



# Breaking the DB2 Platform Barrier

Comparing the Architectural  
Differences Between DB2 UDB for z/OS  
And DB2 UDB LUW

Rich Wolfson  
Quest Software

# Agenda

- Basic Components
- Terminology Differences
- Storage Management
- Object Comparisons
- Administrative Differences

# Components of DB2

## **z/OS**

- Subsystem
- VCAT/Volume
- Stogroup
- Database
- Tablespace
- Creator (Owner/ Schema)
- Table
- Alias
- Synonym
- Index
- View
- Package
- Plans

## **LUW**

- Instance
- Container
- N/A
- Database
- Tablespace
- Schema
- Table
- Alias
- Index
- View
- Package

# Installation

## z/OS

- Only one version of DBMS

## LUW

- Enterprise Edition (EE)
- Enterprise Extended Edition (EEE)
- Enterprise Server Edition (ESE - V8)
  - Optional DB partitioning feature

# System Catalog

## **z/OS**

- **SYSIBM.xxxx**
  - Most optimizer related fields are updateable

## **LUW**

- **SYSIBM.xxxx**
- **SYSCAT**
  - Read-only views defined for catalog base tables
- **SYSSTAT**
  - Updateable set of views
  - Primarily used for access path manipulation


# Accessing DB2

## z/OS

- DB2I
  - DB2 tool set (3270 based)
    - SPUFI
    - DCLGEN
    - Bind/Rebind
    - Command Processor
    - Utilities
    - Defaults
- Control Center

## LUW

- Control Center
  - Basic administration
  - Command center
  - Command line processor
  - Command window
  - Script center
  - Visual Explain
- Health Center
  - Storage
  - Memory
  - Maintenance



# Common Terms, Different Meanings

## z/OS

- SMS
  - System Managed Storage
    - Software for managing all the DASD in an S/390 environment
- Extent
  - Physical extension of a dataset based on a secondary allocation.

## LUW

- SMS
  - System Managed Space
    - Type of space management for TS
- Extent
  - A block of pages within a tablespace
  - Similar to SEGSIZE in OS/390

# Different Terms, Similar Meanings

## z/OS

- **Subsystem** – Logical database environment
  - System databases
    - DSNDB06
    - DSNDB01
    - DSNDB04
    - DSNDB07
  - Memory Structures
  - Database Configuration
    - DSNZPARM
  - Many databases

## LUW

- **Instance** – Logical database server environment
  - Also referred to as a NODE
  - 1 to many databases
  - Database Manager Configuration File

# Bufferpools

## z/OS

- 80 bufferpools available
- Shared by all objects in subsystem

## LUW

- Defined within a database
- Can only be used by objects within database
  - Typically 1 per database
- Create, Delete, Alter without stopping DB (V8)



# Databases

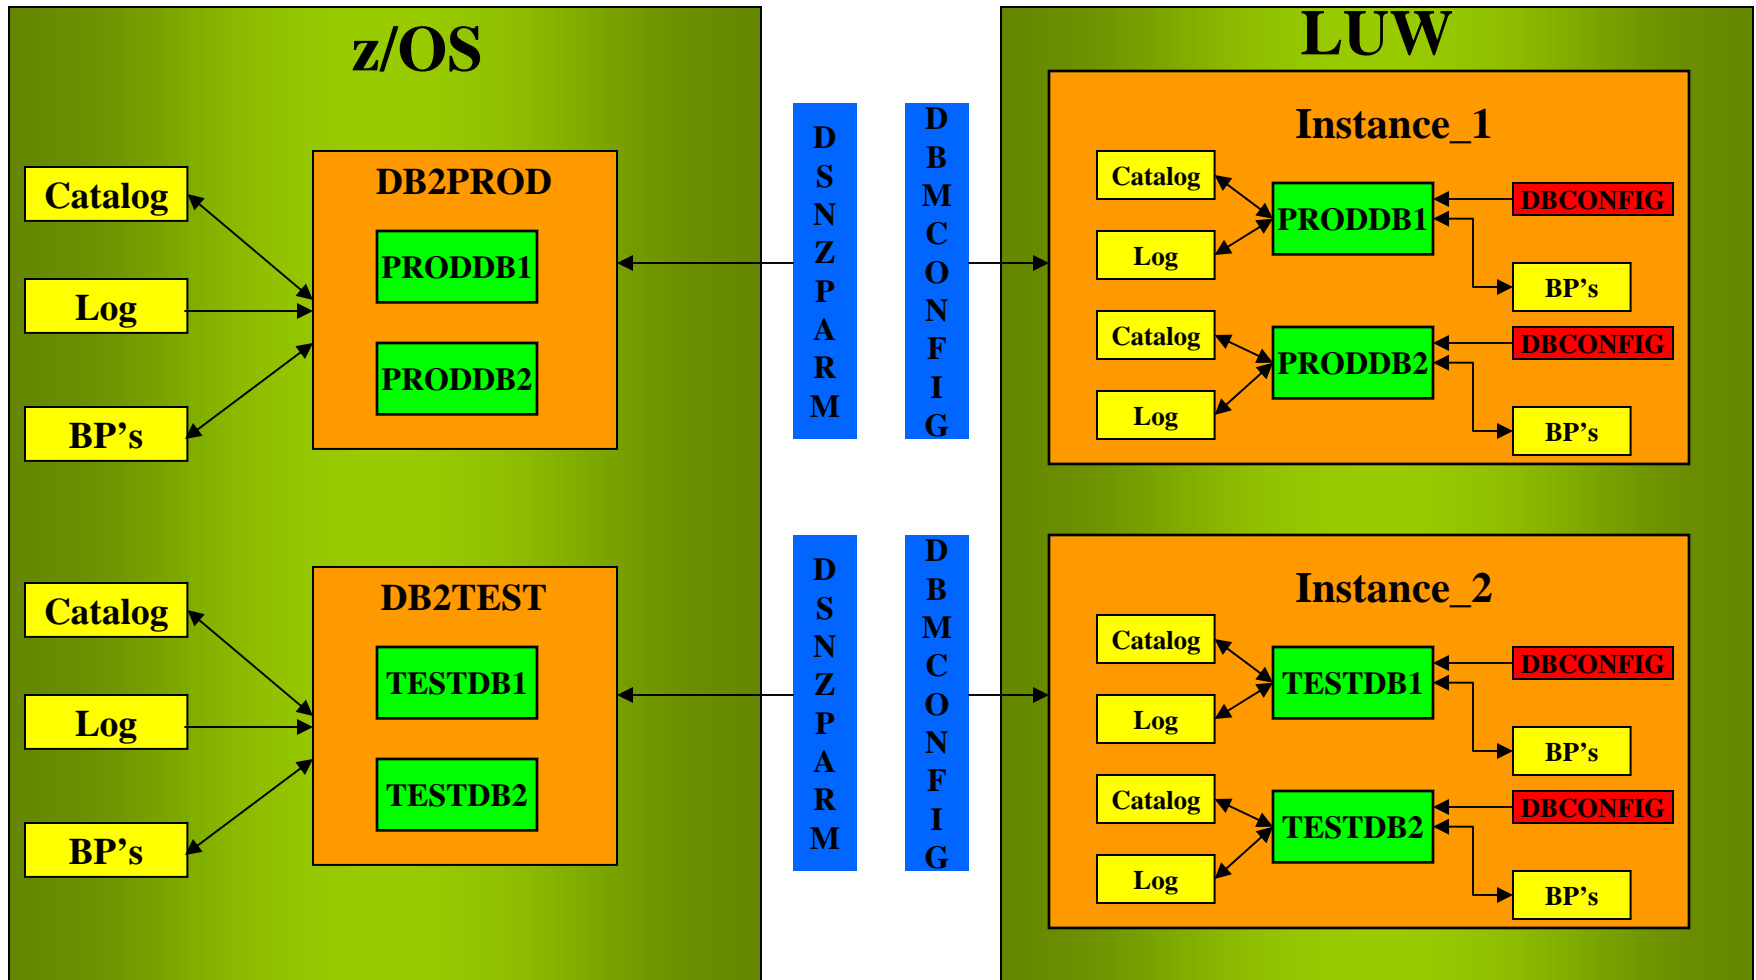
## z/OS

- Logical grouping of DB2 objects
  - Does not consume resources
- Many DBs in subsystem
- Metadata for all DBs stored in 1 system catalog

## LUW

- Logical grouping of DB2 objects
- Typically 1 database/instance
- More like an OS/390 Subsystem
  - Catalog for each database defined within database
    - SYSCATSPACE
    - TEMPSPACE
    - USERSPACE
  - Bufferpools defined in database
  - Database configuration file

# Subsystem vs. Instance



# Biggest Difference: Storage Management

## z/OS

- Volume
  - Device
- **Storage Group**
  - Group of volumes

## LUW

- Container
  - Device
  - Directory
  - File

# Volumes

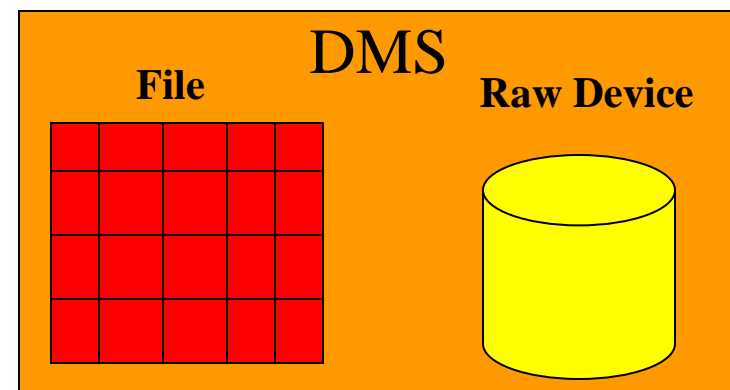
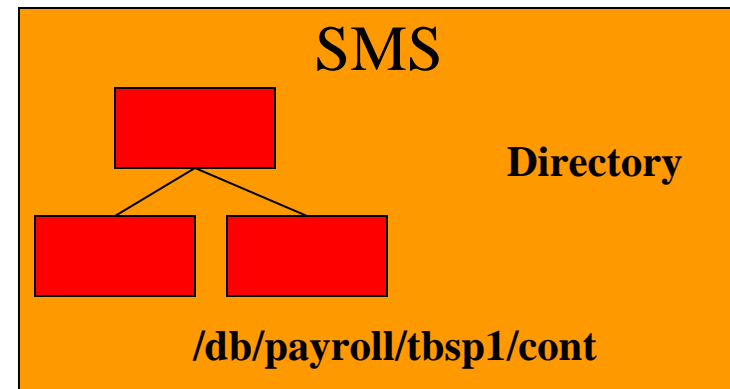
Physical storage device for DB2 z/OS. A volume can contain 1 or many tablespaces or indexspaces

## Terminology

- **DASD** – **D**irect **A**ccess **S**torage **D**evice
  - logical disk drives
- **VolSer** – Volume serial. This is a name identifying the disk pack i.e. DB2001.
- **Storage Group**
  - Defined object
    - A logical grouping of volumes
    - Can be used by more than one TS or IS
    - N/A on LUW.

# Containers

- Physical storage device for DB2 UDB LUW
- A container is assigned to an individual tablespace
  - **Directory name**
    - SMS Only
    - D:\MYTS
  - **Raw Device**
    - DMS Only
    - E:
  - **File name**
    - DMS Only
    - D:\SODADB\SODA.UTILITY.DMS



# Tablespaces: z/OS

4 types of tablespaces can be defined

- Simple
- Segmented
- Partitioned
- Large (DSSIZE)

2 types of allocation methods

- VCAT
- Stogroup

# Tablespaces: z/OS

## Simple

- 1 To many tables
- Table pages commingled
- Smallest unit of recovery is the tablespace

## Segmented

- Can contain multiple tables but rows are not commingled
  - Space is divided into groups of pages called Segments
    - Each segment contains rows for only one table
    - Each table can have different locking strategy
    - Reads only relevant pages during TS scan
  - Automatically reclaims space after drop table
  - Much more efficient for mass deletes

## Tablespaces: z/OS

When a tablespace is created, a VSAM file is defined with the following format:

```
VCAT.DSNDBC.DBNAME.TSNAME.I0001.A001  
VCAT.DSNDBD.DBNAME.TSNAME.I0001.A001
```

### Where:

VCAT – Typically the subsystem name

DBNAME – Database name

TSNAME – Tablespace name

A001 – Partition or dataset number (A001, A002, etc.)

# Tablespaces: DB2 LUW

- 1 type of Tablespace
  - 3 Categories
    - Regular
    - Temporary
    - Large
- Extents
  - A unit of grouped pages
    - 2 – 256 pages
    - Similar to SEGSIZE in OS/390
- 2 Allocation Methods
  - SMS – System Managed Space
    - Directory – SMS only
  - DMS – Database Managed Space
    - File
    - Device

# Tablespaces: DB2 LUW

- **System Managed**
  - No finite storage specified
    - Operating System's file manager allocates space as needed
    - Good for small tables
    - Cannot Add/delete containers after creation
- **Database Managed**
  - Space is pre-allocated
  - Better suited for large tables
  - LOBs must be DMS
  - ALTER to add containers

SMS    DMS

Add Containers to TS		X
Separate Indexes from Data		X
Space allocated as Needed	X	
High performance in heavy OLTP		X
High performance in decision support		X
Ease of administration for small tables	X	
Flexibility of Administration		X
File or Device containers		X

# Partitioning

## z/OS

- Partitioned TS Definition
  - Single table
  - 1-254 partitions
    - 4096 in V8
  - Partitioning key range controls the partition in which the data resides
  - Each partition can be on separate device

## LUW: EEE Only

- Hardware (not data) driven
  - Partitions usually on separate machines
- Controlled via Database Partition Groups
  - Formerly NODEGROUPS
  - DB Part = Node
    - Data
    - Indexes
    - Config files
    - Logs
- Specify key but not data values

## Database Partition Group

A set of one or more database partitions-

Tablespace(s) exists within a Nodegroup

- More than 1 table can be in a Nodegroup
- Rows are distributed across partitions of Nodegroup
  - Partitioning Map controls data placement
    - Hash function places rows on a given partition
  - Data will be evenly distributed across Nodes in Nodegroup

# Tables

## z/OS

- 1 to many tables defined in simple or segmented tablespaces
  - Tables and Indexes are independent of each other

## LUW

- 1 to many tables can be defined within a tablespace
  - Indexspace directly tied to table definition and can exist in same tablespace

# Indexes

## z/OS

- Unique
- Non-unique
- Clustering\*
- Partitioning

## LUW

- Unique
- Non-unique
- Clustering
- Multi-Dimensional Clustering (V8)

\*- Non partitioned TS only



# Indexes: z/OS

## Indexspace created when **CREATE INDEX** executed.

- No CREATE INDEXSPACE DDL like tablespaces
- Only 1 index per indexspace
- VSAM dataset name can be a little cryptic for indexes

VCAT.DSNDBC.DBNAME.IXNAME.I0001.A001

VCAT.DSNDBD.DBNAME.IXNAME.I0001.A001

Where:

VCAT – Typically the subsystem name

DBNAME – Database name

IXNAME – 8 character representation of IX name

A001 – Dataset number (A001, A002, etc.)

## 2 types of allocation methods

- VCAT
- Stogroup

## Indexes: DB2 LUW

Indexes are dependent on tables. The indexspace must be specified when the table is created.

- All indexes for table use 1 tablespace
- Indexspace is predefined before IXs are created
- Indexes can be defined in same tablespace as table
  - Required for SMS



# Administration


# Optimizer

## z/OS

- Fixed optimization
- HINTS in V6 allow for some flexibility
  - Mainly used to maintain old access path
  - Must be turned on at install time
  - Need to modify PLAN\_TABLE
  - Manual Process

## LUW

- Much more flexible than z/OS
  - 7 levels of optimization
  - Adjusted based on query complexity




# Optimization Class: Guidelines

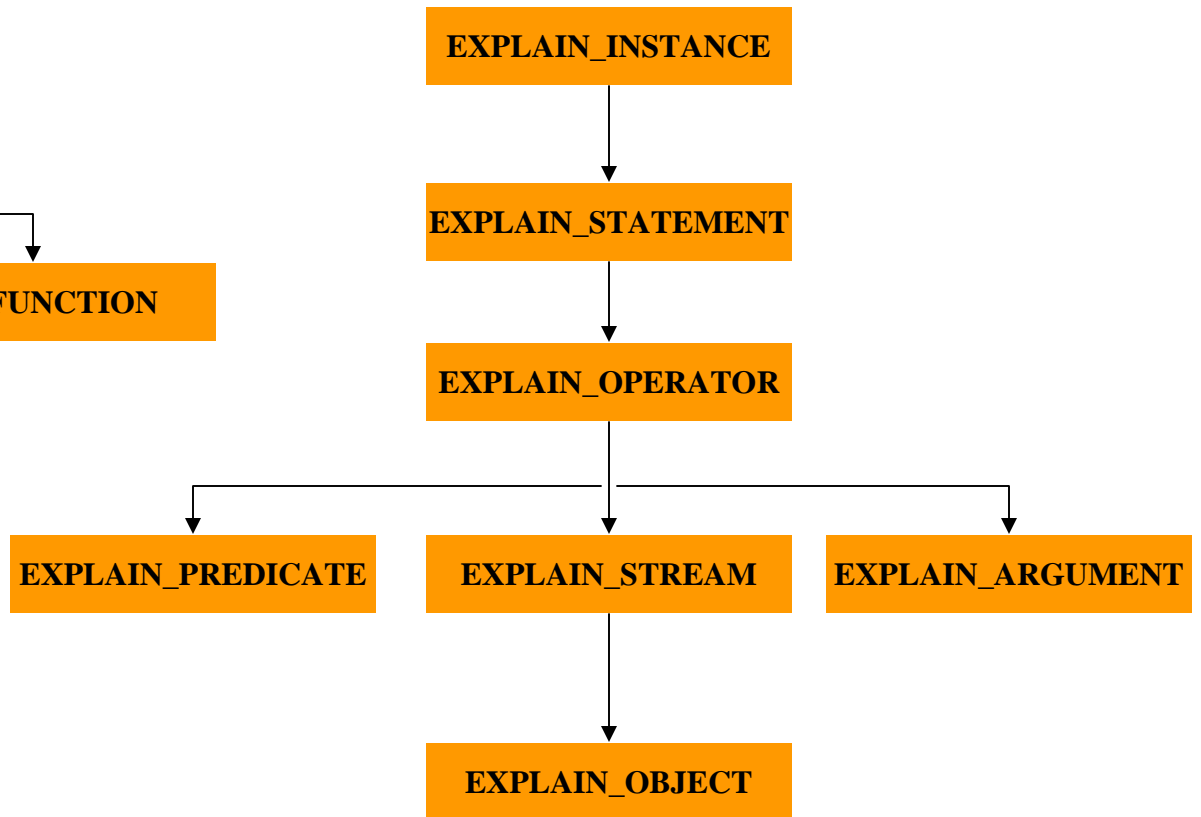
Level	Recommendation
0	Minimal amount of optimization. Only recommended for very simple SQL accessing well indexed tables. Only nested loop joins and IX scans enabled.
1	Similar to 0 except Merge Scan and TS scan enabled.
2	Recommended for very complex queries which are infrequently executed in a decision support or OLAP environment.
3	Closest to OS/390 optimizer. Recommended for queries with 4 or more joins.
5	DEFAULT – Most cost effective method for mix of simple and complex queries. Optimization will be automatically reduced for complex dynamic SQL if optimizer determines that the resources are not necessary.
7	Same as 5 except optimization not reduced for complex dynamic SQL
9	Used to determine whether more comprehensive optimization can generate better access plan for very complex long running queries using large tables

# Explain

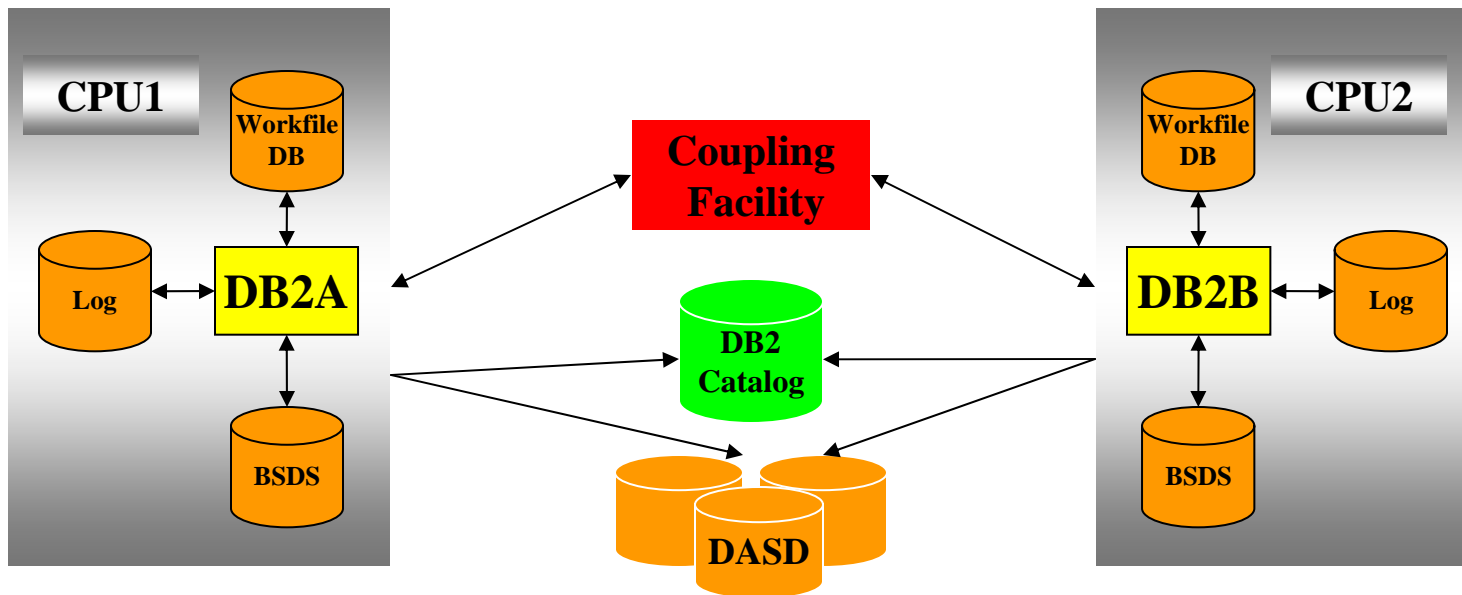
## z/OS



## LUW



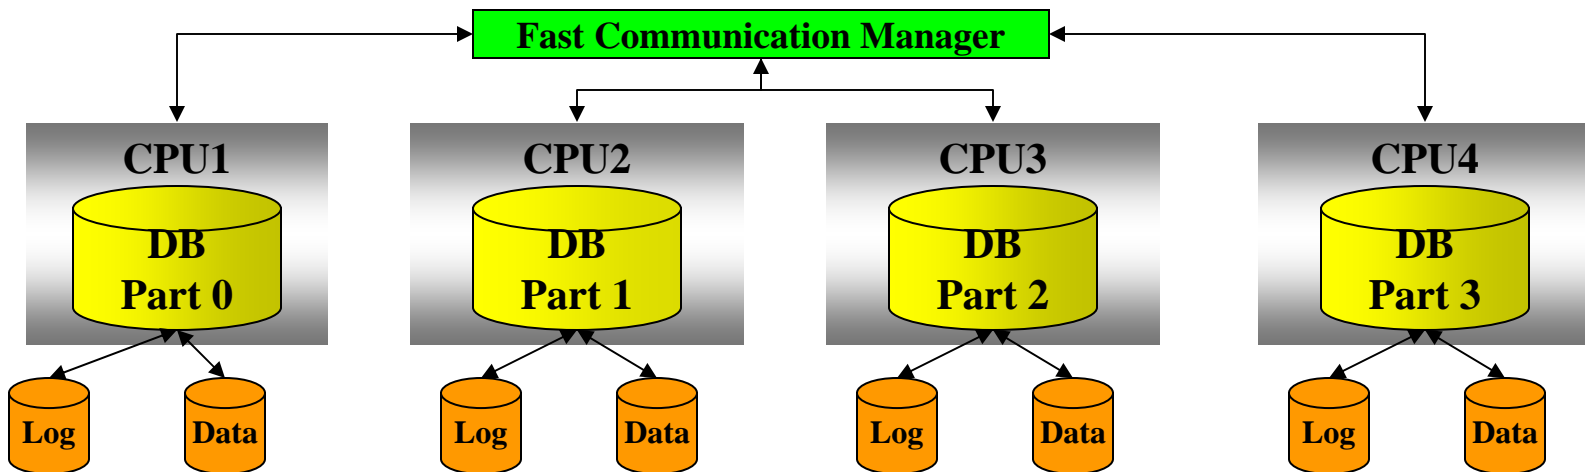
# Parallelism: z/OS



# Parallelism: DB2 LUW

Enterprise Extended Edition

MPP




# Types of Parallelism

## z/OS

- I/O
  - Partitioned TS
- CPU
  - Query processed in multiple tasks in parallel
- SYSPLEX

## LUW

- I/O
  - Multi-container TS
- Query
  - Intra-partition (SMP)
    - Parallelism within single partition
  - Inter-partition (EEE/MPP)
    - Parallelism across multiple partitions




# Monitoring

## z/OS

- **Instrumentation Facility Component (IFC)**
  - **Statistics**
    - Global statistical data
  - **Accounting**
    - Detail info for specific application
  - **Audit**
    - Table access audits
    - Requires AUDIT keyword on table definition
  - **Performance**
    - Most detailed \$\$\$
    - Only use for short periods
  - **Monitor**
    - Makes trace data available for monitoring applications

## LUW

- **Snapshot Monitor**
  - Show status of database for an instant in time
- **Event Monitor**
  - Historical status over time
    - Databases
    - Tablespaces
    - Connections
    - Tables
    - Statements
    - Transactions
    - Deadlocks



# Utilities


## z/OS

- COPY
- DSNTIAUL/Fast Unload
- LOAD
- RECOVER
- REORG (TS,IX)
  
- RUNSTATS
  - “Real-Time Statistics” (V7)
- QUIESCE
- MERGECOPY
- CHECK DATA

## LUW

- BACKUP
- EXPORT
- LOAD/IMPORT
- RESTORE
- REORG (Table)
  - REORGCHK
- RUNSTATS
- QUIESCE
- Set Integrity

# Backup and Recovery



# Backups

## z/OS

- Tablespace
- Index
- Components
  - Full copy
  - Incremental Copy
  - Copy to Copy (V7)
  - Active/Archive Logs
  - BSDS
  - SYSLGRNX

## LUW

- Database
- Tablespace
- Components
  - Backup Image
  - Incremental Copy
  - Backup History File
  - Active Logs
  - Archive Logs

# Recovery Information

## z/OS

### SYSCOPY

- **Updated**
  - Full IC
  - Incremental IC
  - Quiesce
  - LOAD
  - REBUILD IX
  - RECOER TOCOPY
  - RECOVER TOLOGPOINT
  - Reorg
- **Contains**
  - Full/Incr copy
  - Log RBA
  - Copy to Copy

## LUW

### Recovery History File

- **Updated**
  - Backup of full DB or TS
  - Restore of full DB or TS
  - Load of a table
  - Quiesce TS
- **Contains**
  - Part of DB which was copied
  - When DB was copied
  - Location of the copy
  - Time of last restore

# Logging

## z/OS

- Logs apply to entire subsystem
  - Active
  - Archive
- Active logs are automatically archived when full
- Dual Logging




## LUW

- Defined at database
  - **Circular**
    - No roll-forward recovery
  - **Archival**
    - Fully recoverable
    - Similar to z/OS
    - 3 log files
      - Active
      - Online Archived
      - Offline Archived
- On Demand Archiving
  - Close and archive an active log at any time
- Dual Logging

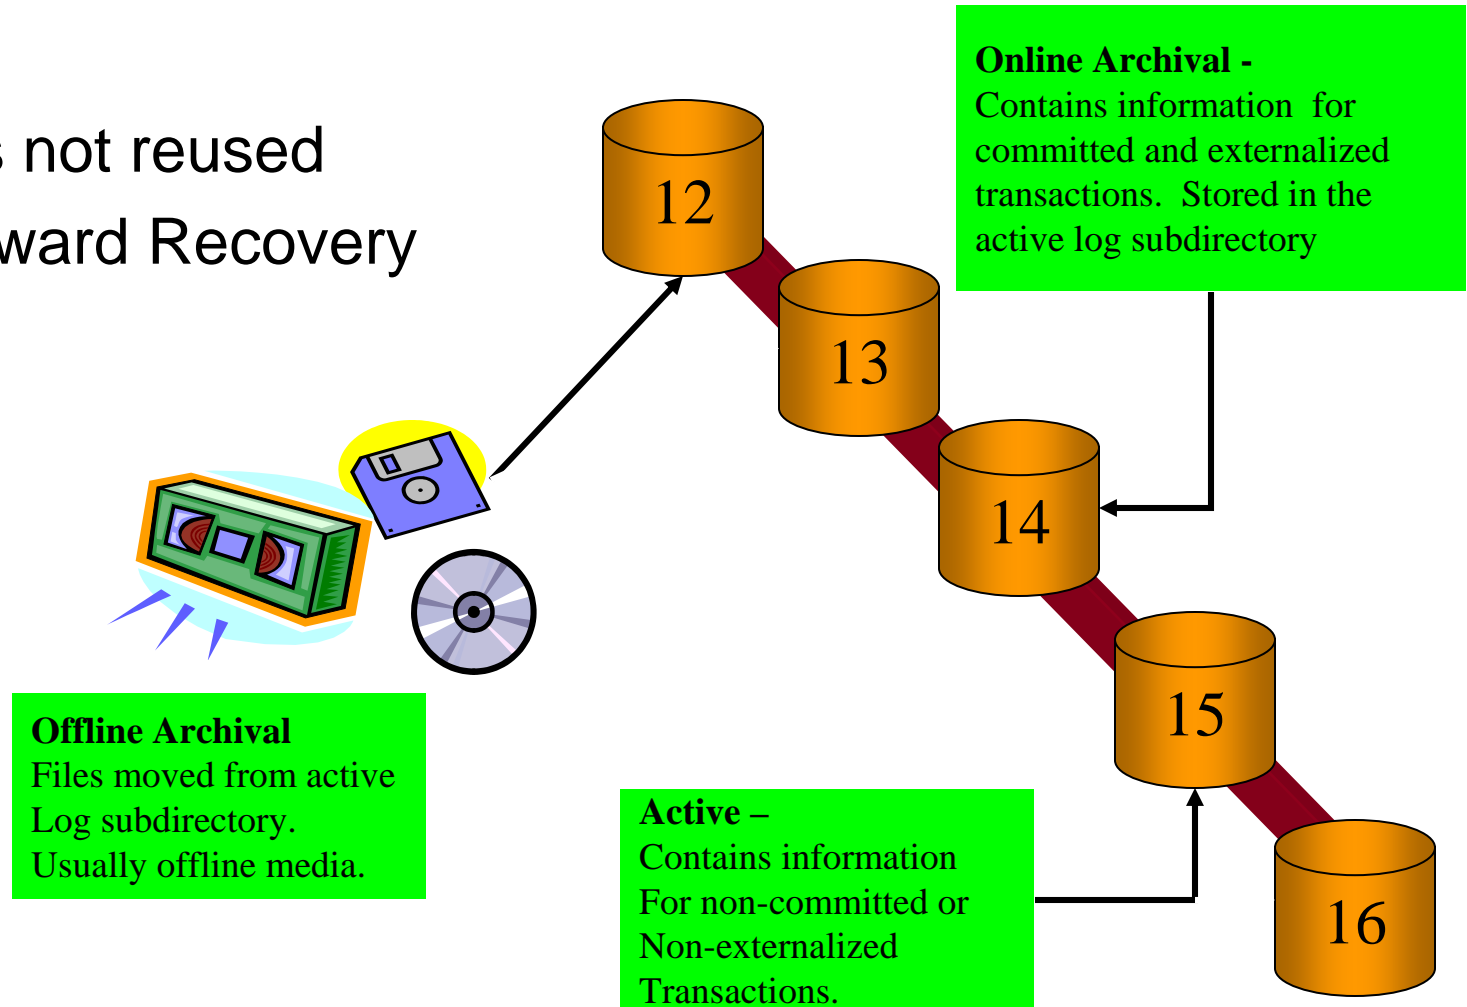
# Circular Logging: DB2 LUW

- **Non-recoverable databases**
- **Log files are reused**
- **Uses active logs only**
  - Secondary used for overflow
- **Roll-forward recovery not possible**
- **Default method for new DBs**



# Archival Logging

- Log files not reused
- Roll Forward Recovery



# Types of Recovery

## z/OS

- Crash
  - DB2 restart
- Roll-Forward
  - IC plus log apply
  - LOGONLY
- Point in Time
  - IC only (TOCOPY)
  - To RBA

## LUW

- Crash
  - Uses logs to recover from power interrupts or application ABENDS
- Roll-Forward
  - Image copy plus log apply
  - LOCAL TIME
    - More flexible than TORBA
- Version
  - Image copy (TOCOPY)

# Reorganizing Data

## z/OS

- Tablespace
  - Log Yes/No
  - Unload Pause
  - Shrlevel
- Index
- Online
  - SHRLEVEL CHANGE

## LUW

- Table
- Index
- REORGCHK
  - Determines when Reorg is required
- Online
  - “In Place”
  - Can require large log space

# Unloading Data

## *z/OS*

- DSNTIAUL
  - IBM sample program
- REORG UNLOAD  
PAUSE
- UNLOAD Utility (V7)
  - Table
  - Image Copy

## *LUW*

- EXPORT
  - Accessed via Control Center or CLP
  - Rename columns
  - Multiple output formats

# Loading Data

## z/OS

- **Load Utility**
  - Resume/Replace
  - Log YES/NO
  - Runstats/Copy
  - Sophisticated SQL processing
- **ONLINE**
  - SHRLEVEL CHANGE (V7)

## LUW

- **Load**
  - Locks single table in TS (V8)
  - Insert/Replace
  - RUNSTATS
  - Good for large amounts of data
  - READ ACCESS for rows not being loaded (v8)
  - Load directly from a SQL query (V8)
- **ONLINE (V8)**
- **Import**
  - Can dynamically create table
  - Insert process
    - Update
    - Replace
  - Good for small amounts of data

# Summary

## **To be a successful cross-platform DBA:**

- Have a sound foundation of relational principles
- Understand the nuances of the individual platforms
- Understand that the principles of database administration are similar regardless of the RDMS



**Thank you**