IBM GLOBAL SERVICES

B79

HALDB Database Administration

Rich Lewis

IMS
Technical Conference

Sept. 27-30, 2004

Orlando, FL

RETURN TO INDEX

Agenda

- Review of database data sets
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test databases
- Secondary indexes

- A more extensive version of this presentation including notes for each page is available on the web at:
  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS842
- Redbook: The Complete IMS HALDB Guide, All You Need to Know to Manage HALDBs
  SG24-6945

HALDB Database Data Sets

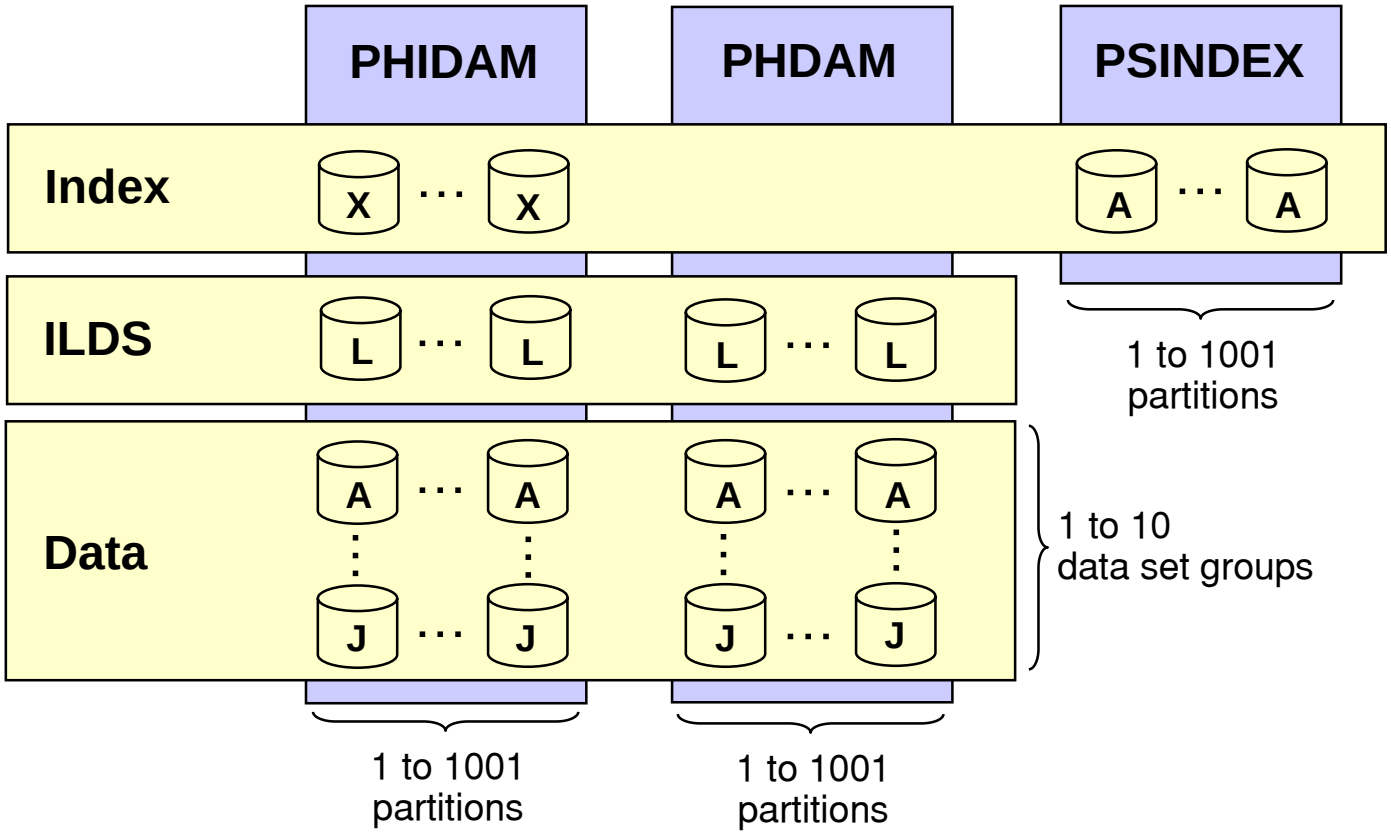

The data sets in a partition have generated data set names and DDNAMEs.

Letters are used to distinguish them.

X : PHIDAM primary index

A – J : Data data sets

L : ILDS

A : PSINDEX

HALDB Database Data Sets

- Each PHDAM or PHIDAM partition requires an ILDS (L)
  - ILDS is empty if there are no sec. indexes or log. Relationships
- Each PHIDAM partition has an index data set (X)
- Each PHDAM or PHIDAM partition has an A data set
  - Root segments are in the A data sets
- Each PHDAM or PHIDAM partition may have B-J data sets
  - Used for multiple data set groups
- Each PSINDEX has an A data set

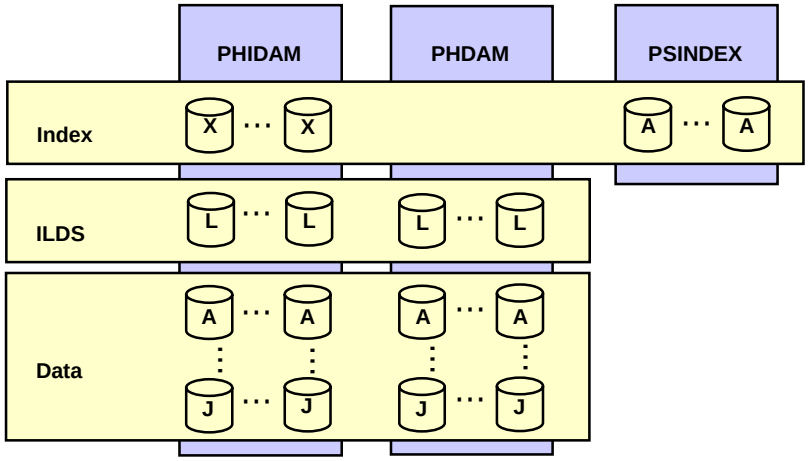

Partition Names and IDs

- Each partition has a name
  - Unique in RECONs
    - Partitions in different databases cannot have the same names
    - Partitions cannot have the same names as databases
  - Choices:
    - Name signifies the data in the partition
      - Could cause problems when partitions are modified
    - Name is arbitrarily chosen
- Each partition has an ID
  - Number assigned by IMS when partition is defined
    - Assigned in creation order within the database
    - Not in key sequence
    - Not reused

HALDB Database Data Sets

- Data set names
  - Begin with data set name prefix for the partition
    - Up to 37 characters assigned by the user
  - Letter and Partition ID are used as suffix
    - X for PHIDAM index
    - L for ILDS
    - A for PSINDEX
    - A through J for data
  - Example:
    - Partition data set name prefix: IMSP0.DB.INV23
    - Partition ID: 00004
    - Data set names:
      - PHIDAM index: IMSP0.DB.INV23.X00004
      - PHIDAM ILDS: IMSP0.DB.INV23.L00004
      - PHIDAM first data data set: IMSP0.DB.INV23.A00004

- Each partition in a database may have the same data set name prefix.
  - Partition IDs make data set names unique

HALDB DDNAMEs

- DDNAMEs
  - Begin with the partition name
    - Up to 7 characters assigned by user
  - Letter is used as suffix
    - X for PHIDAM index
    - L for ILDS
    - A for PSINDEX
    - A through J for data
  - Example:
    - Partition name: LBAD112
    - DDNAME for PHIDAM index: LBAD112X
    - DDNAME for PHIDAM ILDS: LBAD112L
    - DDNAME for first data data set: LBAD112A

Dynamic Allocation

- Dynamic allocation uses RECON information
  - All HALDB databases are registered in RECONs
  - DFSMDA members are never used for HALDB
- If you use a DD statement:
  - If DD statement conflicts with RECON information, allocation fails
  - If DD statement matches RECON information, allocation succeeds
    - It works as if you had not used a DD statement
- THEREFORE, do not include DD statements for HALDB

Agenda

- Review of database data set
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test databases
- Secondary indexes

Partition Initialization

- Partition initialization
  - Prepares partition data sets for use
  - Ensures that partitions with no data are usable
  - Initialization is done by
    - HALDB Partition Data Set Initialization utility (DFSUPNT0)

      or
    - Database Prereorganization utility (DFSURPR0)
  - Database is specified to the utility
    - Partitions with ‘partition initialization required’ DBRC flag (PINIT) are initialized
      - Exception: DFSUPNT0 has unconditional partition initialization function
        - Invokes initialization for all partitions in the database with or without flag set
        - Specified with INITALL statement in DFSOVRDS DD data set
        - Introduced b PQ49638 (IMS V7) and PQ55002 (IMS V8)

Partition Initialization

- Partition initialization process
  - Makes high-used RBA non-zero
    - Writes and erases a record in PSINDEX
  - Writes reorg number and partition ID in PHDAM and PHIDAM
  - Creates first bit map block in PHDAM and PHIDAM
  - Writes high-key (x'FF…FF') record in PHIDAM

Partition Initialization

- Partition initialization is only required in three cases:
  1. Before initial load (PROCOPT=L) of partition
  2. Before migration reload of partition
     - Input to reload was created by unload of non-HALDB database with MIGRATE=YES or MIGRATX=YES option
  3. Before a partition may be used without containing any data
     - Initial load or reload does not insert any segments in the partition
- Partition initialization is not required with reorganizations
  - Not required even when data sets are deleted and redefined
    - Unless the partition is empty
- 'Partition Initialization Required' flag in RECONs
  - Turned 'on' by partition definition or DBRC command
  - Turned 'off' by partition initialization or DBRC command
  - Authorization fails if flag is 'on'

Number of Partitions and Their Sizes

- Things to consider when choosing the number of partitions
  - Number of partitions affects the sizes of partitions
  - Time required to reorg partitions in parallel
    - Smaller partitions shorten the process
  - Time required to image copy and recover partition data sets
    - Smaller partitions shorten these processes
  - Smaller partitions may avoid multivolume data sets
    - Especially important with OSAM
  - Management of the data sets
    - More data sets require more management
- Multiple data set groups
  - May be advantageous to have only one data set per partition
  - May be advantageous to have multiple data sets per partition

Multiple Data Set Groups

- HALDB supports multiple data set groups
  - Multiple data set groups place different segment types in different data sets
  - Should you use them?
- Multiple data set groups were used for two reasons with non-HALDB
  1. Avoid data set size limitations
     - Not required with HALDB
  2. Place infrequently used segments in another data set
     - Also applies to HALDB

- Example:
  - Place D segments in a data set group
  - Increases likelihood that E will be in the same block with A.

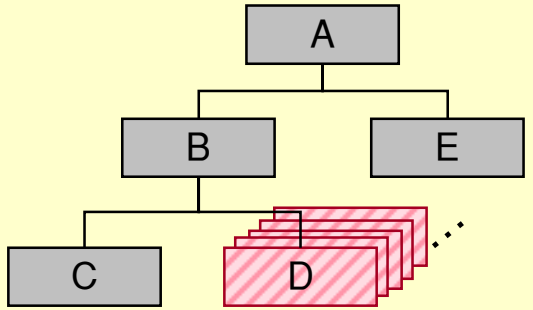

Database Compression

- HALDB supports segment edit/compression routines
  - Should you use them for compression?
- Reasons to use compression with HALDB
  - Saves DASD space
  - May improve performance
    - Reduces I/Os required to retrieve and write data
- Reasons not to use compression with HALDB
  - Not needed to avoid the data set size limitation
  - May hurt performance
    - CPU costs for compression and expansion of segments
  - DBA effort to manage compression routines

Adding, Deleting and Changing Partitions

- Databases change over time
  - The sizes of partitions may change over time
    - Data added or deleted
  - The high keys of partitions may need to be adjusted over time
    - Different amounts of data added or deleted to different partitions
      - Example: Root keys based on date
- Databases need to be adjusted over time
  - Partitions may need to be split, consolidated, created, or deleted
  - Partition boundaries (high keys) may need to be adjusted

HALDB Migration Aid Utility

- HALDB Migration Aid utility (DFSMAID0)
  - Reads HDAM, HIDAM Secondary Index databases
    - Provides sizing and high key information for migration planning
    - Secondary index support
      - Provides key range boundaries and numbers of records
        - Reports of ‘bytes’ and ‘prefix-incr’ information are inaccurate for secondary indexes
        - Number of segments and high key values are accurate in the report
        - Sizes are easily calculated from the number of records
    - Read PHDAM, PHIDAM, and PSINDEX databases
      - Provides sizing and high key information for repartition planning

HLADB Migration Aid Utility

- Sample report:

```
partition 1 :

   minimum key =

      +0000  d2c1c1f1 f1f2f3f4     |KAA11234 |

   maximum key =

      +0000  d2f2f3f9 f9f2f3f4     |K2399234 |

                    segments         bytes   prefix-incr   length-incr
   1) 'PRODUCT '       31567       4040576        252536             0
   2) 'INVENT'        103781       8094918        830248             0
   3) 'ORDQTY'        171182      10955648       1369456             0
   4) 'MFGSPECS'       51115      10938610        408920             0
 SUM)                 357645      34029752       2861160             0
```

segments - number of segments
bytes - number of bytes for the segments
prefix-incr - additional bytes due to increased prefix size
length- incr - additional bytes required for paired logical relationships

HALDB Migration Aid Utility

- Using the HALDB Migration Aid utility
  - You may specify one of the following
    - Number of equal sized partitions
    - Number of bytes per partition
      - Bytes of segment data
      - Does not include free space, bit maps, RAPs or FSEAPs
    - High keys for partitions
  - Report for each partition and the entire database
  - Bytes in the reports do not include free space, bit maps, RAPs or FSEAPs
    - You must adjust for these

Splitting a Partition

- If partition B with high key 4000000 needs to be split
  - Unload partition B
    - HD Unload or HP Unload
  - Define new partition E
    - Initialize partitions B and E
  - Reload partitions B and E
    - HD Reload or HP Load
  - Partitions A, C, and D are not affected

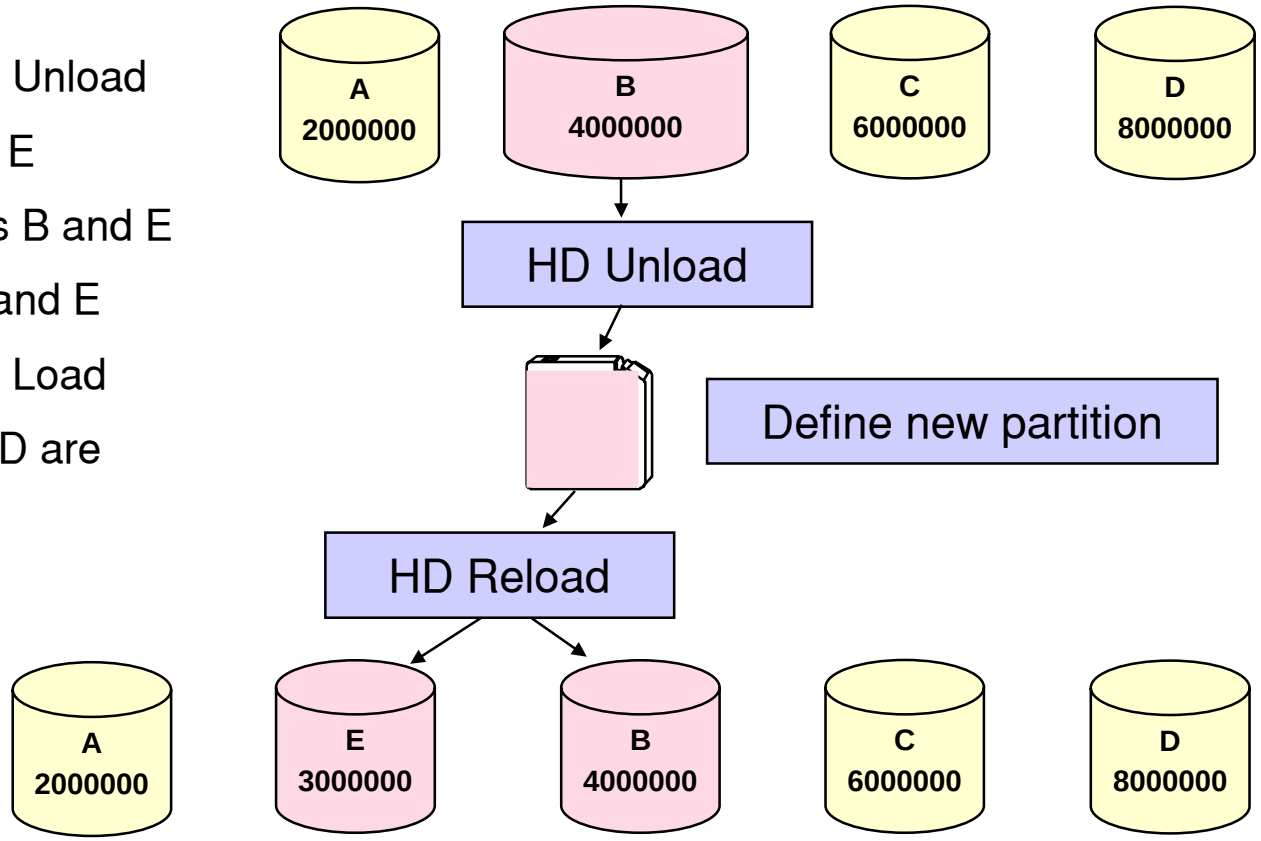

Modifying Partition Boundaries

- If records need to be moved from partition B to C
  - Unload partitions B and C
    - HD Unload or HP Unload
  - Change high key for partition B
    - From 4000000 to 3000000
  - Reload partitions B and C
    - HD Reload or HP Load
  - Partitions A and D are not affected

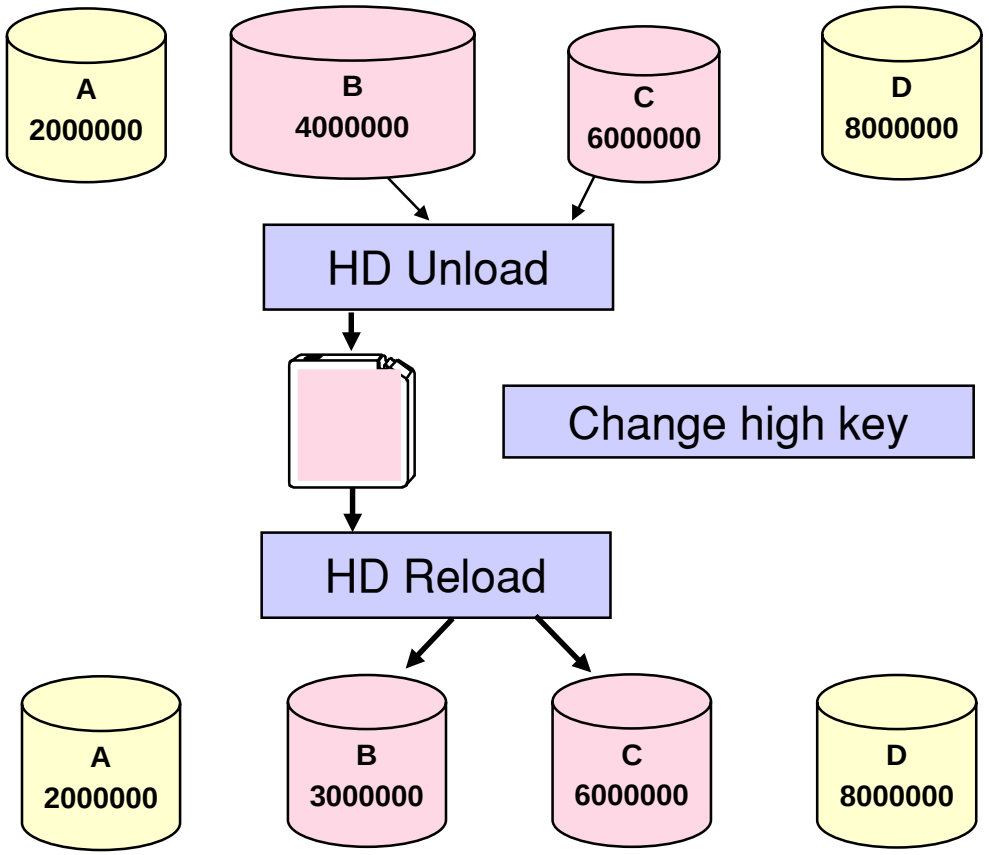

Databases with Dates for Keys

- Some databases have dates as the high-order part of the key

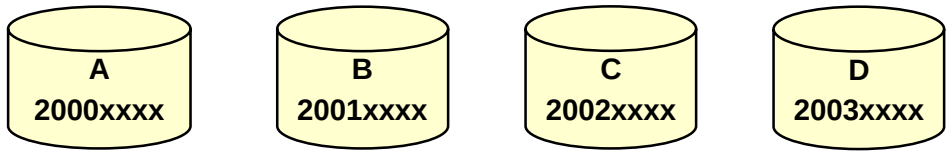

  - To add a partition for a set of dates (higher keys)
    - Define it and initialize it

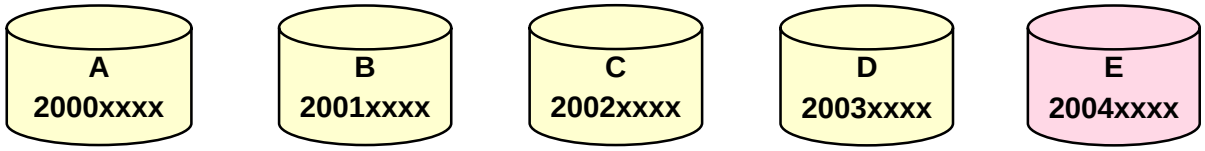

  - To delete the partition with the lowest dates (keys) and all of its data
    - Delete the partition definition

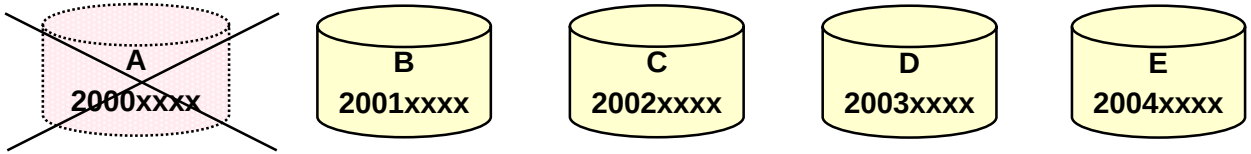

- Unloads and reloads are not required for these changes

Disabling and Enabling Partitions

- Disabling and Enabling partitions was introduced by APARs
  - PQ48421 for IMS V7
  - PQ73858 for IMS V8
- Disabling partitions
  - Definitions and information remain in RECONs
    - Includes partition IDs, DSN prefixes, and recovery information
  - Partitions are not used
    - Partitions are ignored
- Disabled partitions may be enabled
  - Enabled partitions are made active
  - Enabled partitions are marked ‘recovery needed’

Disabling and Enabling Partitions

- Use of disabling and enabling partitions
  - Disabling is normally done prior to deleting a partition
    - Keeps recovery information, partition ID, DSN prefix, etc.
  - If testing is successful, partition is deleted
    - Deletion removes all information
  - If testing is not successful, partition is enabled
    - Partition is recovered and becomes active
      - Other partitions may require timestamp recovery
- Partition Definition Utility (PDU) support
  - New 'Partition status' flag on 'Change Partition' panel
- DBRC command support:

```
CHANGE.PART DBD(dname) PART(pname) DISABLE
CHANGE.PART DBD(dname) PART(pname) ENABLE
```

Agenda

- Review of database data set
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test databases
- Secondary indexes

Non-HALDB Reorganizations

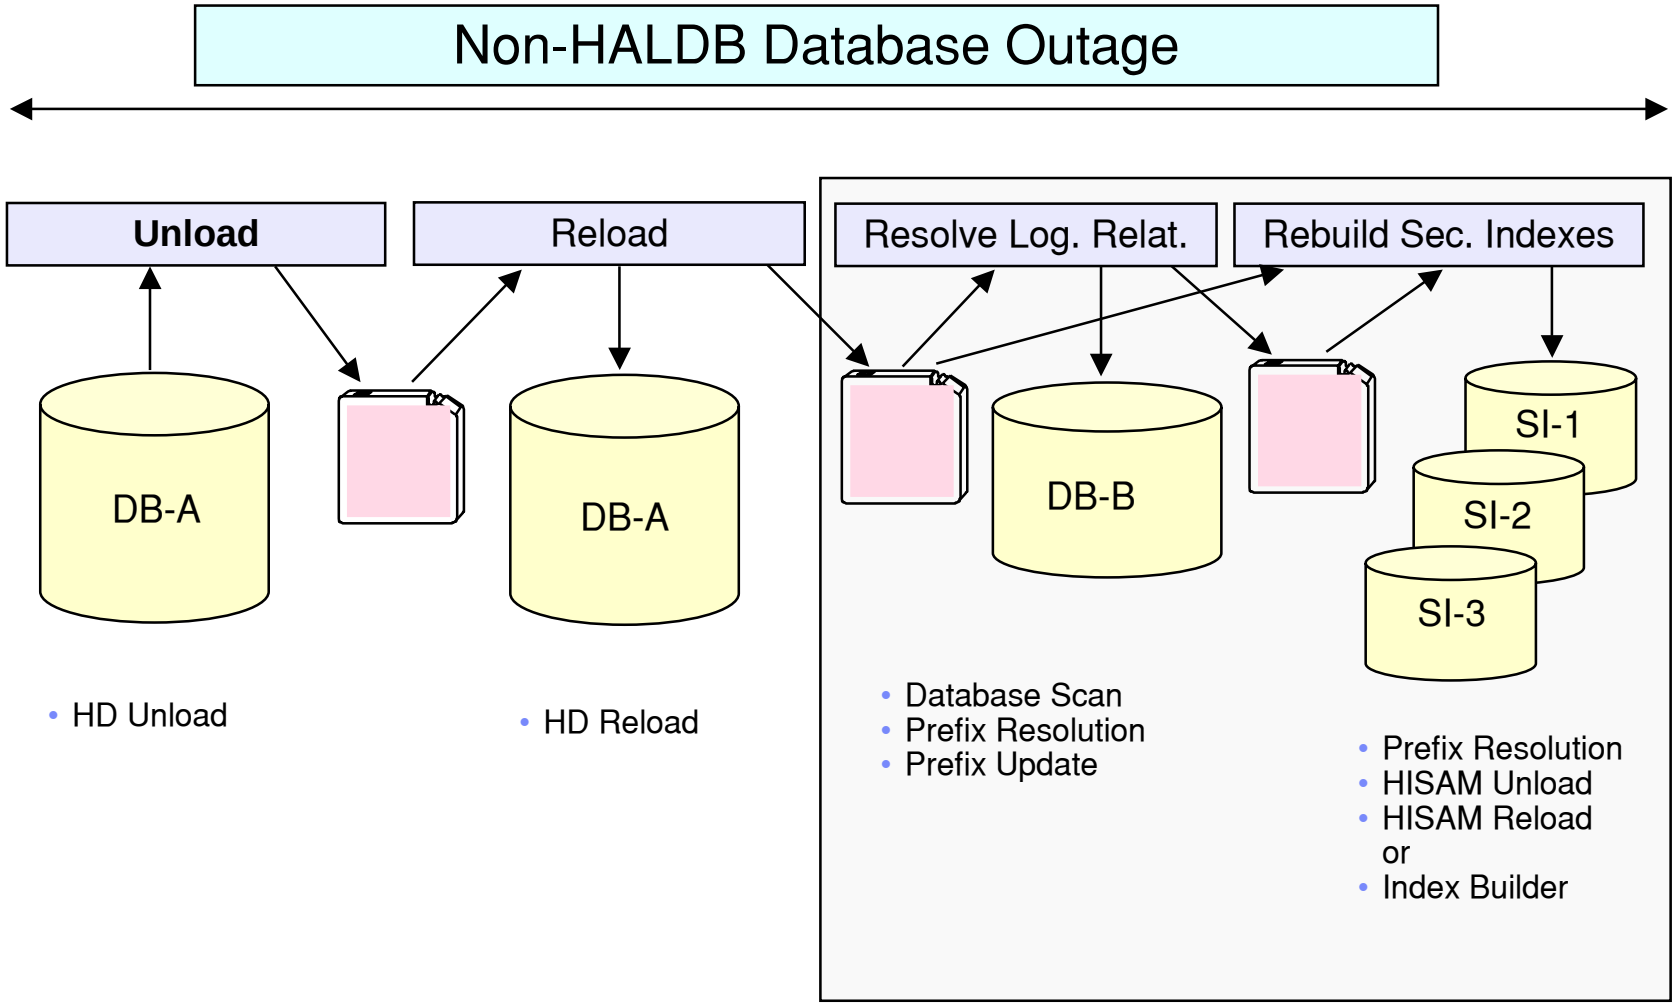

HALDB Reorganizations

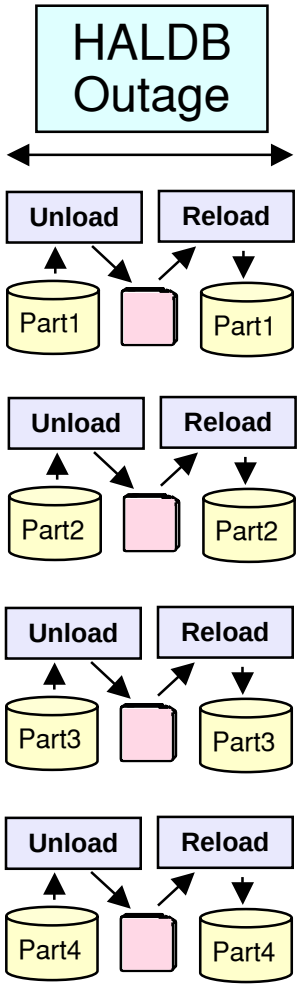

- Shorten the reorg time to your window
  - Create enough partitions to meet your requirements
- Reorg partitions in parallel
- Eliminate rebuilds of secondary indexes
  - Prefix Resolution, HISAM Unload, and HISAM Reload, or Index Builder are not required
- Eliminate updates to logical relationships
  - Database Scan, Prefix Resolution, and Prefix Update are not required

Healing Pointers After Reorgs

- After a reorganization sec. index and log. rel. pointers are "broken"
  - Normal processing heals them efficiently
    - Only heals pointers that are used
    - Reads of pointers are "free"
      - They are being read for normal processing
    - ILDS reads are efficient
      - ILDS CIs hold many entries
      - ILDS CIs are maintained in the buffer pools
  - Optionally, you can heal them
    - Extends the reorganization process
    - Typically, uses more resources
      - Heals all pointers
      - More total I/Os
    - HALDB Conversion and Maintenance Aid includes pointer healing utility
- My recommendation: Let normal processing heal the pointers

Data Set Delete and Define for Reorgs

- HALDB database data sets may be reused
  - Delete and redefine are not required for reorganization
    - VSAM REUSE attribute is honored by HD Reload
      - Non-HALDB VSAM required DELETE and DEFINE
    - OSAM allows reuse with both HALDB and non-HALDB
  - Delete and define are required to move data sets
- REUSE attribute is required for HALDB VSAM data sets
  - Except ILDS
    - Parameter is allowed but not honored for ILDS
      - ILDS will not be reused by Index/ILDS Rebuild utility (DFSPREC0)

Partition Initialization During Reorgs

- Partition initialization is not required during reorganizations
  - Data sets may be deleted and redefined without partition initialization
    - Exception: A partition which contains no data must be initialized

- Reorganization steps:
  - Unload partition
  - Delete partition data sets (optional)
  - Define partition data sets (optional)
  - Reload partition

Reorganizations and Secodary Indexes

- Reorganization of a HALDB database does not require rebuild of its secondary indexes
  - Self-healing pointer scheme eliminates this requirement

Good!

- Many installations never reorganize non-HALDB secondary indexes
  - They are rebuilt (and organized) with every reorg of the indexed databases

- HALDB secondary indexes may become disorganized
  - They may require reorganization

This is a change in procedures

Reorganization Alternatives

- HD Unload and HD Reload
  - Standard IMS utilities
- High Performance Unload and High Performance Load
  - IMS tools
- IMS Parallel Reorg
  - IMS tool
  - Unload and reload are done in parallel
- IMS V9 Online Reorganization
  - Standard IMS utility
  - Absolutely no outage for reorganization

Agenda

- Review of database data set
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test databases
- Secondary indexes

Backup and Recovery

- HALDB A-J data sets (not the ILDS or PHIDAM index)
  - Standard IMS full function processes
    - Image Copy, Image Copy 2, Online Image Copy, tools
    - Logging
    - Change Accumulation
    - Database Recovery utility
    - Database Recovery Facility (DRF) tool
    - DBRC support
      - GENCJCL.IC
      - GENJCL.CA
      - GENJCL.RECOV

Backup and Recovery

- ILDS (L) and PHIDAM index (X) data sets
  - Backup
    - No image copies
  - Updates are not logged
    - ILDS is only updated by reorganization
    - PHIDAM Index is treated like a non-recoverable database
  - Recovery
    - Index/ILDS Rebuild Utility (DFSPREC0)
      - Rebuilds the data set(s) from the database
  - DBRC
    - GENJCL.USER MEMBER(DSPUPJCL)
      - May be used to generate DFSPREC0 JCL to rebuild ILDS or PHIDAM index

Timestamp Recoveries

- All data sets in a partition must be recovered to the same time
  - PHIDAM index must be rebuilt
    - Rebuilt from the “A” data set with Index/ILDS Rebuild utility
  - ILDS may need to be rebuilt
    1. If secondary indexes or logical relationships are used
       and
    2. If recovery is to time before last reorganization
       - ILDS is only changed by reorganizations
    - ILDS may be rebuilt wit the Index/ILDS Rebuild utility
  - Alternative for ILDS
    - After reorganization
      - Copy ILDS with REPRO
    - If ILDS needs to be restored
      - Use copy produced by REPRO

Timestamp Recoveries

- Must all partitions of a database be recovered to the same time?
  - Almost always
  - User must understand when this is not required
    - For example, offending program updated only one partition
  - Secondary index implications
    - Usually, database with secondary index forces recovery of all partitions to the same time
      - All partitions of the indexed database
      - All partitions of the secondary indexes
  - Logical relationship implications
    - Usually, database with logical relationships forces recovery of all partitions to the same time
      - All partitions in the logically related databases

Agenda

- Review of database data set
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test Databases
- Secondary indexes

Test Databases

- Non-HALDB test databases
  - Often, not registered in RECONs
  - Each programmer may have one or more versions of a database
- All HALDB databases are registered in RECONs
  - Multiple versions of a database must be defined in different RECONs
    - DBRC does not allow multiple databases with the same name
  - Multiple test versions of a database require multiple RECONs
  - Plan your batch test environments

Defining Test Databases

- Use the same DBD as production
  - DBD does not include partition or data set information
  - Place in test DBDLIB and ACBLIB
- Create test partition definitions
  - Define partitions for test environments

  or
  - Use Partition Definition Utility EXPORT and IMPORT functions
    - Moves partition definitions between RECONs
      - They may be modified after IMPORT
        - Data set name prefix, RAA, etc.
      - APARs PQ48421 (V7) and PQ73858 (V8) maintain partition IDs

Creating Test Databases

- Alternatives for creating test database from a production database
  - Unload and reload
    - HD Unload (HP Unload) of production
    - HD Reload (HP Load) to test
      - You may create a different partition configuration
        - Partition IDs will generally be different
        - Partition names may be changed
        - Partition boundaries may be changed
  - Image Copy and restore
    - Export and import partition definitions
      - Maintains partition IDs (with APARs PQ48421 and PQ73858)
    - Image copy production databases and restore to test
      - Partition IDs are stored in database data sets
    - Change database data set names of test database
      - Change data set name prefixes
  - Use application programs

Agenda

- Review of database data set
- Partitions
  - Initialization
  - Sizing
  - Adding, deleting, and modifying partitions
- Reorganizations
- Recoveries
  - Timestamp recoveries
- Test databases
- Secondary indexes

Secondary Indexes

- Plan the partitions for the secondary indexes
  - How many partitions do you need?
    - Space requirements
      - HALDB secondary index entries are much larger than those for non-HALDB secondary indexes
        - Pointers are larger
        - Root key of target is stored in the entry
    - Will they need to be adjusted during the life of the database?
      - Keys based on date, etc.
- Plan to reorganize them
  - They are not rebuilt with each reorganization of their indexed database
- Don't make the non-recoverable
  - Unless you have a tool to rebuild them (e.g. Index Builder)
  - They are not rebuilt by IMS utilities

Some Things to Remember

- HALDB Migration Aid utility can analyze existing HALDB databases
  - Useful when planning repartitioning
- Deleting a partition definition deletes its recovery information
  - Disabling a partition keeps its recovery information
- Secondary indexes may require reorganizations
  - They are not rebuilt when the indexed database is reorganized
- Secondary index cannot be rebuilt from database with IMS utilities
  - Don't make them non-recoverable unless you have a tool like the IBM Index Builder
- PHIDAM indexes and ILDSs have a different recovery process
  - They are rebuilt with Index/ILDS Rebuild Utility (DFSPREC0)
- Plan your scheme for creating HALDB test databases
  - DBRC registration is required for all databases

A more extensive version of this presentation including notes for each page is available on the web at:
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS842