

May 6-10, 2007

San Jose Convention Center

San Jose, California, USA

Session: A13

## Back to Basics: A Day in the Life of a DB2 Page

IDUG® 2007

North America

Pete Sexton  
*J P Morgan Chase*

May 10, 2007 10:40 a.m. – 11:40 a.m.

Platform: DB2 for z/OS



2007 JPMorgan Chase & Co.  
All Rights Reserved.  
TO JPMORGAN CHASE & CO.

GoFurther



## Legal Disclaimers

**2007 JPMorgan Chase & Co.  
All Rights Reserved.  
TO JPMORGAN CHASE & CO.**

## Agenda

DB2 Page layout

DB2 subcomponents

A "Getpage"

Reading vs. Updating

The afterlife

Miscellaneous

JPMorganChase 

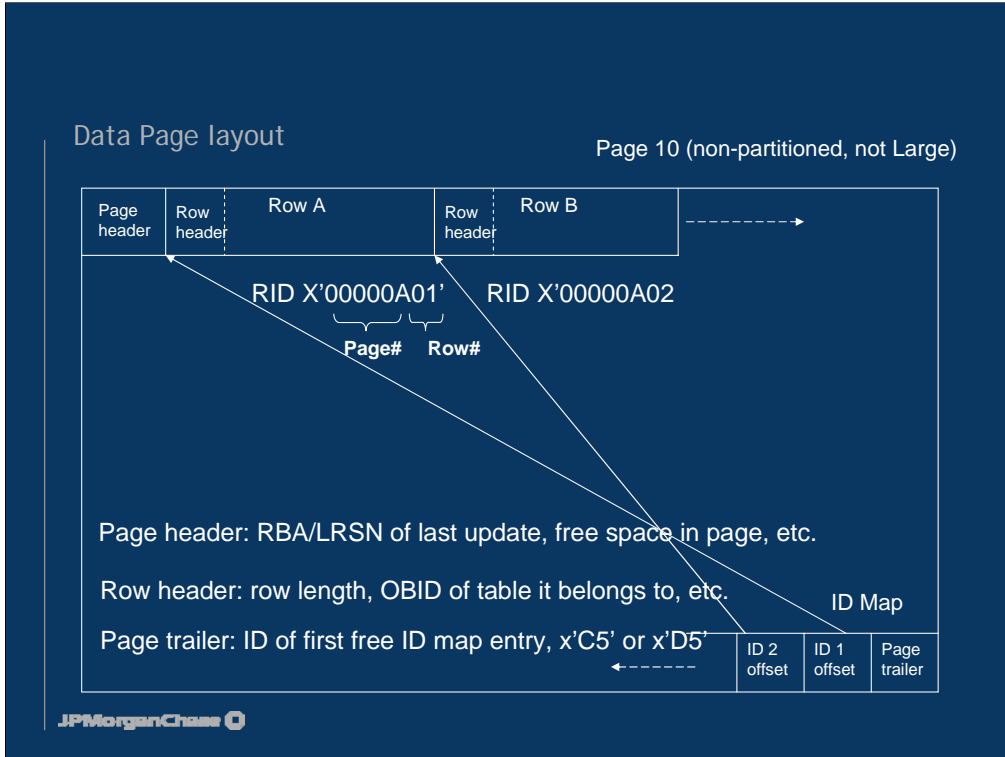
Geared toward new DB2 professionals

## Some terms

**Page Set** - a collection of pages comprised of different types.  
Can be Simple, Segmented, or Partitioned.

**Page type** - a physical format of a page. Examples are: header, space map, data, index, dictionary, etc.

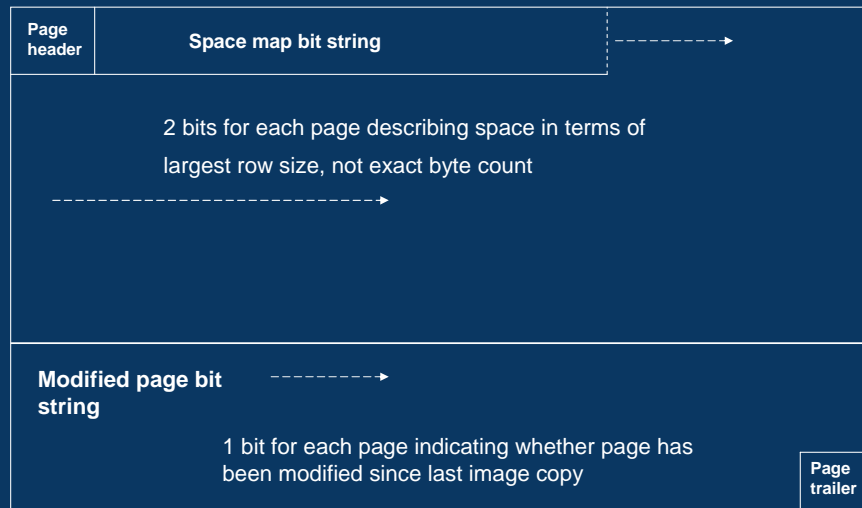
**RID** - Row identifier.



Broken pages

Latching


## Space Map Page layout



JPMorganChase

Segmented TS a little different format. Mostly just data to keep track by segment. Page Header/Trailer similar to data page – LRSN of last update, X'C5' or X'D5'

## Pageset/Partition



Header: DBID/OBID, RBA/LRSN of last page update,  
RBA/LRSN if RECOVER utility run, etc.

Space map: Free space information for pages that follow this SMP.

## DB2 subcomponents

### Buffer Manager (BM)

- moves pages between DASD and storage buffers
- writes updated pages
- opens/closes pagesets/partitions

### Data Manager (DM)

- finds and updates requested data
- requests locks and logging functions
- maintains pageset space and indexes

### Relational Data System (RDS)

- parses SQL statements
- checks authorizations
- selects access paths
- executes SQL statements in executable format

JPMorganChase 

DM – actual locking is done by the IRLM and the Recovery Log Manager does the log writes/reads (as stated later)



## DB2 subcomponents - continued

### Recovery Manager (RM)

- controls units of recovery
- coordinates across different recovery managers

### Recovery Log Manager (RLM)

- reads/writes the DB2 log records
- maintains Boot Strap Datasets

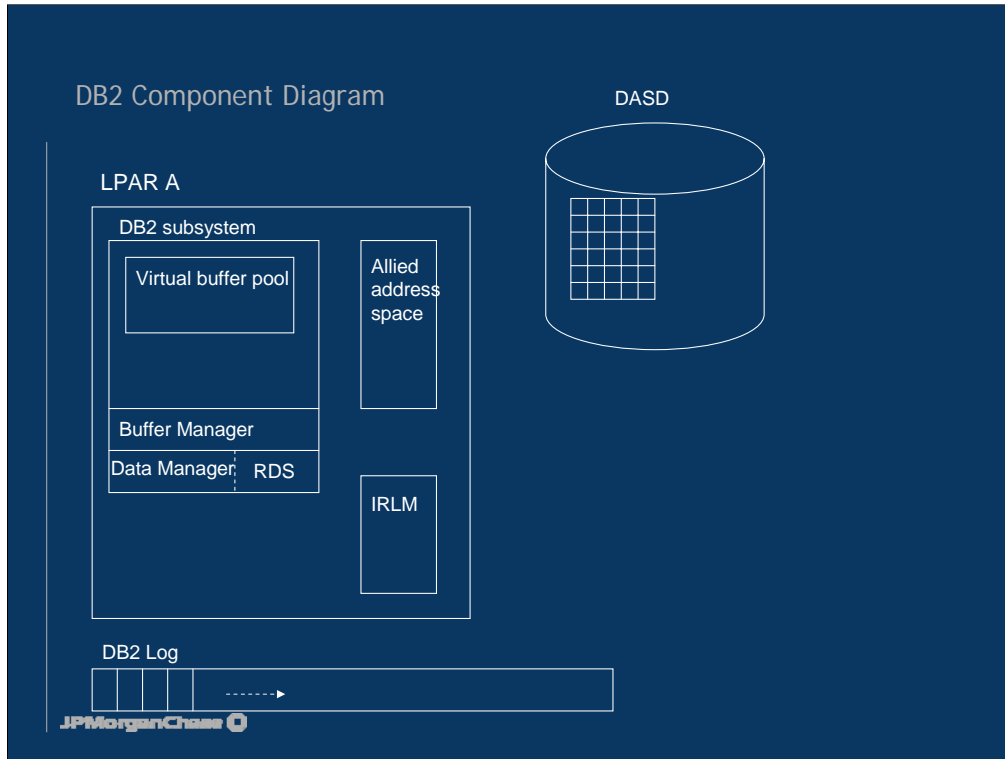
### Internal Resource Lock Manager (IRLM)

- grants lock requests from requestors (DB2 and IMS)
- performs Deadlock processing
- communicates across data sharing group via Notices

## DB2 subcomponents - continued


### Allied Address Space (User)

- a connector to DB2 – CICS, Batch, TSO, Distributed, etc.
- requests services from DB2
- CPU is charged here for synchronous reads/writes, locking requests, thread creation, etc.



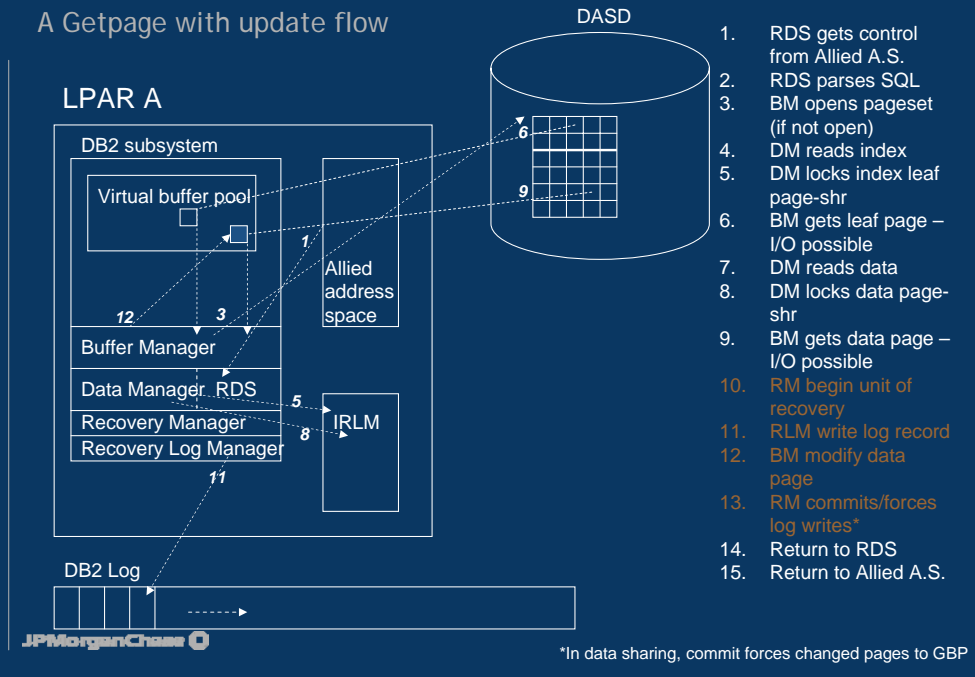
Shows where each component resides relative to the DB2 subsystem.

## A Getpage



1. RDS gets control from Allied A.S.
2. RDS parses SQL
3. BM opens pageset (if not open)
4. DM reads index
5. DM locks index leaf page-shr
6. BM gets leaf page – I/O possible
7. DM reads data
8. DM locks data page-shr
9. BM gets data page – I/O possible
10. Return to RDS
11. Return to Allied A.S.


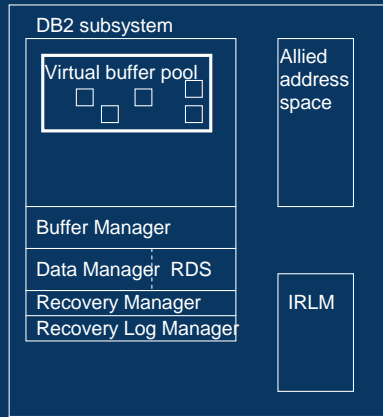
## A Getpage with update flow



Read only

DASD

LPAR A



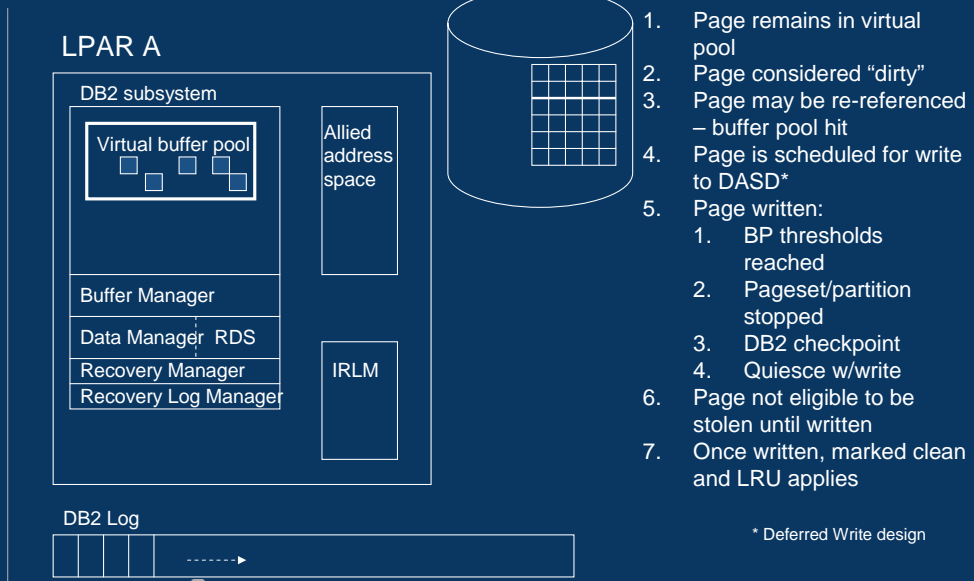
1. Page remains in virtual pool
2. Page considered "clean"
3. Page may be re-referenced – buffer pool hit
4. Least Recently Used algorithm
5. LRU for page steal
6. Data sharing implications – de-registration

DB2 Log



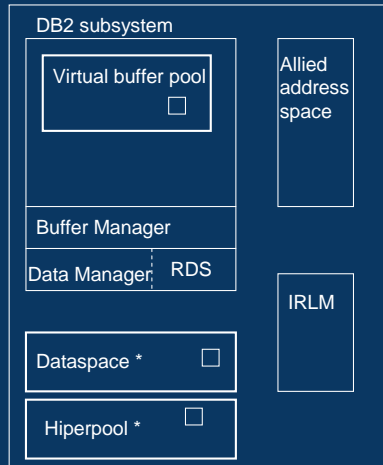
JPMorganChase

## Page update



# Where is the Page ?


## LPAR A



\* V7 only



## DASD



## Coupling Facility



1. Virtual pool
2. Hiperpool
3. Group Buffer Pool
4. DASD



## Miscellaneous

Force log buffers when page written to DASD before commit

Auxiliary Storage – Page Fault

JPMorganChase 

Can DASD copy be more current than log ?

MVS Paging – DB2PM reports reads/writes that cause MVS Page-In operation: On  
Statistics Report Long - Page-ins required for read/write

Diagnosis Guide and Reference

# Pete Sexton

JP Morgan Chase

Peter.C.Sexton@JPMCHASE.COM

