast-replication options to reduce background copy time and resources
- Consider using one set of volume targets to support multiple database systems – next slide
  - Saves fast-replication target volume storage requirements
- Consider using space efficient fast-replication methods like EMC VDEVs to save space – later slides
- Consider cloning database systems to space efficient volumes using a full volume clone or SLB as the source – later slides

One Set of Backup Volumes for Multiple DB2 Systems

- Backup DB2–1
  - SLB-1 created on disk
  - Archive SLB-1
  - Backup volumes are available after archive completes
- Backup DB2–2
  - SLB-2 created on disk
  - Archive SLB-2
  - Backup volumes are available after archive completes
- Repeat for DB2-1
- Repeat for DB2-2

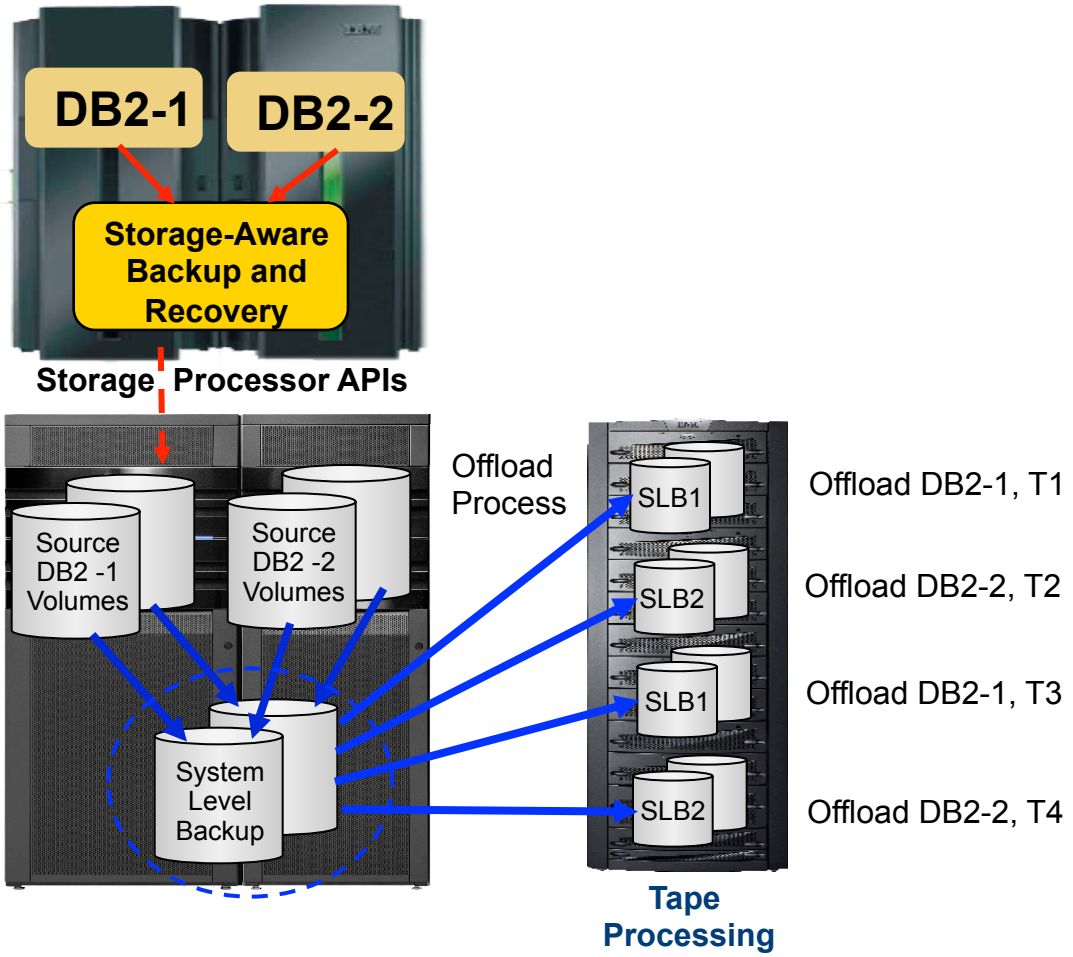

Create SLBs and Clone DB2 Systems
Full and Space Efficient Volumes

- EMC TimeFinder example
- TimeFinder/Clone
  - Full volume copy
  - Relationship can be retained with production volume
  - Allows TimeFinder/Clone incremental resynchronization
- TimeFinder/Snap Virtual Device (VDEV)
  - Space efficient copy
  - Allows TimeFinder/Snap incremental restore
  - Can have multiple TimeFinder/Snap volumes associated with production volume

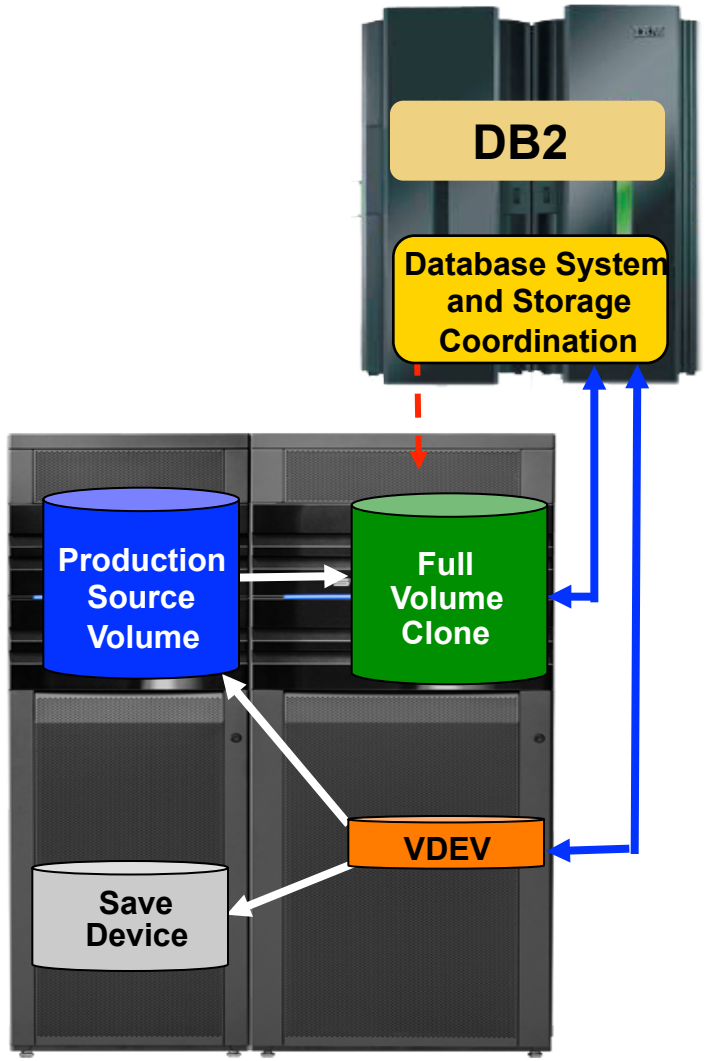

TimeFinder/Snap VDEV Operation Overview

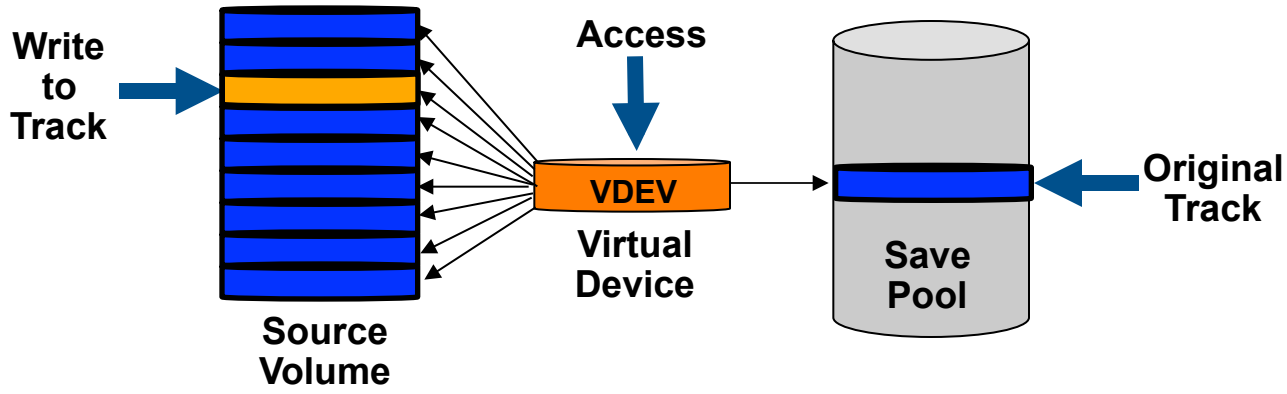

- The Snap target is accessible when the copy session is activated
- The first time a track on the source volume is written to:
  - Original data on the source volume is copied to a save volume (pool)
  - Pointer on the VDEV device is changed to point to the save pool
  - The host write is written onto the track of the source volume in cache
- The track on the source volume is then updated
- Unchanged data stays in place on the source volume

Space Efficient Usage Economics
Enable Frequent DB2 SLB or Clone Copies

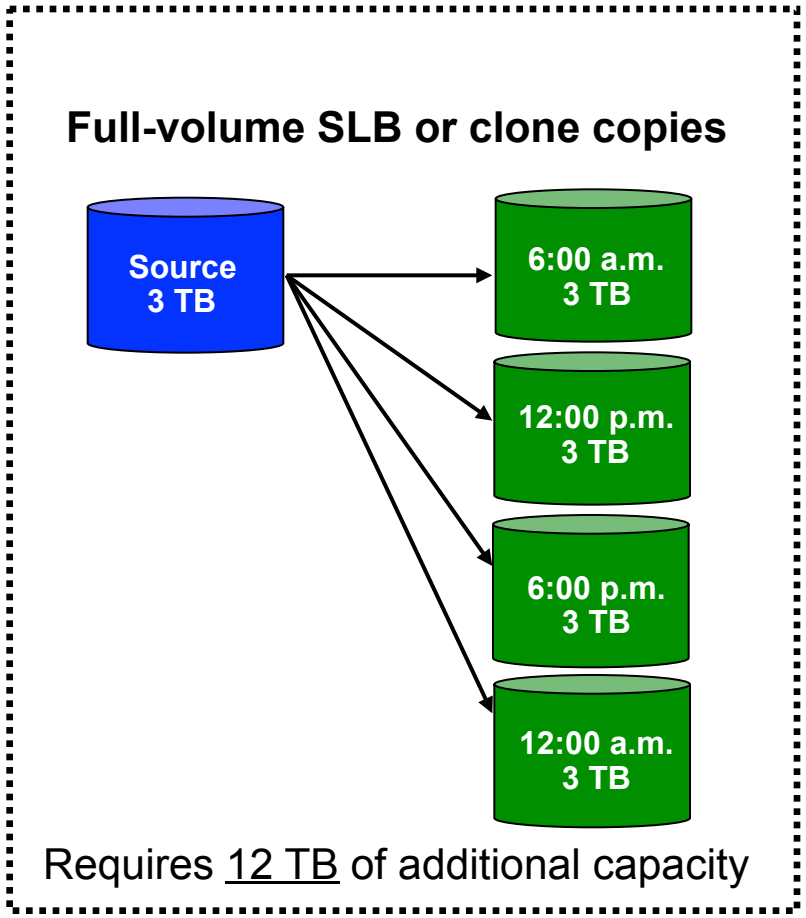

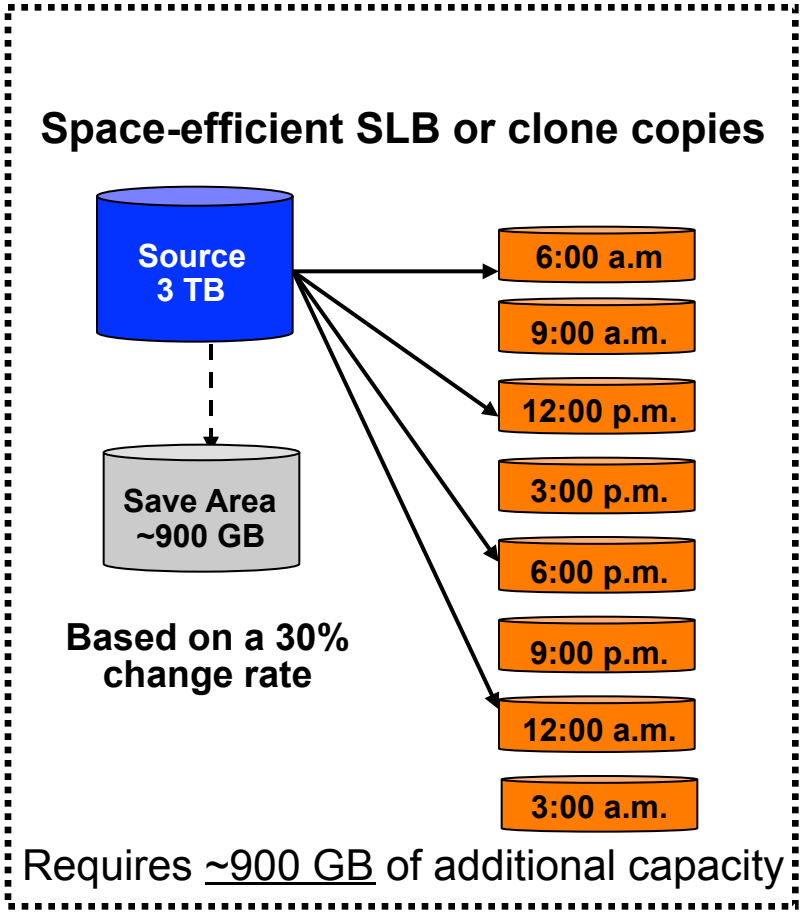

Full Volume and Space Efficient
Usage Example

- Full DB2 system-level backup created using full volume fast-replication
- DB2 clone operations performed using SLB backup volumes as source
- Cloned DB2 systems use virtual storage devices (VDEVs)
  - DB2 SLB volumes are used to service I/O for DB2 clone access
  - DB2 clone writes (few) go to save pool
  - DB2 SLB writes (none) go to save pool
- Storage-aware database tools provides infrastructure and metadata to manage DB2 and storage processor coordination

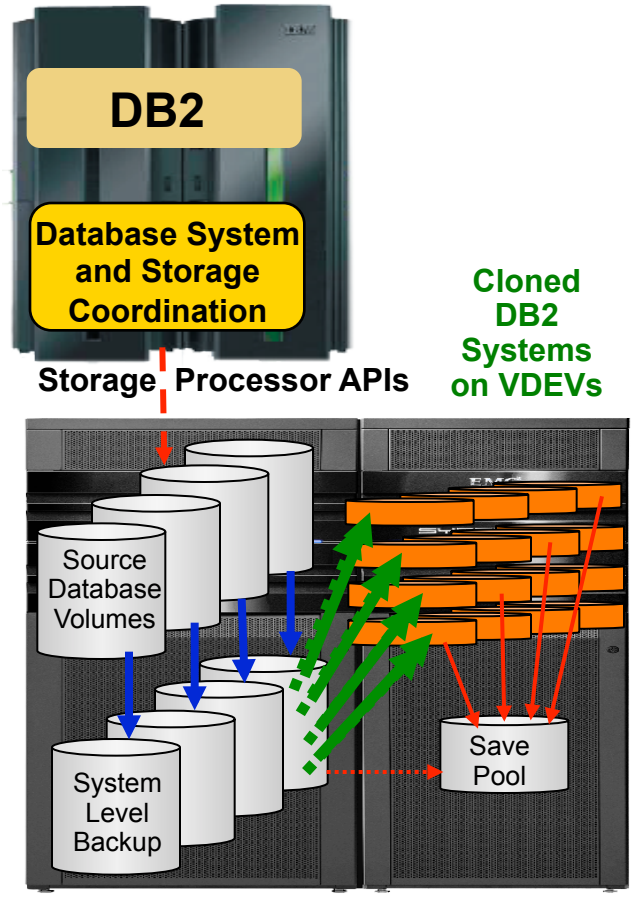

Full Volume and Space Efficient
Usage Example (2)

- Perform full volume DB2 cloning automation
  - Requires same amount of space as the source
- Perform space efficient clone operations
  - Use full volume clone as the source
  - No real space used for space efficient clones unless they are updated
- Operational automation may be required to re-instantiate space efficient clones when the full volume clone is re-instantiated

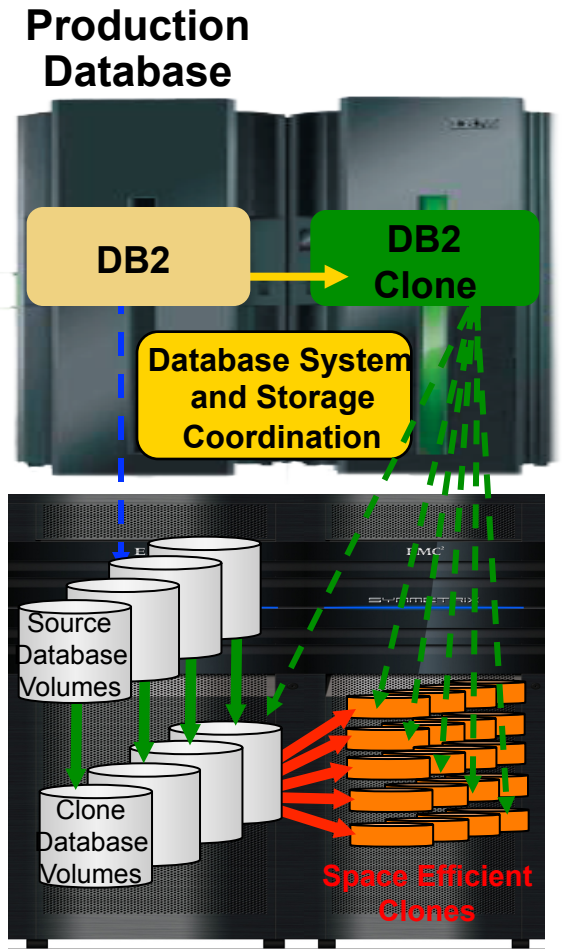

Implementation Planning
DB2 Disaster Restart Considerations

- DB2 SLB should contain database system data only
  - Can contain other data that is restarted together
    - Recovering database and other data together may require using a storage based consistency function to create the SLB
    - Cannot roll forward if database and other data require consistency
- Use disaster recovery profiles to prepare for roll-forward recovery at the DR site
  - Disaster recovery profiles specify options on how to copy log data for DR site, etc.
  - Ensure disaster recovery metadata is taken offsite with archive logs and image copies (Example DBR for DB2 DR PDS)
  - Reduces recovery point objectives (RPO)

Using DB2 SLBs for a Tertiary DR Site

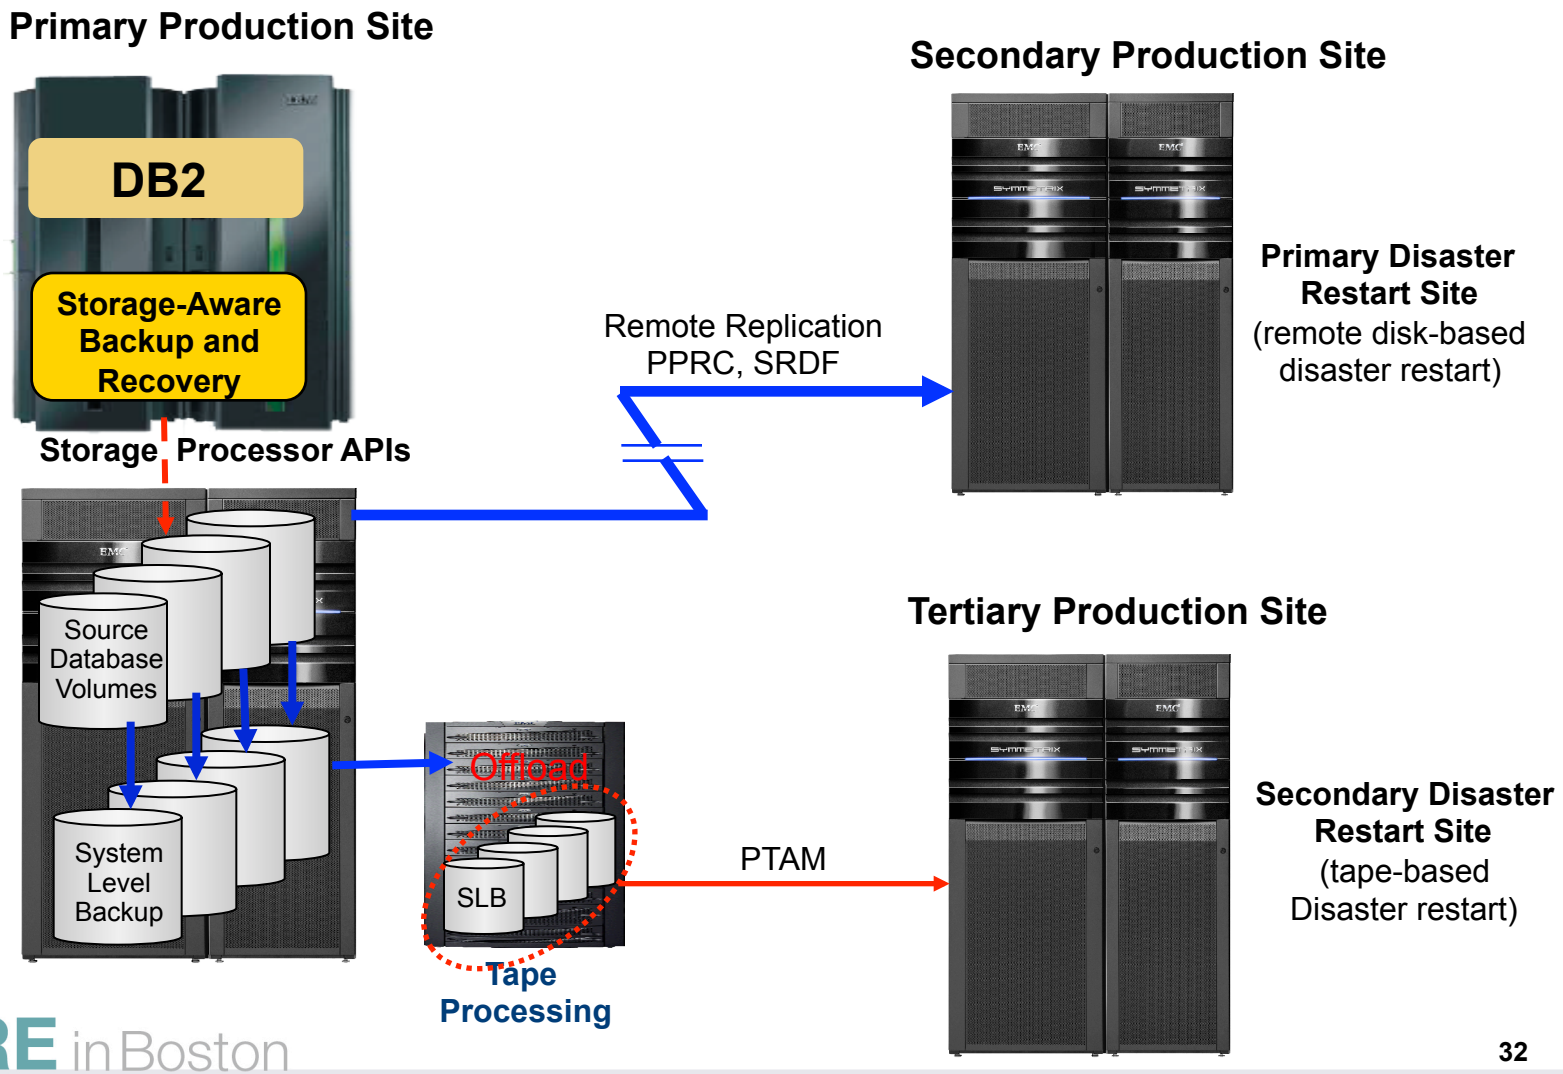

DB2 SLBs with PPRC Remote Pair FlashCopy

- Storage Aware Backup/Recovery and “Remote Pair FlashCopy” Support
  - FlashCopy to PPRC Primary volume **while maintaining Full Duplex**
  - FlashCopy Metro Mirror implementations only
- Preserve Mirror support option specified in installation ParmLib (FCTOPPRCP)
  - N - Do not allow the PPRC primary to become a FlashCopy target
  - Y - The pair can go into a duplex pending state
  - P - It preferable that the pair does not go into a duplex pending state.
  - R - It is required that the pair not go into a duplex pending state
- Copy Blade Support
  - DB2 BACKUP SYSTEM blade
  - FlashCopy Blade
  - DFSMSdss copy blade

System Level Backup Without Remote Mirror FlashCopy

- SLB causes backup volume data to be copied through PPRC link
- SLB can cause PPRC duplex pending state
- SLB restore can cause PPRC duplex pending state

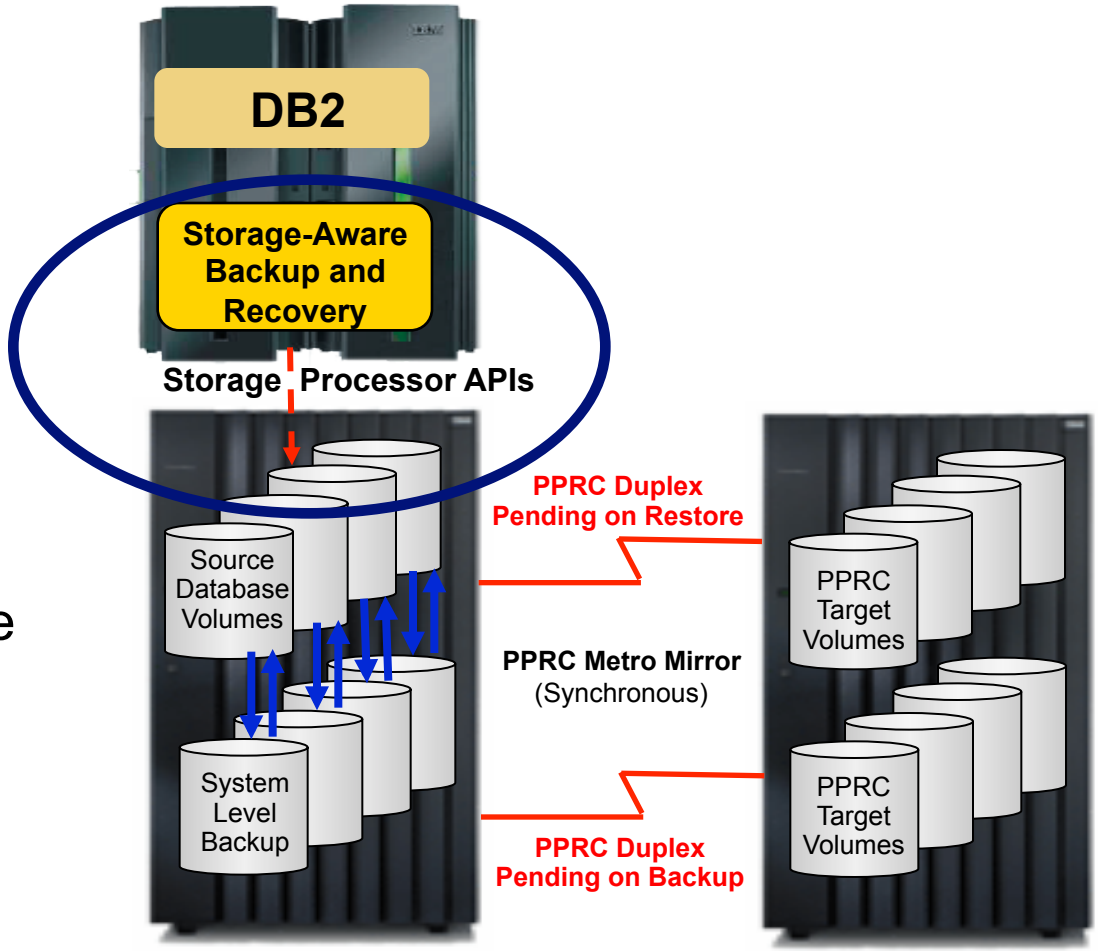

System Level Backup With Remote Mirror FlashCopy

- FlashCopy data is not copied over PPRC links
- SLB drives remote pair FlashCopy operation
  - Remote PPRC production volumes Flashed to remote PPRC SLB volumes
- System level restore drives remote pair FlashCopy operation
  - Remote PPRC SLB volumes Flashed to remote PPRC production volumes

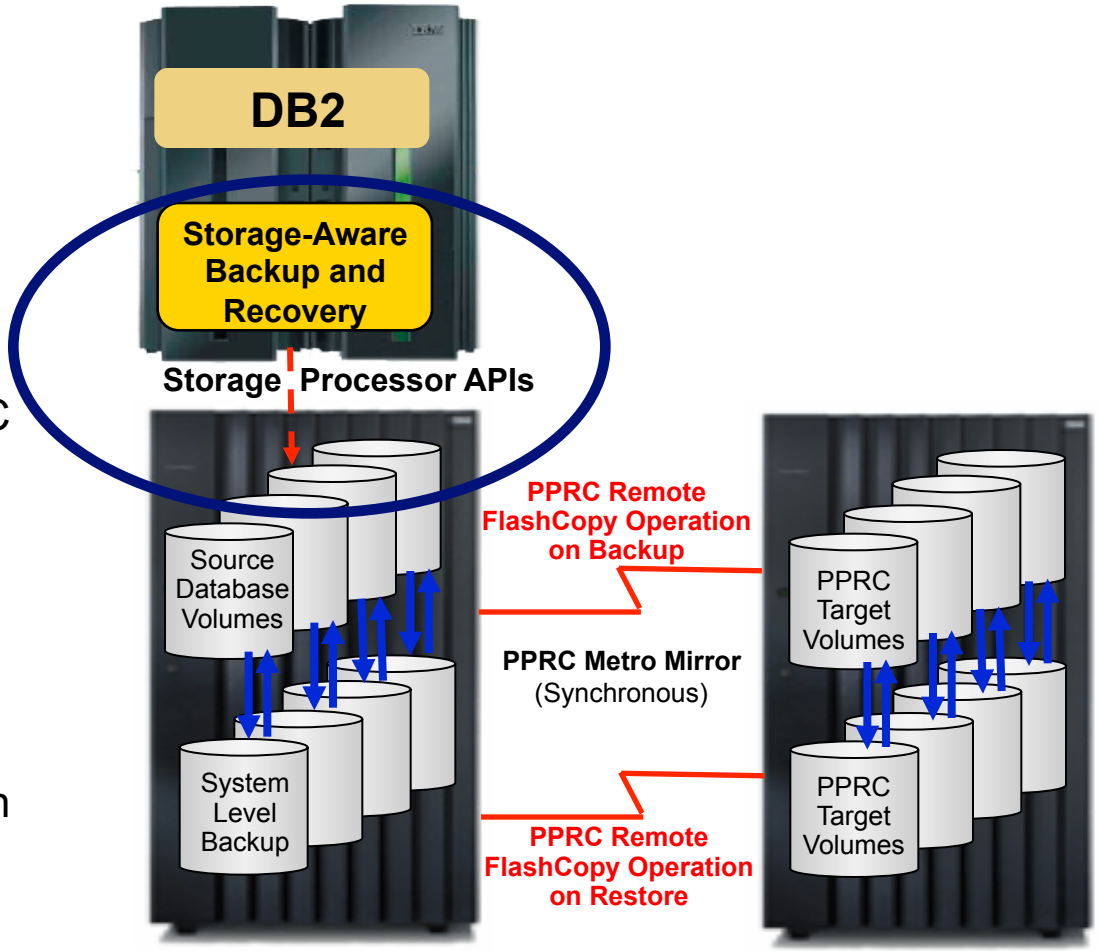

DB2 SLBs with XRC and PPRC without Remote Pair FlashCopy

- Assume DB2 volumes are primary volumes in a PPRC metro mirror or XRC relationship
- Backup target volumes must not be in a PPRC or XRC relationship
- Backup volumes cannot be used for DB2 system recovery without duplex pending state
- DB2 application and object recovery allowed
  - DBr for DB2 performs application and object recovery by copying data sets from the backup volumes to the source volumes
  - DFSMSdss used to copy data sets
    - Fast Replication Preferred option used to copy data
    - DFSMSdss uses slow copy methods as data sets cannot be Flashed to source PPRC or XRC volumes.

Implementation Planning
Copy Blade Selection

- Know your storage processing infrastructure
  - What storage processors are used (EMC, IBM, HDS)
  - What fast-replication facilities are licensed and preferred
- Determine storage blade and fast-replication facilities to use
  - DB2 Backup System Blade
  - DFSMSdss Blade
  - IBM FlashCopy Blade
  - EMC TimeFinder Blade
  - HDS ShadowImage Blade
- Determine which type of consistency function is best for your environment
  - Database suspend, storage-based consistency

IBM Copy Blades

- IBM DB2 BACKUP SYSTEM Blade
  - Provide support for DB2 Backup System
- IBM FlashCopy Blade
  - Provides support for IBM FlashCopy V2
  - Data set FlashCopy support for fast object / application recovery
  - ANTRQST calls issued to drive FlashCopy volume commands (fast performance)
  - Requires Database Log Suspend
  - Supports IBM, EMC, HDS FlashCopy products
- IBM DFSMSdss Copy Blade
  - ADRDSSU utility invoked to perform volume copies
  - Fast replication (preferred) is used – Will support non-fast replication volumes
  - Data set FlashCopy support for fast object / database (IMS) / application recovery
  - Slower than using ANTRQST in native FlashCopy blade
  - Requires Database Log Suspend
  - Supports FlashCopy (IBM, EMC, HDS), SnapShot (STK, RAMAC Virtual Array)

EMC Copy Blades

- EMC TimeFinder Blade
  - TimeFinder/Mirror
  - TimeFinder/Clone Mainframe Snap Facility
  - TimeFinder/Snap Virtual Devices
    - Allows multiple backups with reduced storage utilization
  - Incremental copy support for all copy methods
  - EMC Consistency Technology support for all volume copy methods
    - Reduce the need for database suspend functions
  - TimeFinder Data Set Snap facility to perform fast replication application / object restores
  - DB2 Log Suspend performed on one data sharing member when backing up a data sharing group

Hitachi Data Systems Copy Blades

- ShadowImage Blade
  - Supports HDS native ShadowImage volume copy processes
  - Invoked using FlashCopy backup profile
    - Checks *shadow_image* field in backup product parameter library
      - *N – DBR for DB2 drives FlashCopy*
      - *Y – DBR for DB2 drives ShadowImage*
  - Incremental copy support
  - Requires a database log suspend operation
  - Can support database systems that span HDS and IBM storage using native methods (ShadowImage and FlashCopy)
  - HDS data set FlashCopy emulation used for fast replication object / application restores

Session Summary

- Storage-aware DB2 utilities provide storage integration to simplify DB2 administration tasks
- DB2 system backup solutions leverage storage-based fast-replication facilities and investments
  - Fast and non-intrusive backup operations with less administration
  - Reduces host CPU, I/O and storage utilization
  - Backups can be used for system, application, disaster restart
  - Parallel recovery reduces system and application recovery time
- DB2 system cloning automaton allows production data to be leveraged easily and effectively
- DB2 table and index spaces refreshed easily
- Less skills required to implement advanced backup, recover, disaster recovery, and cloning solutions
- Implementation planning is important to optimize the benefits

Complimentary SHARE Sessions

- Simplify and Improve IMS Administration by Leveraging Your Storage System
  - Session Number 7987
  - Wednesday August 4, 2010
    - 3:00 – 4:00 PM