#SHAREorg

SHARE
Technology · Connections · Results

Rethinking Backup and Recovery Strategies with IMS Recovery Expert

Ron Bisceglia
Rocket Software

March 13, 2012
Session 10805

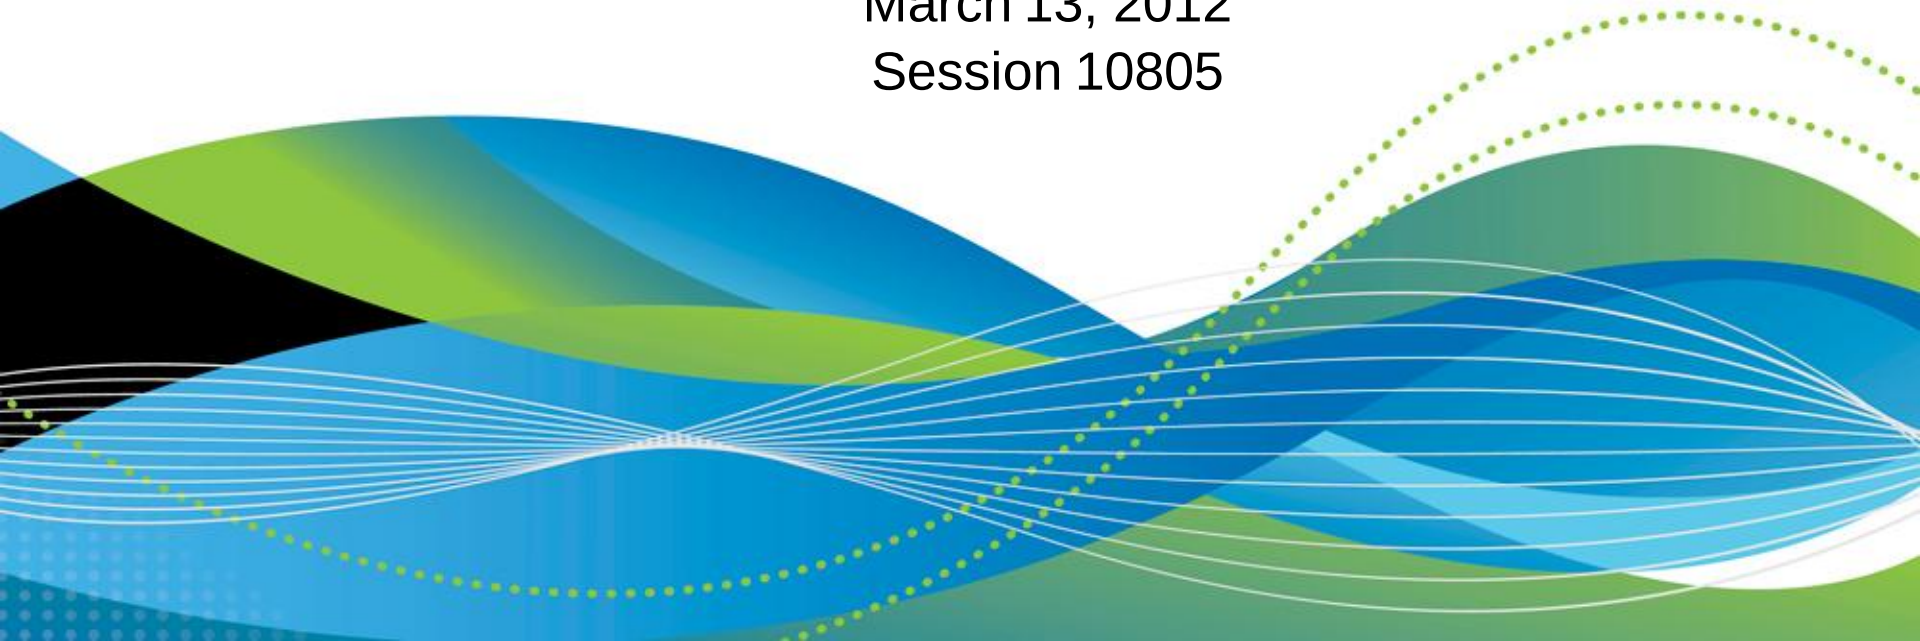

Topics

- Cost of Downtime
  - How often do we backup our data?
  - How do we backup our data?
- Fast Replication Storage
  - Host-based data copy options
  - What is Storage-based Fast Replication?
  - Why are Storage-aware Fast Replication Tools Needed?
- Reducing Costs to Create Backups
  - Leveraging fast replication
  - System level backups
- Improving Recoverability
  - Guaranteed backup
  - Simplified and automated recovery processes
  - Reducing recovery times

Trivia

- 2 TB hard drive = $79.00

- How many 3390 mod-1s = 2 TB?
  - Kilobyte 10_3 2_{10}
  - Megabyte 10_6 2_{20}
  - Gigabyte 10_9 2_{30}
  - Terabyte 10_{12} 2_{40}

- Tie Breaker : What year was 3390 disk drive first introduced?

Database Downtime Drives Up Costs

Most organizations spend an extra $1.5M USD per year because of unplanned database downtime

Yearly Cost Metrics	Best-in-Class	Industry Average	Laggards
Business interruption events	.9	3	3.5
Time per business interruption event (hours)	1.3	4.7	8.4
Total disruption (hours)	1.2	14.1	29.4
Average cost per hour of disruption	$60,000	$110,000	$98,000
Total cost of business interruption events	**$72,000**	**$1,550,000**	**$2,880,000**

Source: Aberdeen Group, Month 2010

Very few organizations have perfect or near-perfect datacenter uptime

- ✓Only 3% of organizations have uptimes of 100%
- ✓ Only 4% of organizations have uptimes of 99.9%

What Causes Outages

"Based on extensive feedback from clients, we estimate that, on average, unplanned application downtime is caused: 20 percent of the time by hardware (e.g., server and network), OSs, environmental factors (e.g., heating, cooling and power failures) and disasters; 40 percent of the time by application failures including "bugs," performance issues or changes to applications that cause problems (including the application code itself or layered software on which the application is dependent); and 40 percent of the time by operator errors, including not performing a required operations task or performing a task incorrectly (e.g., changes made to infrastructure components that result in problems and incur unexpected downtime).

Thus, approximately 80 percent of unplanned downtime is caused by people and process issues, while the remainder is caused by technology failures and disasters. Improving availability requires a different strategy and set of investment choices for each of the three unplanned downtime categories."

-- *Gartner Group*

Database Downtime Affects Your Business

Business

- Average cost of database downtime $1.5M USD/year
- Revenue at risk
- Customer satisfaction declines
- Missed services level agreements
- Brand damage and loss of goodwill

IT

- Time consuming, rarely used, manual backup/recovery procedures don't scale as data volumes grow
- Inability to backup all data because of shrinking maintenance windows and growing data volumes
- Difficult to get complete database backup without production impact

Evolution of Image Copies

- Batch Image Copies
  - Database is unavailable
  - No uncommitted updates
- Concurrent Image Copies
  - Database is still available for update
  - Logs need to be applied to make transactionally consistent
- Incremental Image Copies
  - Logs applied to image copy
  - Does not affect database availability
  - Could be ‘expensive’ depending on number of logs
- SMS Image Copies
  - Can use fast replication, reducing host CPU and I/O
  - Recovery utility must support this type of backup

Factors Affecting Cost to Backup Our Data

- Number of Database Data Sets
- Image Copy Frequency
  - High update activity can affect recovery time
- Number of Image Copy Generations
  - How far back we may have to recover to
  - Fall back in case more recent IC is bad
- Change Accumulations
  - Mainly done to reduce recovery time
  - Usually ‘expensive’ in terms of host CPU
  - Do we do this for all databases?
- Database Unavailability
- Backups for Disaster Recovery

Fast Replication Storage

Host Based Data Copy Options

Data copy processes use host based CPU and I/O facilities
More costly and slower than storage-based fast replication

- Volume copy options
  - DFSMSdss (IBM)
  - FDR (Innovation Data Processing)
  - TDMF (IBM)
  - FDRPAS (Innovation Data Processing)
- Data set copy options
  - DFSMSdss (IBM)
  - FDR (Innovation Data Processing)

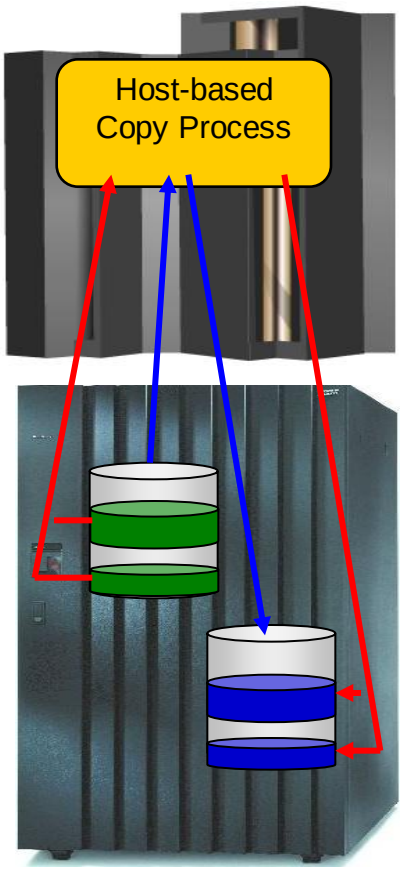

Storage Processors

- Storage
  - Non-volatile
  - Computer components and recording media that retain digital data
- Processor
  - Computing system that executes software programs
- Storage Processors = Storage + Processor
  - Enable data movement through host-based API
  - Allows CPU reduction from host-based I/O

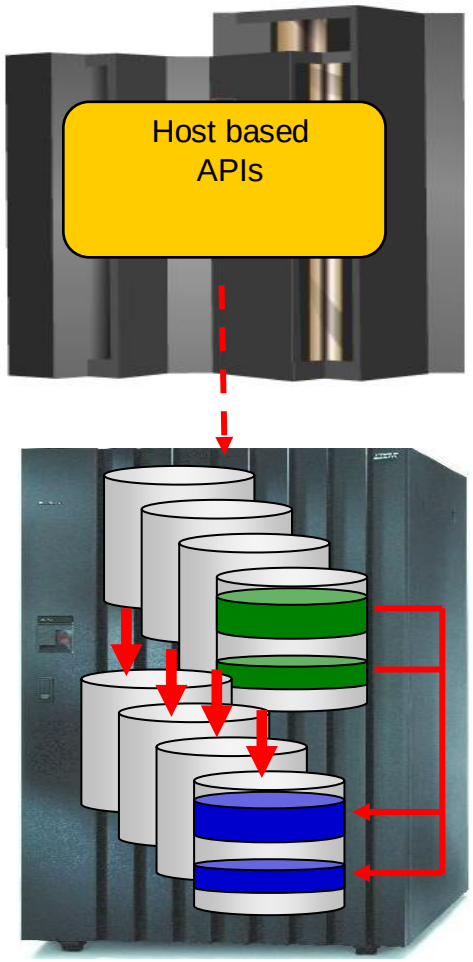

What is Storage-based Fast Replication?

- An instant copy of a volume/data set at a specific point in time
  - Builds a bitmap to describe the source volume
  - After the bitmap has been created, the source and target volume data can be used immediately
- Data movement (CPU and I/O) offloaded to storage processor
  - Frees up resources on host processor
  - No host CPU or I/O costs
- For volume replication a relationship is established between a source and a target
  - Geometrically similar devices
- Consistency Groups
  - Group of volumes copied at exactly the same point in time while maintaining the order of dependent writes

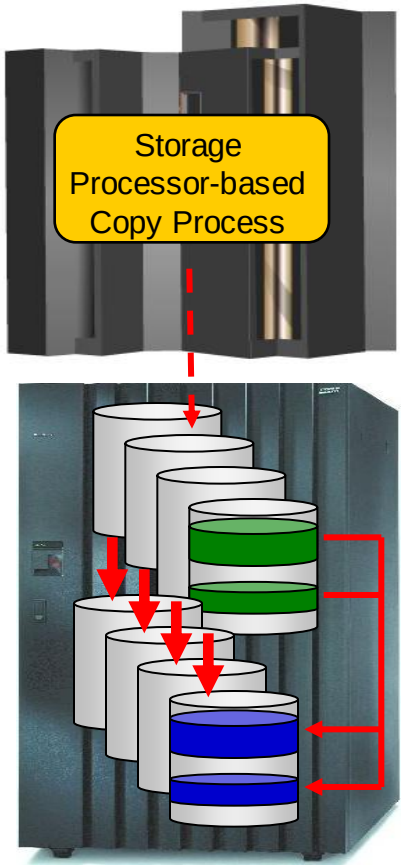

Local Storage-based Fast Replication

- First product availability – (late 1990s)
  - Used to streamline batch processing
  - Speed backup processing
- Data copied using storage processor fast replication facilities
  - Volume based
  - Data set based
- No application or database knowledge
- Typically used by storage administrators
- Examples
  - EMC TimeFinder
  - IBM FlashCopy
  - HDS Shadow Image

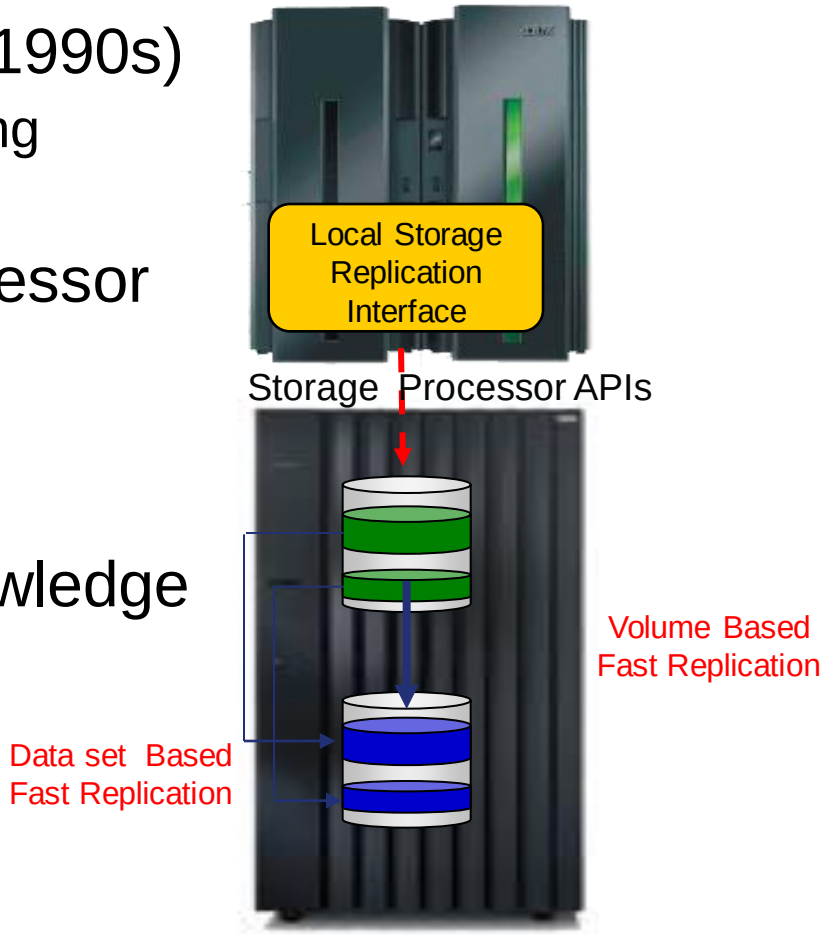

Copying with Consistency

- Dependent writes
  - The start of one write operation is dependent upon the completion of a previous write to a disk in either the same storage subsystem frame or a different storage subsystem frame
    - Transactions updating database and log for example
  - Basis for providing consistent data for copy operations
- Consistency
  - Preserves the order of dependent writes
  - For databases, consistent data provide the capability to perform a database restart rather than a database recovery
    - For databases, consistent data provides the capability to perform a data base restart rather than a data base recovery
    - Restart can be measured in minutes while recovery could be hours or even days

Methods for Creating Consistency Groups

- Time Based
  - Global Mirror (XRC)
- Data Freeze
  - Metro Mirror
  - Consistency Group Flashcopy
- Ordered I/O
- Quiescing/Stopping the Application or DBMS
  - DB2 LOG SUSPEND

Variations of Fast Replication

- BACKGROUND Copy
  - Storage processor keeps track of what has been copied
  - COPY from source may occur on reads to target
- NOCOPY
  - Data read from source
  - COPY and WRITE occurs on source updates
  - Useful for situations where data only needed for temporary amount of time
- PERSISTENT/INCREMENTAL
  - Relationship between source and target maintained
  - Only changed tracks copied
- Space Efficient or Virtual Devices (available 2007)
  - IBM Flashcopy SE
  - EMC Symmetrix Virtual Devices
  - Reduces disk space for target devices
  - Mappings or pointers to source data maintained in storage processors

Application and Database

Storage Integration

Application and Database Management Domain

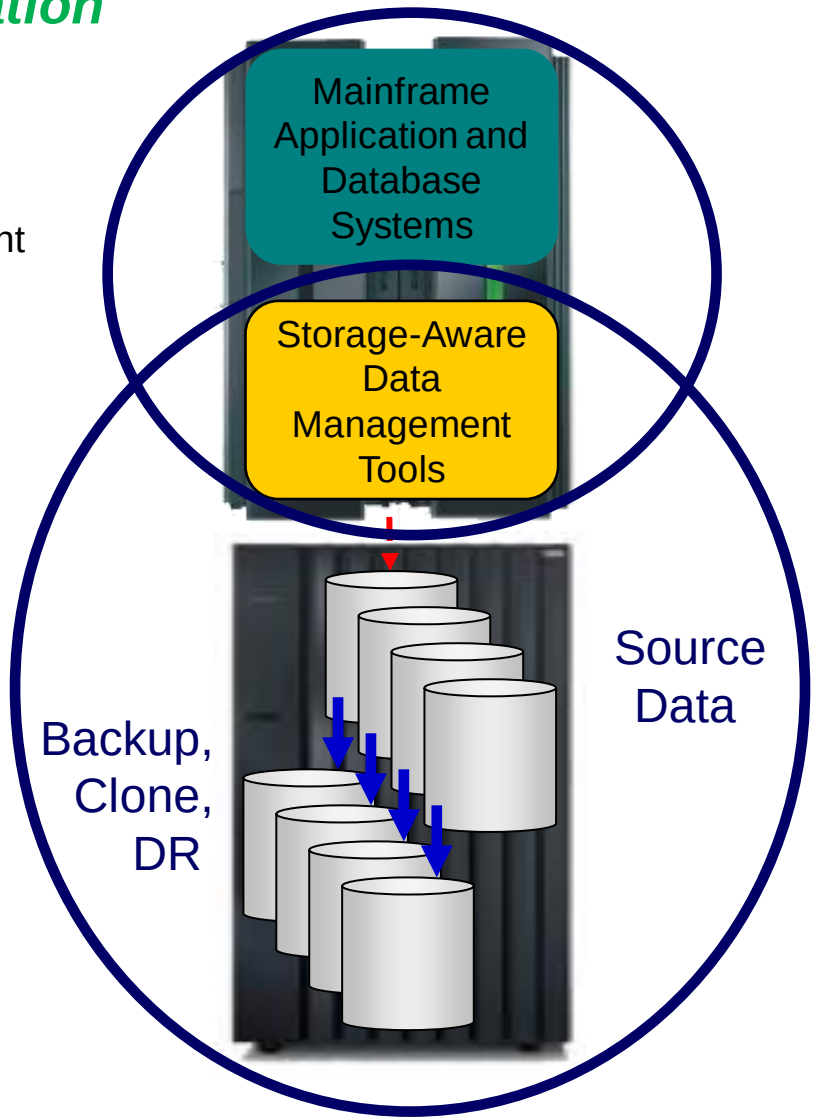

Storage Administration and Business Continuity Domain

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

Reducing Costs of Creating Backups

While Improving Recoverability

Backing Up IMS Data

- IMS Databases
  - Image Copies
  - Change Accumulations
  - Archive Logs
- IMS System and Application Libraries
  - Volume Dumps
  - DFSMSdss data set copies
  - DFSMShsm
- Remote Mirroring/Replication
  - XRC, PPRC, SRDF
  - Disaster Recovery only
- IMS Recovery Expert System Level Backups
  - System Level Backup (SLB)
    - IMS Databases, IMS System Data Sets, Recovery Structures, and Application Libraries in one backup at same point in time
  - Automated offloading and encryption of data
  - Meta-data repository to maintain information on system level backups
  - One backup for multiple uses

IMS Recovery Expert

System Level Backup - the 'Next Level' of Backup

- A System Level Backup is a backup of the entire DBMS environment at a point in time
  - Recorded in IMS Recovery Expert Meta data repository
- Leverages storage-based fast replication to drive a volume level backup
  - Backup in seconds
  - Offloading data copy process to the storage processor saves CPU and I/O resources
  - Significantly faster than data set copies
- Backup DBMS without affecting applications
  - Backup windows reduced by replacing image copies
  - Extends processing windows
- Data consistency ensures data is dependent-write consistent
  - IMS Recovery Expert Log Suspend
  - Storage-based consistency functions
    - **FCCGFREEZE** to perform a FlashCopy consistency group (transparent to the user)
  - Equivalent to a power failure

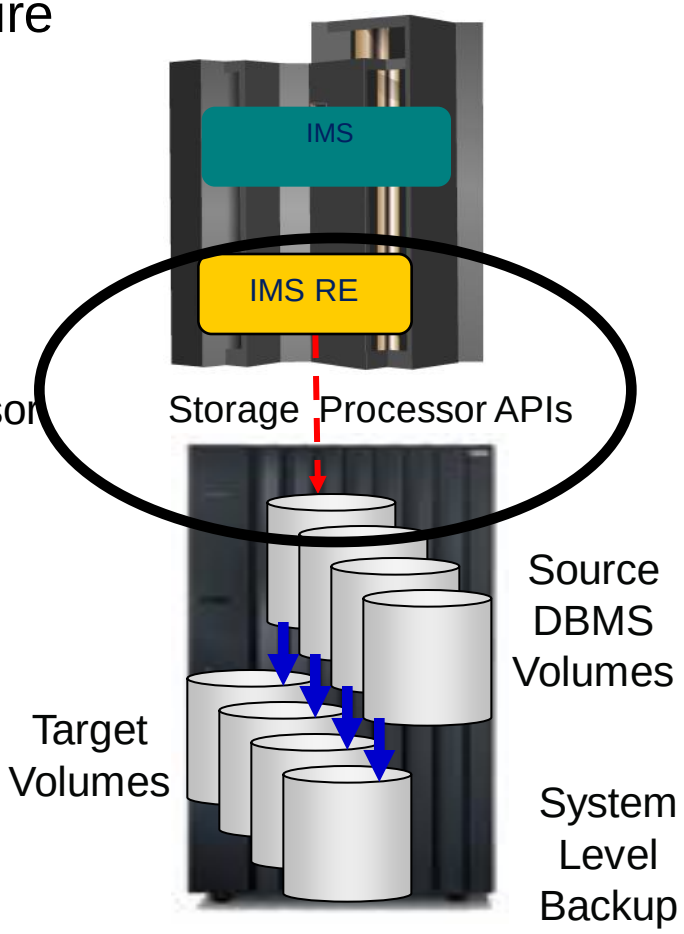

IMS Recovery Expert

System Level Backup - the 'Next Level' of Backup

- Backup validation each time ensures successful recoveries
  - Insurance that a backup is available
- Automated backup offload (archive/recall)
  - Copies system backup from fast replication disk to tape for use at either local or disaster site (or both)
- Can be used in combination with other backups (image copies)

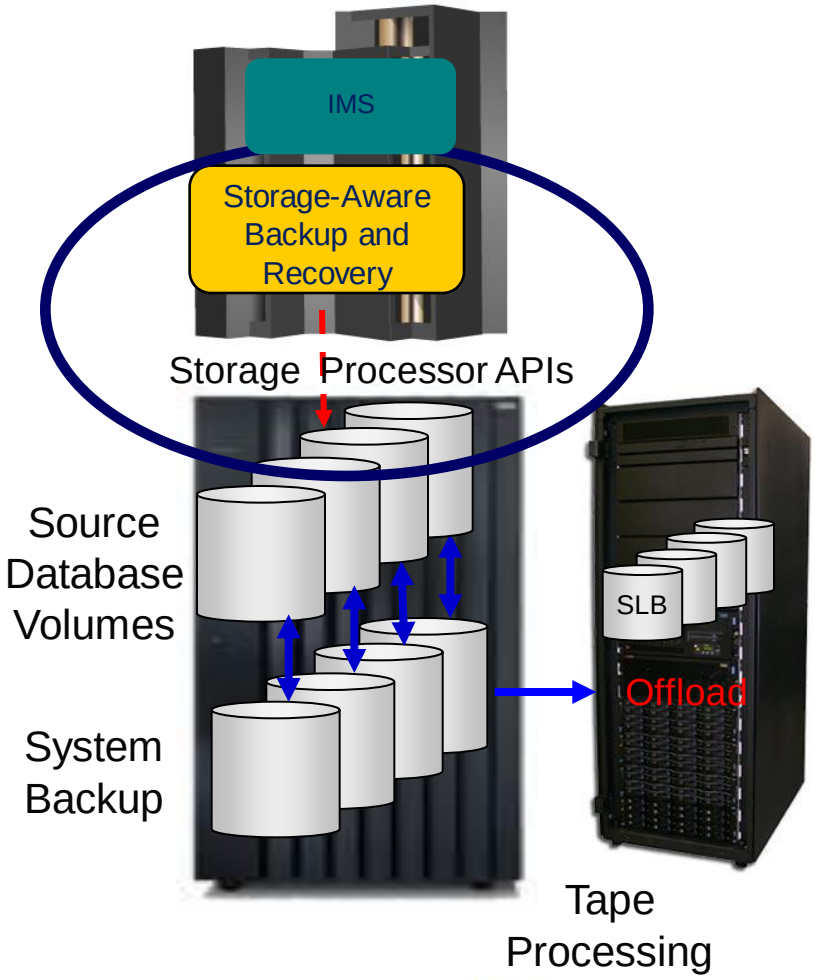

System Level Backup

System and Application Recovery

- Recover IMS systems, applications, or database from disk or tape automatically
- Intelligent Recovery Manager invoked to optimize recovery plans
  - Integrates with traditional recovery tools
- Faster recovery
  - Instantaneous system or application restore process
  - Parallel recovery reduces downtime
- One system backup used for:
  - Database Recovery
  - Application Recovery
  - System Recovery
  - Disaster Recovery
- Basis for coordinated IMS and DB2 Recovery

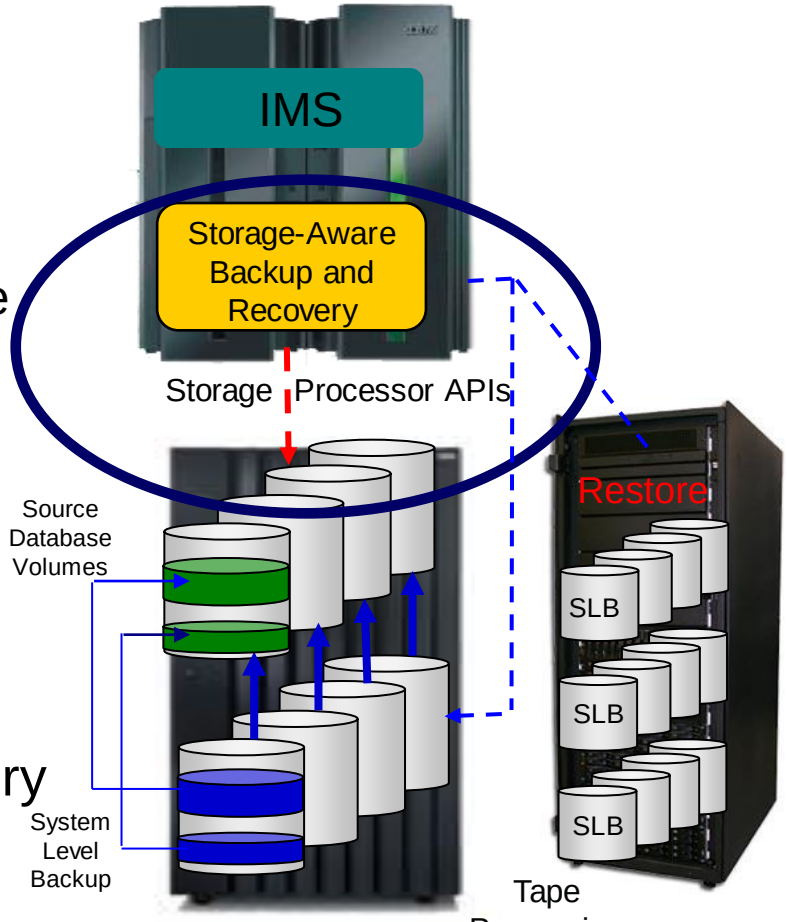

System Level Backup

Disaster Recovery

- 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 or emergency restart process
- 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
  - Disaster recovery is as simple as restarting from a power failure
- Possible tertiary DR option for sites using remote mirroring

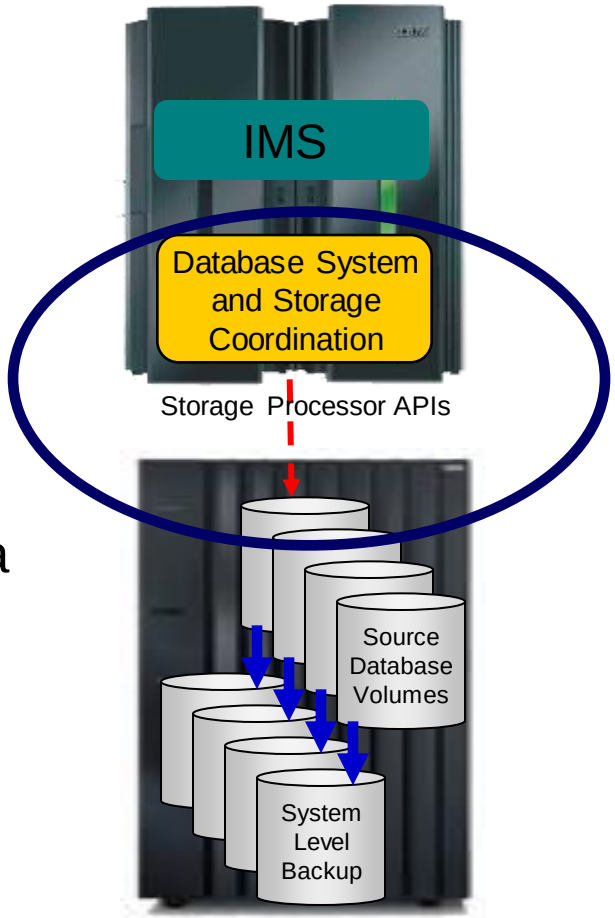

“Restartable DBMS Copy”

Integrating System Level Backups into Recovery

Intelligent Recovery Manager

- Recovers applications, individual databases, or indexes
  - To current, timestamp, or PITR
- Application profile is created in advance
  - Single database or group of databases
  - Logically related databases and indexes can be included automatically
- Determines best recovery method
  - Restores from either IC or SLB
  - Indexes that can not be restored are rebuilt
  - Recovery using log apply needs one pass of the logs
  - Access to DBs is automatically stopped and restarted at end of recovery
- Storage-based fast replication performs restore
  - Performs an instantaneous data set restore process

Intelligent Recovery Manager

Automated Recovery

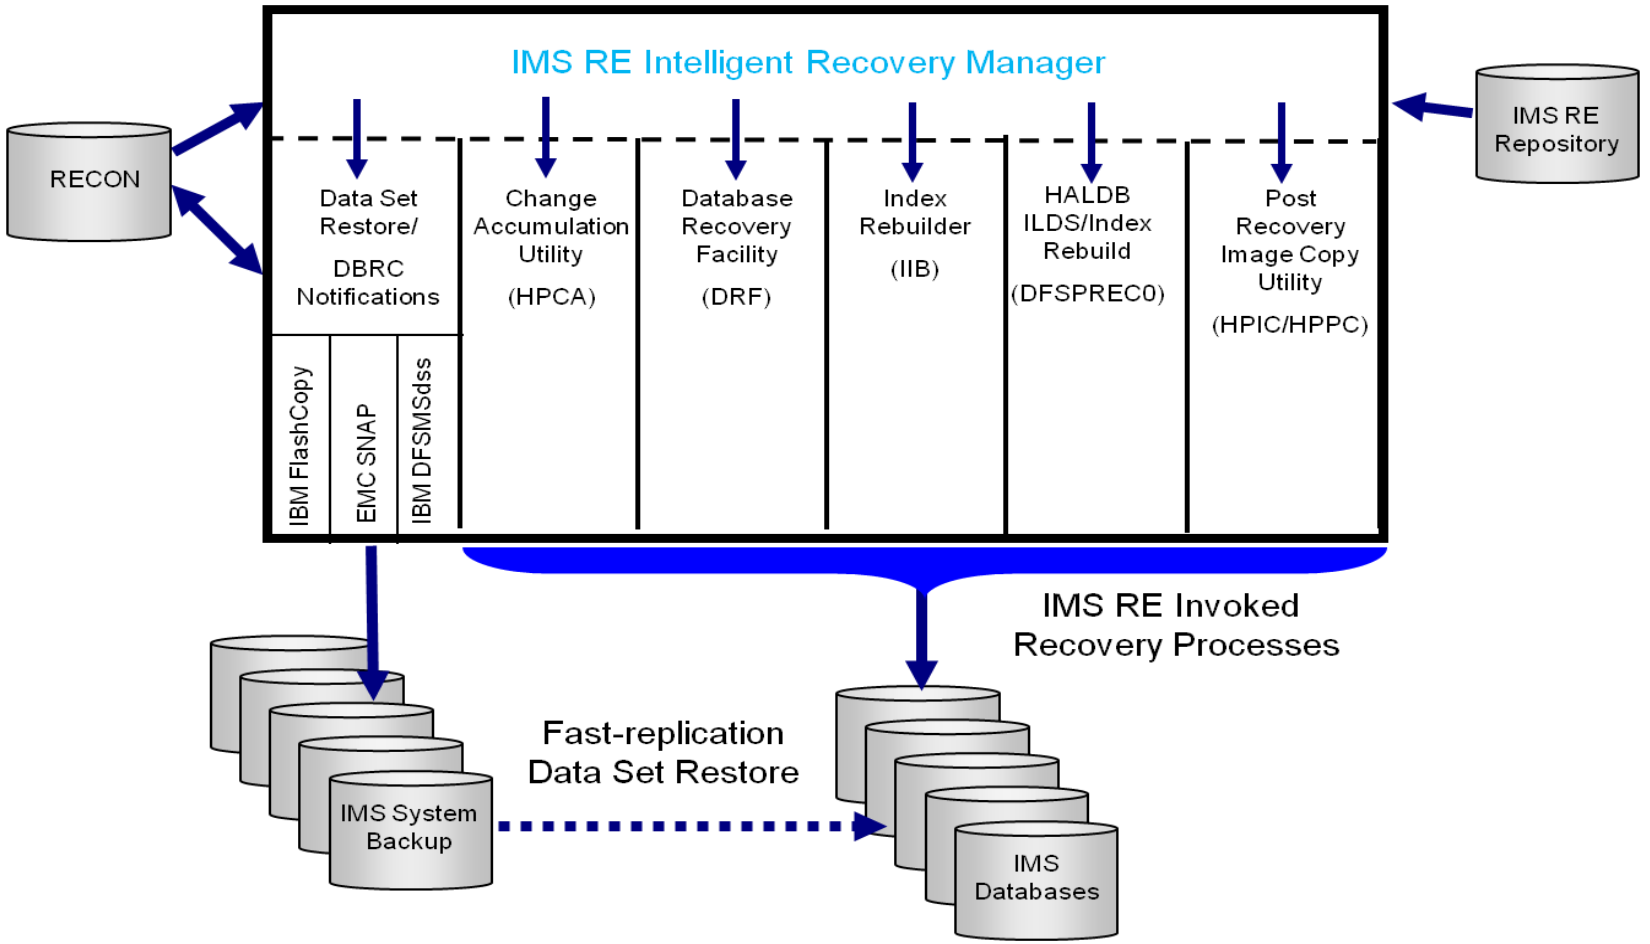

IMS Recovery Expert Main Menu

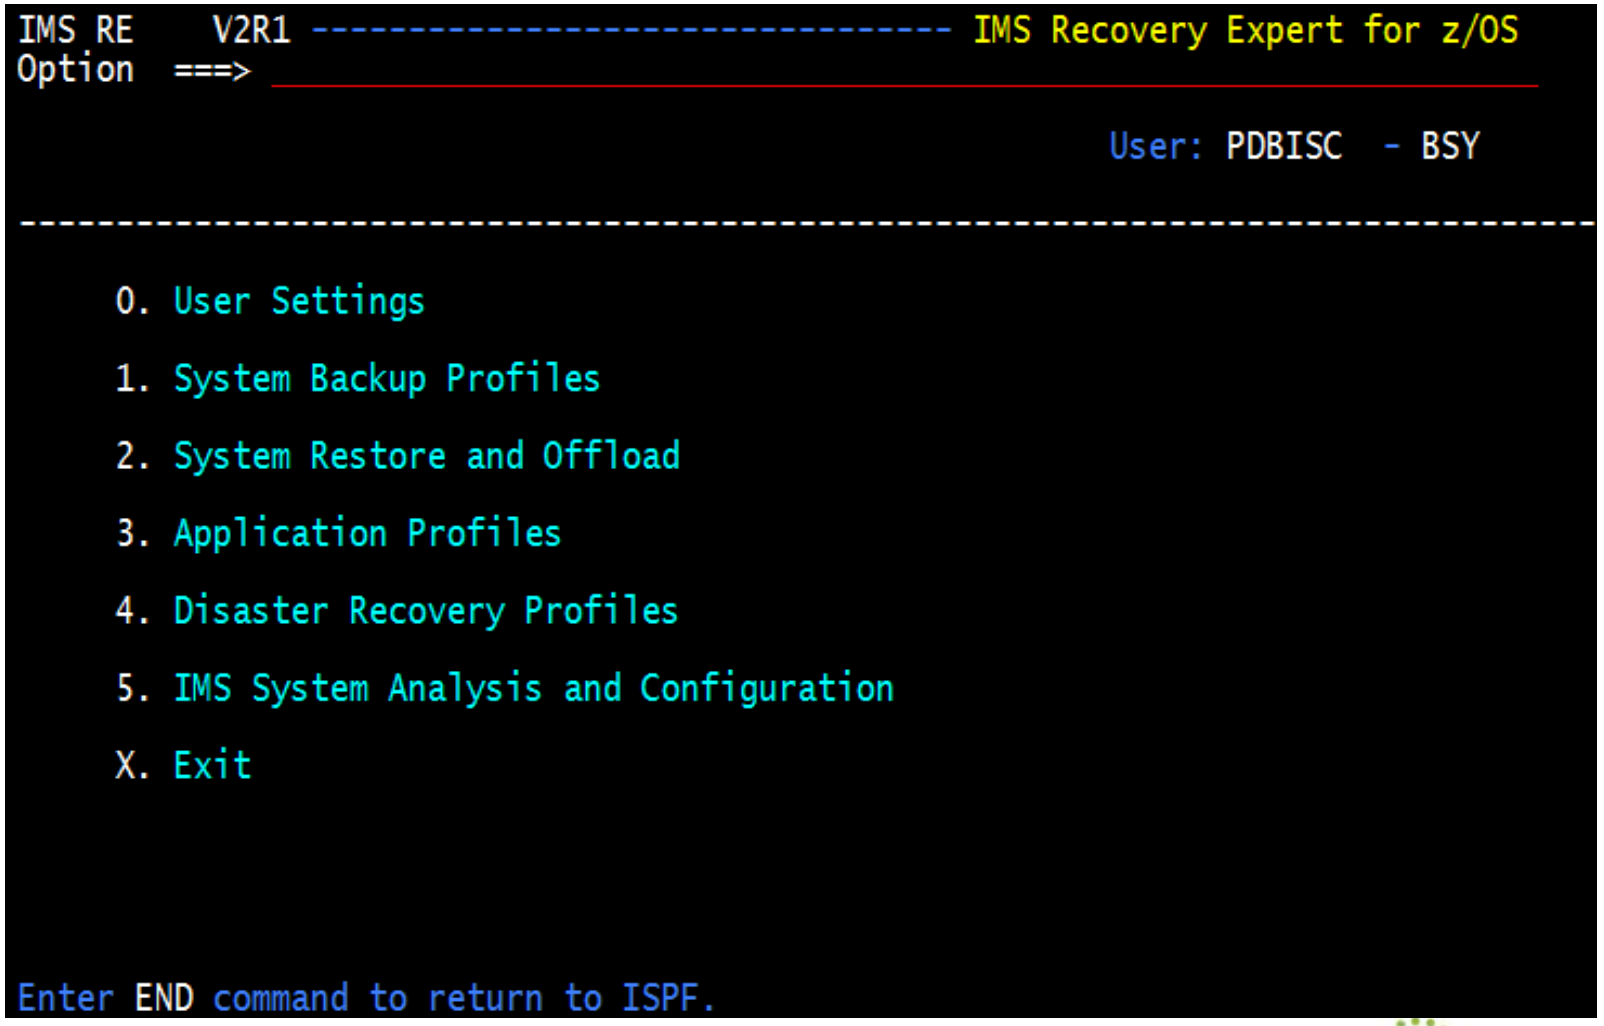

Creating a Backup Profile

```
IMS RE    V2R1 ------------------------------------ Update Backup Profile ----------
Option  ===> ____________________________________________ Scroll ===> PAGE

     Commands: ? - Show all commands
Line Commands: I - Insert  D - Delete  X - Exclude  U - Undo from exclude

--------------------------------------------------------------------------------
 Creator: PDBISC     Name: IAA - AUTO                        SSID: IAA
 Share Option: U  (Upd,View,No)    Description: AUTO MAPPING
---------------------------------------------------- Backup Options ----------
Backup Method     ==> F  (B/S/F/L)          Current Generation==> 01
Backup Scope      ==> F  (Full/Data)        Setup Needed      ==> Y
Backup Generations==> 01 (01 - 99)          Issue Log Suspend ==> N  (Yes/No)
Offload Options   ==> Y  (Yes/No/Update)    Validate IMS Vols ==> Y  (Yes/No)
Target Pool       ==> Y  (Yes/No/Update)    Enable DB Restore ==> Y  (Yes/No)
---------------------------------- Volume Inclusions/Exclusions ----------------
     Source   Dev     Src   Target
Cmd  Volumes  Type    Unit  Units     Message Area

 _   ______
********************************************************** Bottom of Data ********
```

System Level Backup Scope

- 'Full' System Level Backup
  - Includes all Databases, RECONs OLDS, Archive Logs, IMS System Data Sets, ICF Catalogs
  - Used for DR or Local Recovery
- 'Data Only' System Level Backup
  - Includes Databases Only
  - Database isolation required for DR
- 'Partial' System Level Backup
  - Contains subset of Databases
  - Can only be used for local application or database recovery
- Multiple Backup Profiles can be Defined for Each IMS
  - Define based on DR, local, and application recovery requirements

System Level Backup Scope

Partial System Level Backup

- Partial system-level backup (PSLB)
  - Backup volumes representing a subset of the IMS system
  - PSLBs used for database or application recovery only
  - Data set fast replication used to restore data
  - Desired IMS databases 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
  - Reduce disk utilization
  - Support more backup generations
- Partial SLB Following Full SLB
  - Reduce Change Accums for specific applications
  - Improve data recovery capability for specific applications

Defining Offload Options

```
IMS RE    V2R1 ------------------------------------- Offload Options -----
Option   ===>
--------------------------------------------------------------------------
 Creator: PDBISC     Name: IAA - AUTO                        SSID: IAA
 Share Option: U  (Upd,View,No)    Description: AUTO MAPPING
--------------------------------------------------------------------------
Enter the Offload options to associate with this Backup profile:

 Local Primary             ==> Y   (Yes/No/Update)
 Local Backup              ==> N   (Yes/No/Update)
 Recovery Site Primary     ==> Y   (Yes/No/Update)
 Recovery Site Backup      ==> N   (Yes/No/Update)

 Offload Generations       ==> 08  (1 - 99)
 Delete Aged Backup files  ==> Y   (Yes/No)
 Compress Data             ==> Y   (Yes/No)
 Data Mover                ==> D   (Dfsmsdss, Fdr, or fdrInstant)
 Encrypt Data              ==> Y   (Yes/No/Update)
 Number of Tasks           ==> 02  (1 - 99)
```

Benefits of SLBs
over Image Copies and Change Accums

- Creating SLB with Fast Replication is equivalent to:
  - Creating all Image Copies with < 1 second of IMS unavailable time
  - SLB created using storage processor CPU (not Host CPU)
  - Significant CPU cost savings
- Guaranteed Recoverability
  - Validation of IMS configuration each time SLB is created
- Fast Restore with Parallel Log Apply
  - Reduces recovery time and complexity
  - Executes the restore in parallel with the log apply
- Change Accumulations may not be needed
  - System Level Backups can be created frequently
  - Save host CPU and I/O
- Significantly reduce costs by using less CPU and I/O resources
  - Reduce costs to create backups
  - Save cost by reducing number of image copies needed

Implementation Planning

Backup Frequency and Space Utilization

- SLB type: Full, Data only, or Partial
  - Can be combination of different types
- Determine optimal backup frequency
- Determine number of backups to keep online (on disk)
- Determine offline (tape) backup requirements
  - eVaulting and Virtual Tape replication make SLB more accessible
  - “Stacking volumes” reduces number of tapes
- Consider incremental fast-replication options to reduce background copy time and resources
- Consider using space-efficient fast replication methods to save space
  - SLBs can be created more often
- Consider using one set of volume targets to support multiple IMS systems
  - Saves fast-replication target volume (DASD) requirement

One Set of Backup Volumes for Multiple IMS Systems

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

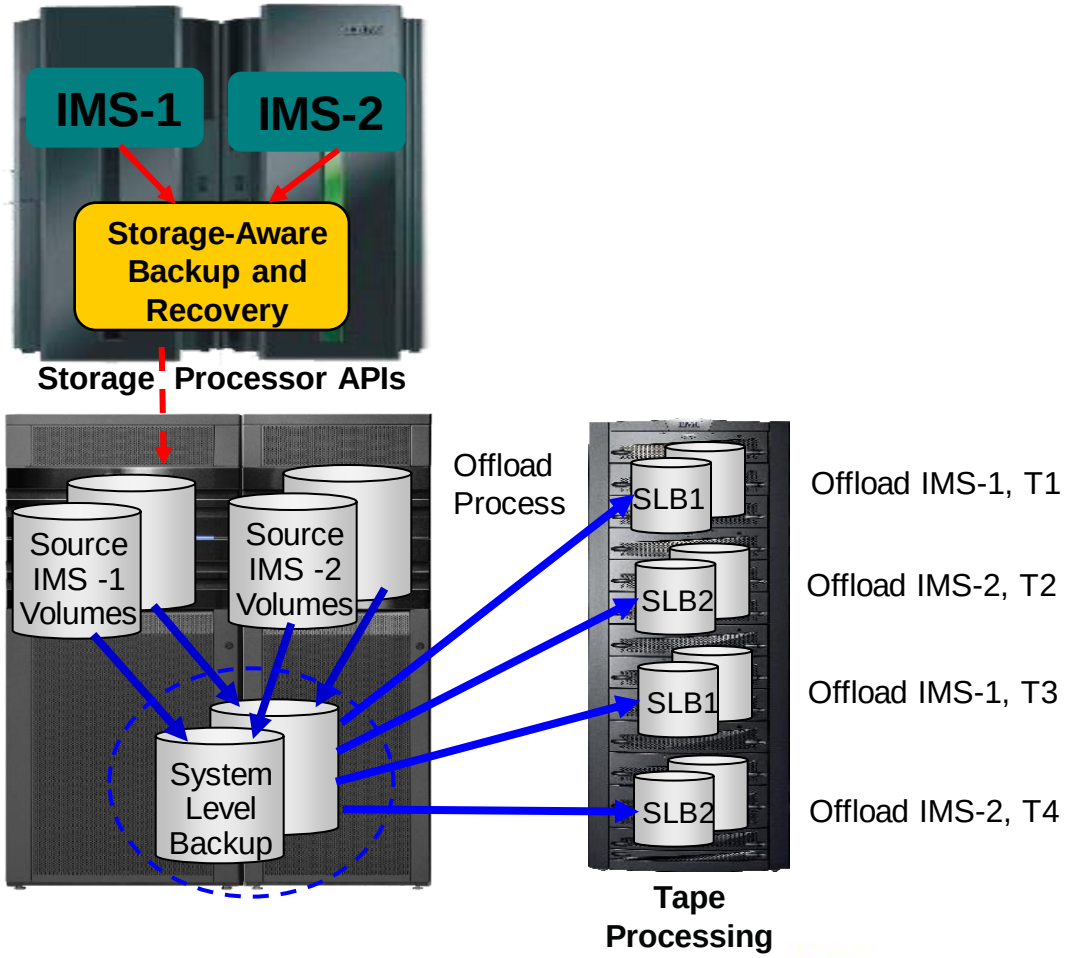

Space Efficient Devices

Operational Overview

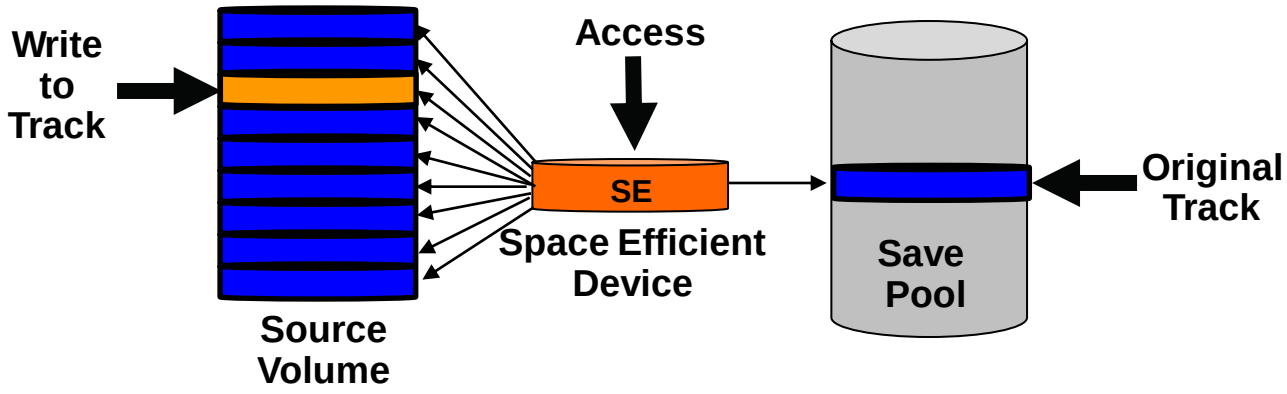

- The target is accessible when the copy session is activated
  - Pointers/Map to tracks created
- The first time a track on the source volume is written to:
  - Original data on the source volume is copied to a save volume (pool)
  - Pointer on the SE device is changed to point to the save pool
  - 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 SLBs

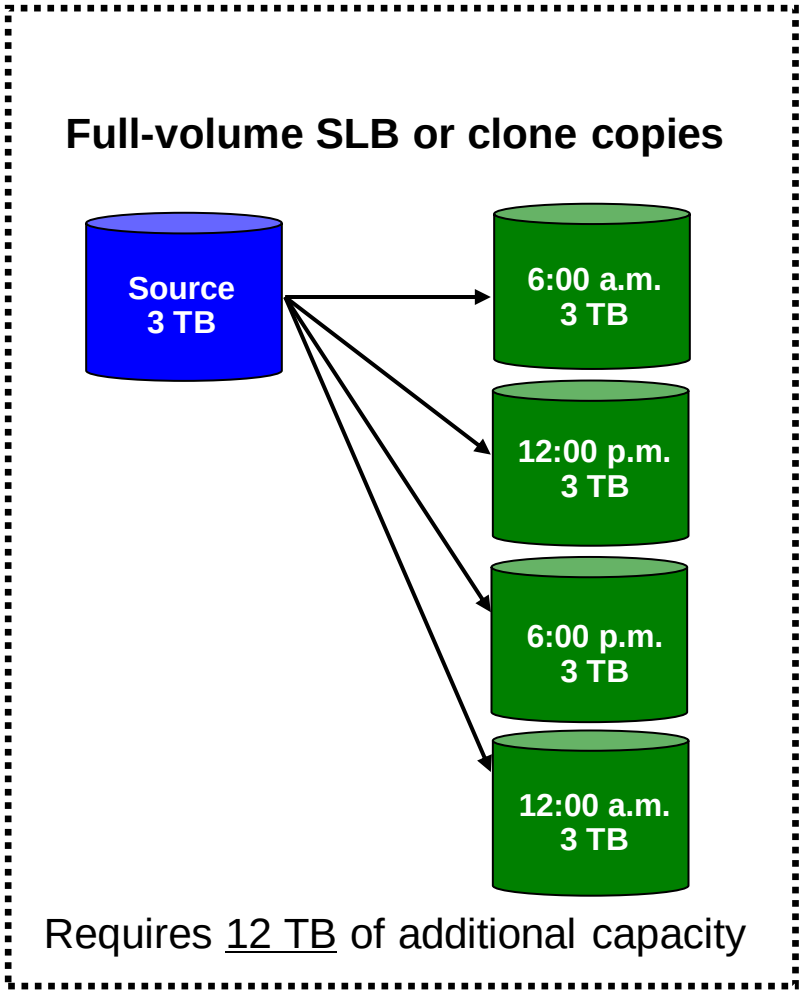

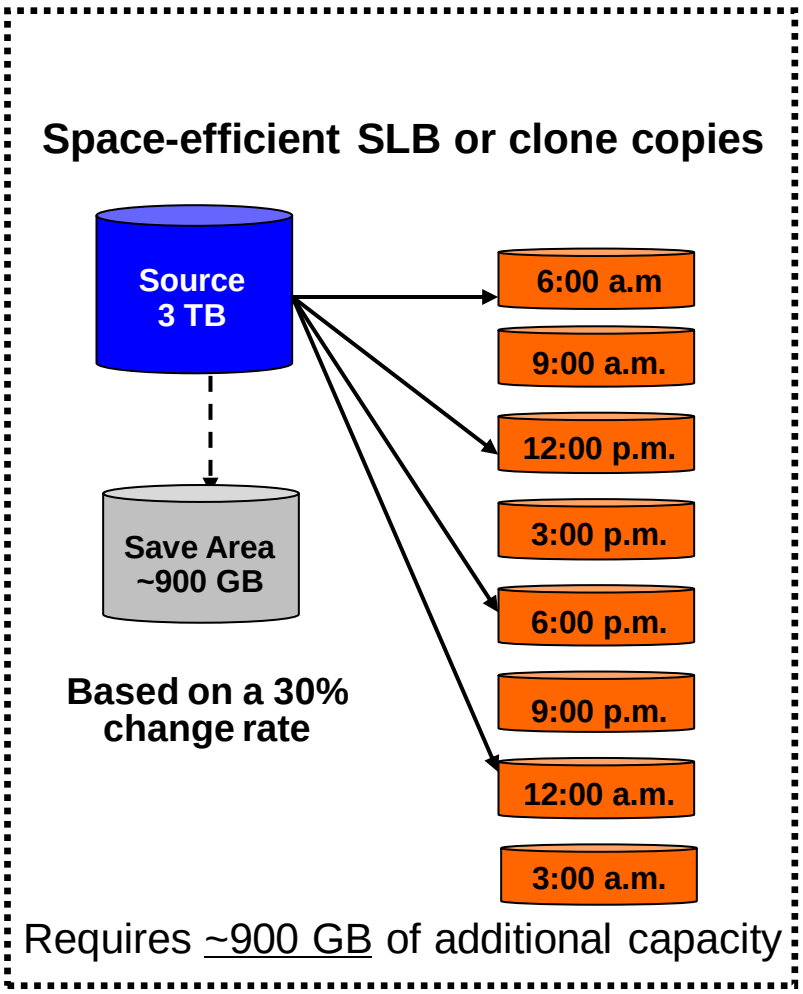

Implementation Planning

Disaster Restart

- SLB should contain IMS data only
  - Can contain other data that is restarted together
    - Recovering IMS and other data together may require using a storage based consistency function to create the SLB
    - Cannot roll forward if IMS 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 IMS REs Disaster Recovery PDS is taken offsite with archive logs and image copies
  - Reduces Recovery Point Objectives (RPO)

Using IMS SLBs for a Tertiary DR Site

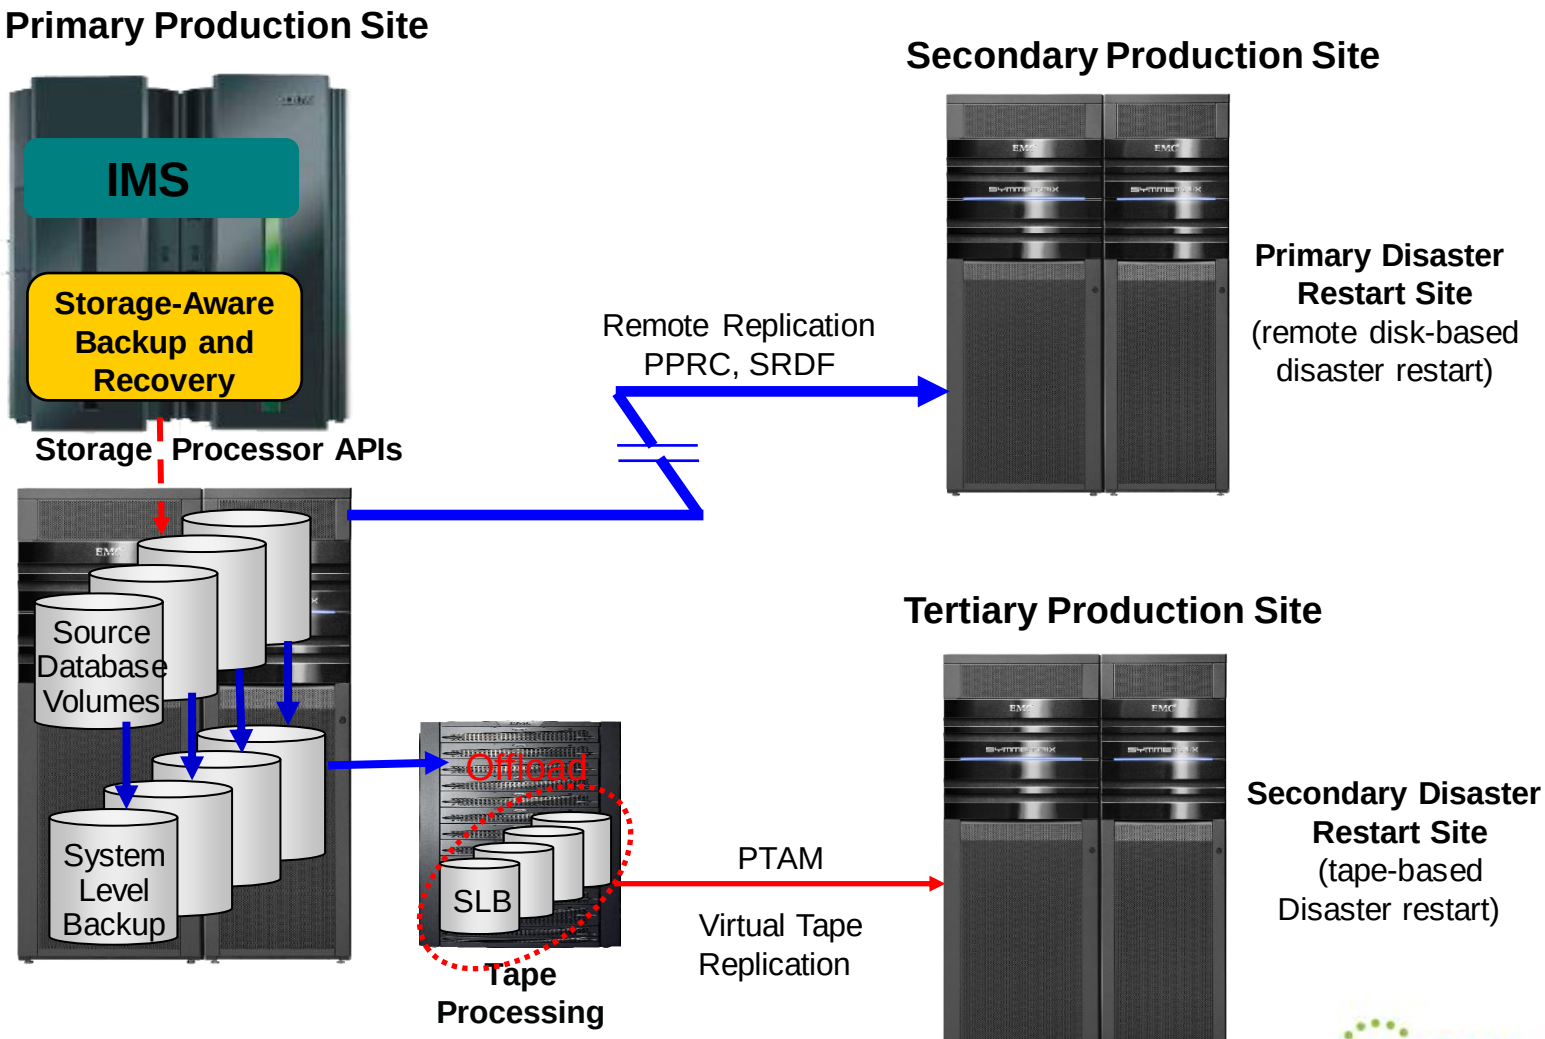

Implementation Planning

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 IMS Recovery Expert storage blade to use
  - DFSMSdss Blade
  - IBM FlashCopy Blade
  - EMC TimeFinder Blade
  - HDS ShadowImage Blade
- Know the type of consistency function is best for your environment
  - IMS RE Log Suspend
  - Storage-based consistency

IBM Copy Blades

- 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)
  - Supports IBM FlashCopy V2 storage based consistency
  - 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 DASD
  - Data set FlashCopy support for fast database / application recovery
  - Slower than using ANTRQST in native FlashCopy blade
  - Requires IMS RE 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
    - IMS RE Log Suspend not required when storage-based consistency technology is used
  - TimeFinder Data Set Snap facility to perform fast replication application / object restores.

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 – IMS RE drives FlashCopy*
      - *Y – IMS RE drives ShadowImage*
  - Incremental Copy Support
  - Requires an IMS RE Log Suspend operation
  - Can support IMS 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

Recovery Readiness Reporting Tool

- What is RRR?
  - RRR is a reporting tool that will report on a customer's potential recovery deficiencies and assess their current backup resource usage
  - It will expose those DB2 or IMS resources that may be missing backups - causing recoveries to fail
  - Recovery resources are analyzed for a given DB2 subsystem or IMS Environment

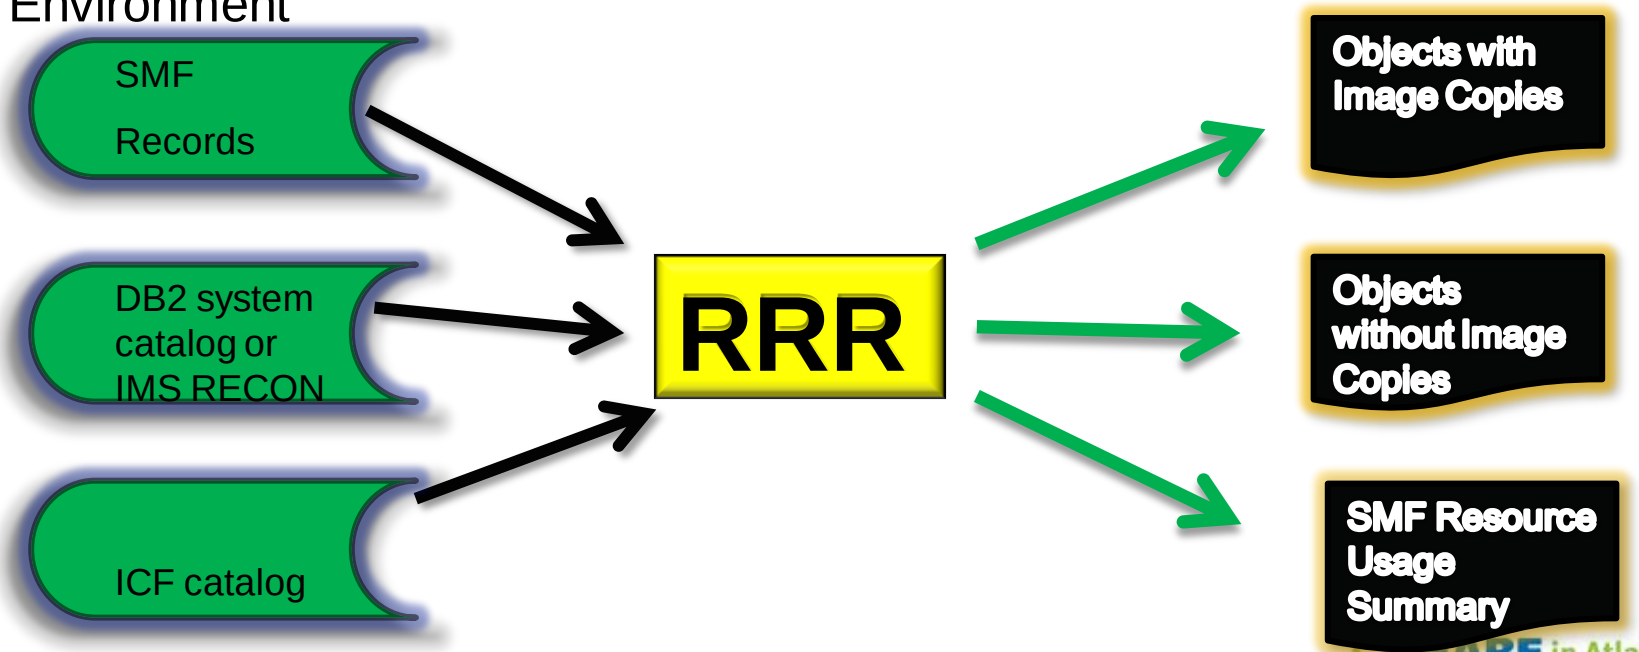

Recovery Readiness Reporting Tool

- Shows Databases that do not have a good image copy within a given time frame

Recovery Readiness Image Copy Report
Start Timestamp 2011-09-11-00.00.00.000000
End Timestamp 2011-09-17-23.59.59.999999
Objects without Image Copies

DBNAME	Area/Part	Type	DDNAME	Recoverable
DBFSAMD3	CUSDB	DEDB		Y
DBFSAMD4		FF	LOAN	Y
DI21PART		FF	DI21PARO	Y
DI21PART		FF	DI21PART	Y

- Show system resources to create backups within a given time frame

Image Copy Resources for RECON
Total IC Steps: 9158
CPU: .....10891 (secs)
EXCP: 163110434
Tape Mounts: 9084
Elapsed Time: 48:48:27

Change Accumulation Resources for System
Total CA Steps: 2126
CPU: .....24162 (secs)
Elapsed Time: 83:59:21

Database and Storage Integration

New Solutions for DBAs to Consider

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

Session Summarization

- IMS storage-aware database utilities provide storage integration to simplify database administration tasks
- IMS 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 database recovery time
- Fewer skills required to implement advanced IMS backup, recover, disaster recovery, and cloning solutions
- Implementation planning is important to optimize the benefits

Trivia Answer

- 3390 bytes/track = 56664
- 36 3390 mod 54s = 2 TB
  - 65,520 cyls, 55.68 gb per volume
- 72 3390 mod 27s= 2 TB
  - 32,760 cyls, 27.84 gb per volume
- 236 3390 mod 9s = 2 TB
  - 10017 cyls, 8.51 gb per volume
- 705 3390 mod 3s = 2 TB
  - 3339 cyls, 2.84 gb per volume
- 2115 3390 mod 1s = 2 TB
  - 1113 cyls, 946 mb per volume
- According to Wikipedia…

IBM 3380 – 1GB Disk Drive

The IBM 3390 Direct Access Storage Device series was introduced November 1989, offering a maximum storage of up to 22 gigabytes in a string of multiple drives. Cost of a storage system varied by configuration and capacity, between $90,000 and $795,000. A Model 3 enhancement to the drive family, announced September 11, 1991, increased capacity by approximately 1.5 and a Model 9 family, announced May 20, 1993, further increased capacity by an additional factor of 3 to minimum capacity of 34 gigabytes in a single drive box.