


Breaking the DB2 Platform Barrier



**Comparing the Architectural
Differences Between
DB2 for z/OS
vs.
DB2 for Linux, Unix, and Windows**

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

- **DB2 for z/OS**
 - Data Server
 - Value Unit Edition

LUW

- **DB2 Express**
 - Low Cost
 - Not scalable
 - 2 CPU/ 4GB RAM max
- **DB2 Workgroup Edition**
 - 4 CPUs/ 16 GB RAM
- **DB2 Enterprise Edition**
 - Unlimited scalability
 - Data partitioning (DPF)
- **DB2 Data Warehouse Edition**
 - Additional B.I. capabilities

System Catalog

z/OS

- **SYSIBM.xxxx**
 - 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**
- **Visual Explain**
- **Data Studio**



LUW

- **Control Center**
 - Basic administration
 - Command center
 - Command line processor
 - Command window
 - Script center
 - Visual Explain
- **Health Center**
 - Storage
 - Memory
 - Maintenance
- **Data Studio**
 - Stored procedure management
 - UDF manager
 - SQL Builder

Common Terms, Different Meanings

z/OS

- **SMS**
 - System Managed Storage
 - Software for managing disk allocation on System z
- **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 z/OS

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
 - **50 – 4K**
 - **10- 8, 16, & 32K respectively**
 - Share across subsystem
- Global Bufferpools
 - Shared across data Sharing group

LUW

- **IBMDEFAULTBP**
automatically created with database
 - Additional pools created with DDL
- Hidden Bufferpools
 - 4k,8k,16k, & 32K
- Share only within individual databases



Databases

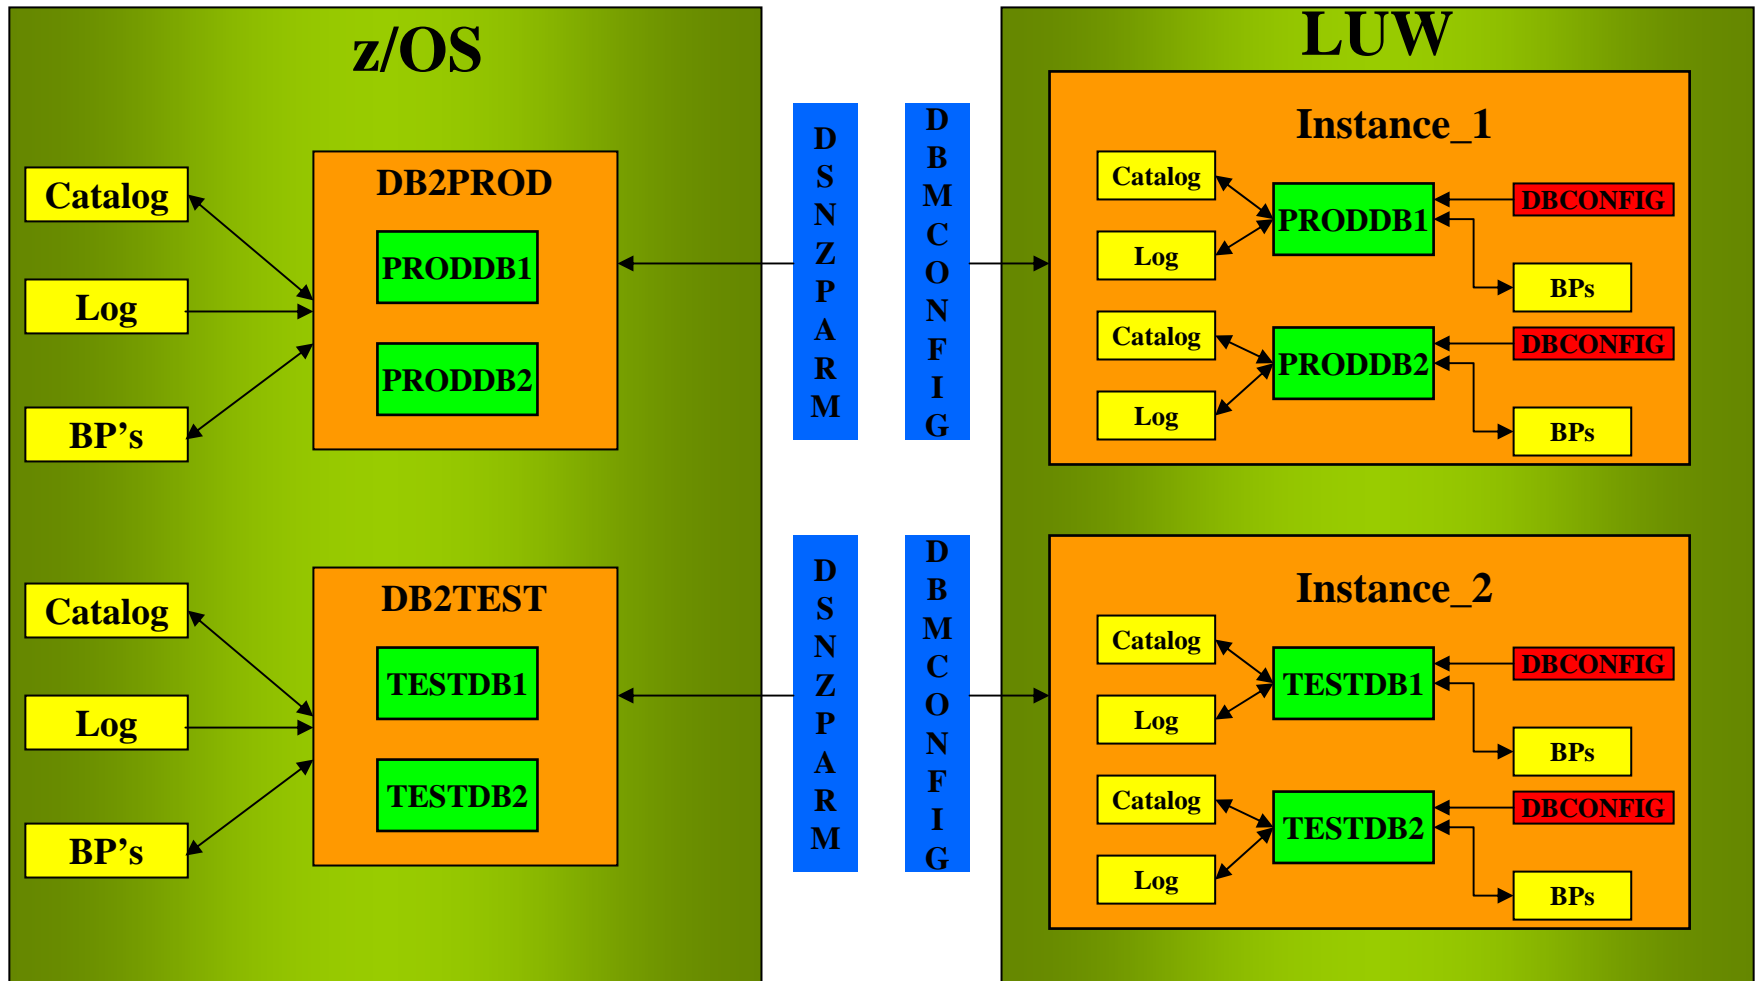
z/OS

- Logical grouping of DB2 objects
 - Does not consume resources
- Many DBs in subsystem
 - 60K
- Metadata for all DBs stored in 1 system catalog

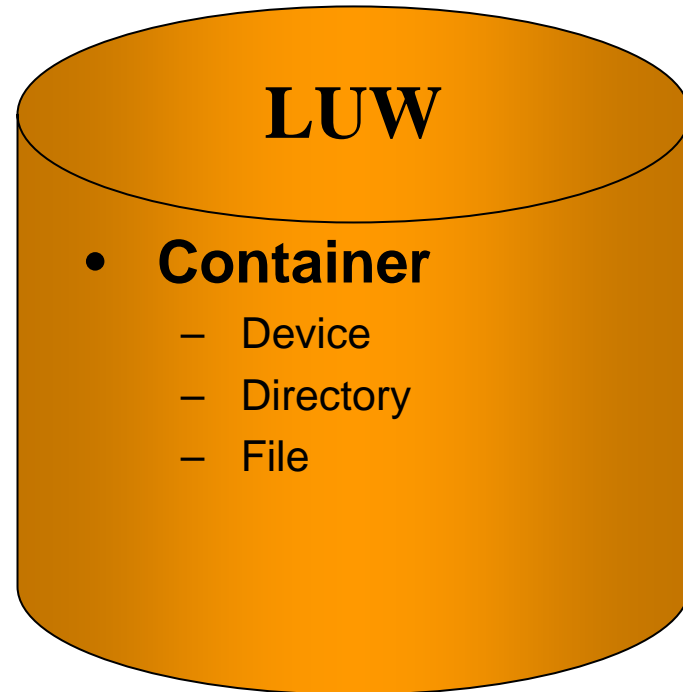
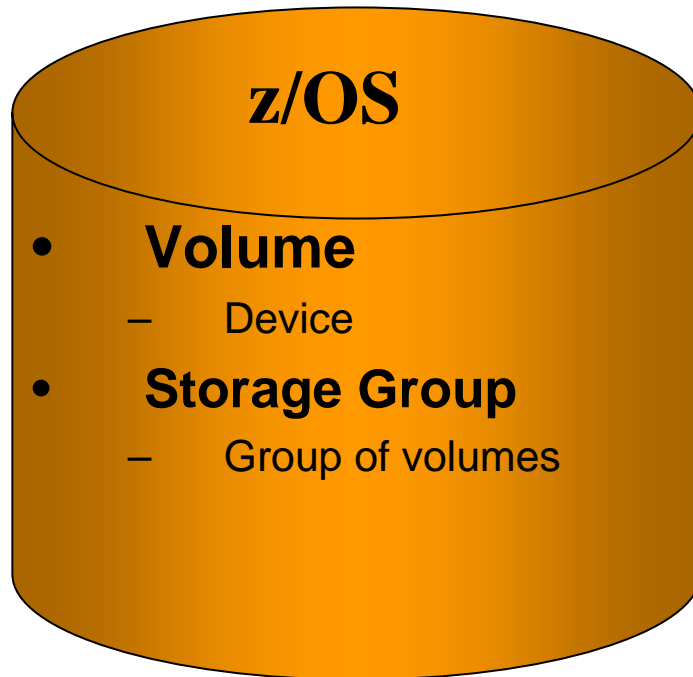
LUW

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

Subsystem vs. Instance



Storage Management



Volumes

Physical storage device for DB2 z/OS. A volume can contain 1 or many table spaces or index spaces

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 or SMS “*”
 - Can be used by more than one TS or IS
 - N/A on LUW.

Containers

- Physical storage device for DB2 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

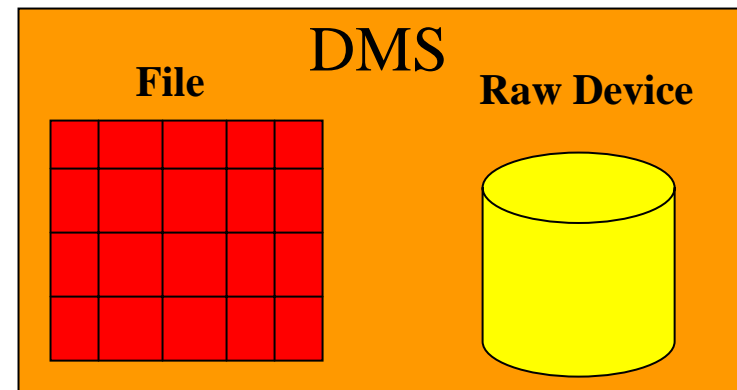
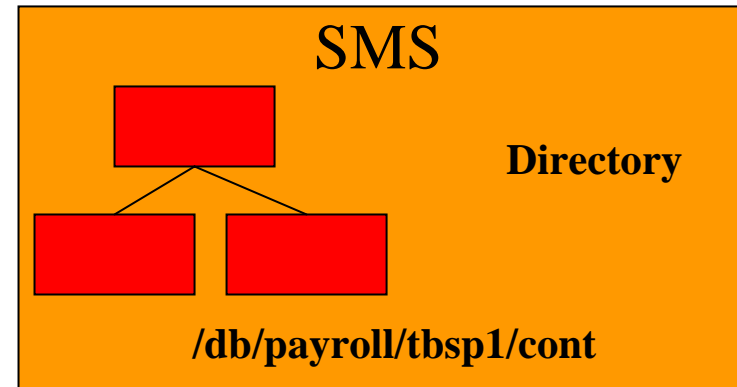


Table spaces: z/OS

5 types of table spaces can be defined

- Simple
 - N/A as of V9
- Segmented
- Partitioned
- DSSIZE (large)
- Universal (UTS)
 - V9

2 types of allocation methods

- VCAT
- Stogroup

Table spaces: 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

I0001 – Instance number changes with cloned tables

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 z/OS
- 2 Allocation Methods
 - **SMS** – System Managed Space
 - Directory – SMS only
 - **DMS** – Database Managed Space
 - File
 - Device
 - **AUTOMATIC**
 - No need to specify container info.

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

	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 Table Space**
 - Single table
 - 1-4096 partitions
 - Partitioning Index key controls the partition in which the data resides
 - Each partition can be on separate device
 - Partitions can be administered independently
 - Universal table space provides better space management
- **Universal Table Spaces**
 - Partition by Growth (PBG)
 - Segmented TS
 - Increased space limits
 - Utility parallelism
 - MAXPARTITIONS
 - Range Partitioned (PBR)
 - Combines benefits of segmented and partitioned tablespaces

LUW

- **Table partitioning**
 - Partitioning part of table DDL
 - Partitions can be in one or multiple tablespaces
 - Partitions can be rolled in/out
 - Partitions can be administered independently
- **Database Partitioning**
 - Database Partitioning Feature (DPF)
 - ESE Edition only
 - Database is partitioned across multiple servers
 - Multiple tables
 - Database Partition Group controls data location
 - Database Partition Group controls data distribution

Tables

z/OS

- 1 to many tables defined in non-partitioned table spaces
 - Typically 1 to 1 however for performance, easier management
 - Tables and Indexes are independent of each other
 - Except for PK which requires unique IX definition

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



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 allow for some flexibility
 - Mainly used to maintain old access path
 - Must be turned on at install time
 - Need to modify PLAN_TABLE
- Visual Plan Hints

LUW

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



Optimizer Class

- **DB2 Optimizer Class**
 - Values are between 0 and 9, default is 5
 - Determines the intensity used by the DB2 SQL Compiler when rewriting SQL
 - Dynamic SQL can't spend time optimizing, use lower class
 - Static SQL optimizes once, use a higher class
 - "dft_queryopt" database setting
 - SET CURRENT QUERY OPTIMIZATION n

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 z/OS 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 Tables

- Native Tables

- PLAN_TABLE
- DSN_FUNCTION
- DSN_STATEMENT

- Tables used by IBM Tools

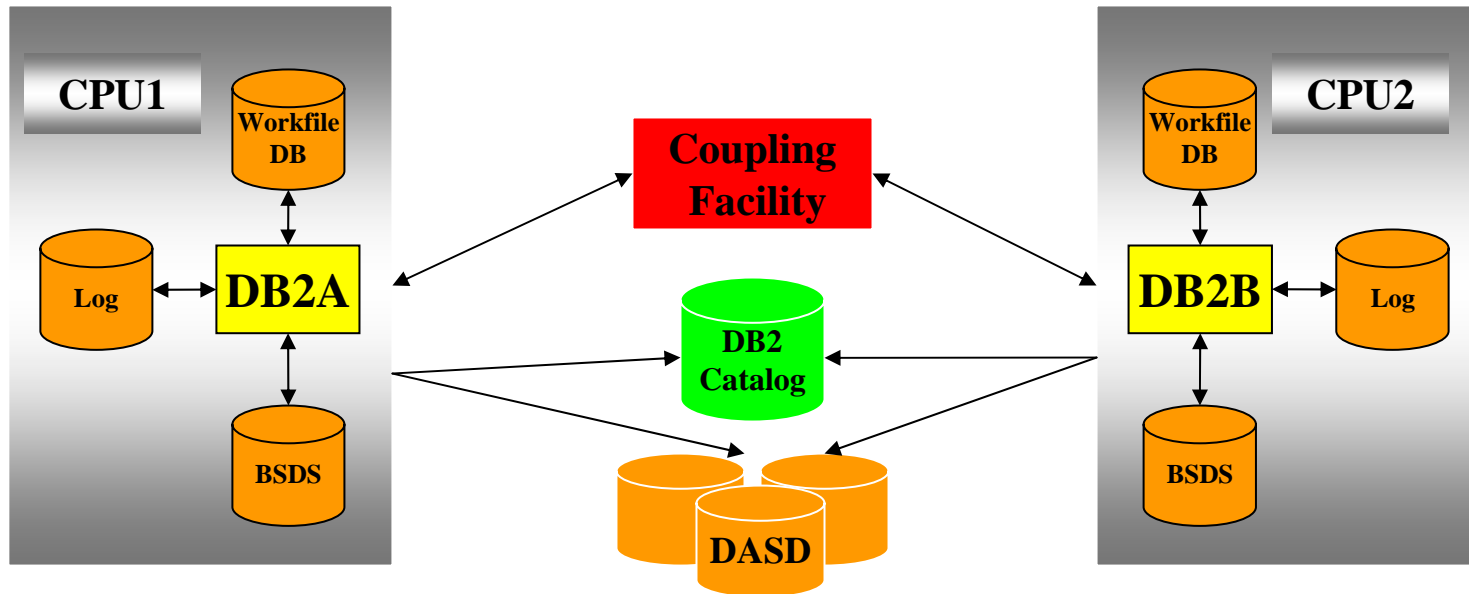
- DSN_PREDICAT_TABLE
- DSN_STRUCT_TABLE
- DSN_PGROUP_TABLE
- DSN_PTASK_TABLE
- DSN_FILTER_TABLE
- DSN_DETCOST_TABLE
- DSN_SORT_TABLE
- DSN_SORTKEY_TABLE
- DSN_PGRANGE_TABLE
- DSN_VIEWREF_TABLE
- DSN_QUERY_TABLE

- Native Tables

- EXPLAIN_INSTANCE
- EXPLAIN_STATEMENT
- EXPLAIN_OPERATOR
- EXPLAIN_PREDICATE
- EXPLAIN_OBJECT
- EXPLAIN_ARGUMENT
- EXPLAIN_STREAM

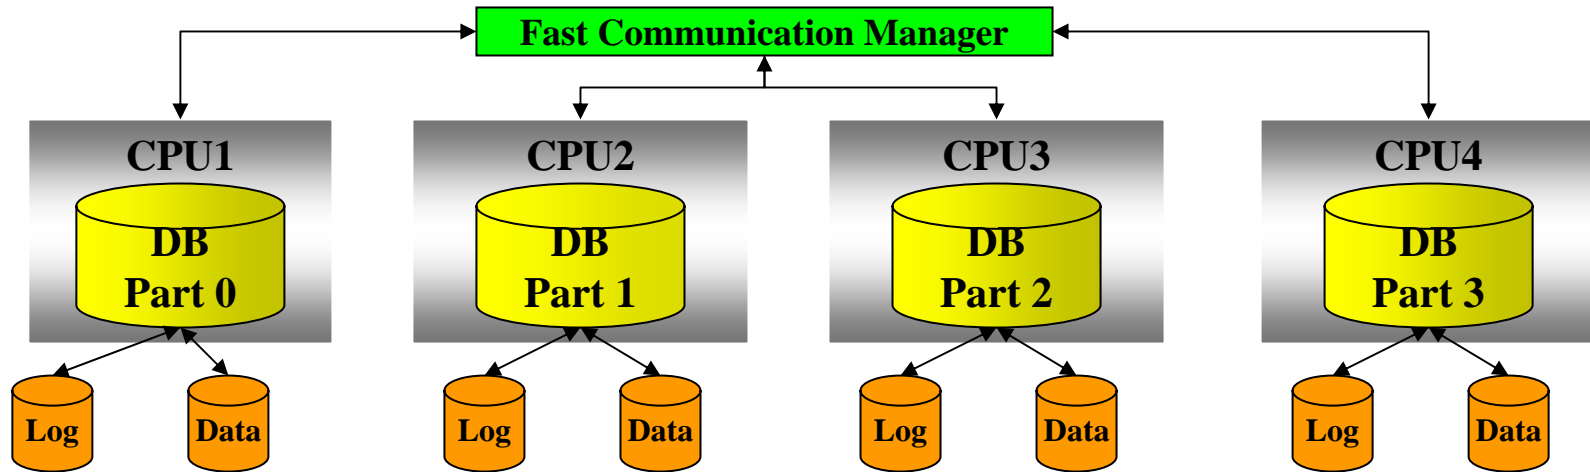
Parallelism: z/OS

Data Sharing SYSPLEX



Parallelism: DB2 LUW

Enterprise Extended Edition MPP




Types of Parallelism

z/OS

- I/O
 - Partitioned TS
- CPU
 - Most common
 - Query processed in multiple tasks in parallel across machines or LPARs
- SYSPLEX

LUW

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




Performance 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
 - 3 types
 - Activity
 - Threshold Violations
 - Statistics



Automated Monitoring - LUW

- Self Tuning Memory (V9)
 - Bufferpools
 - Locklist
 - Package Cache
 - Sort Heap
 - Database Shared Memory

Backup and Recovery



Backups

z/OS

- BACKUP SYSTEM
- Tablespace
- Index
- Components
 - Full Image copy
 - Incremental Image Copy
 - Copy to Copy
 - 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
 - RECOVER 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

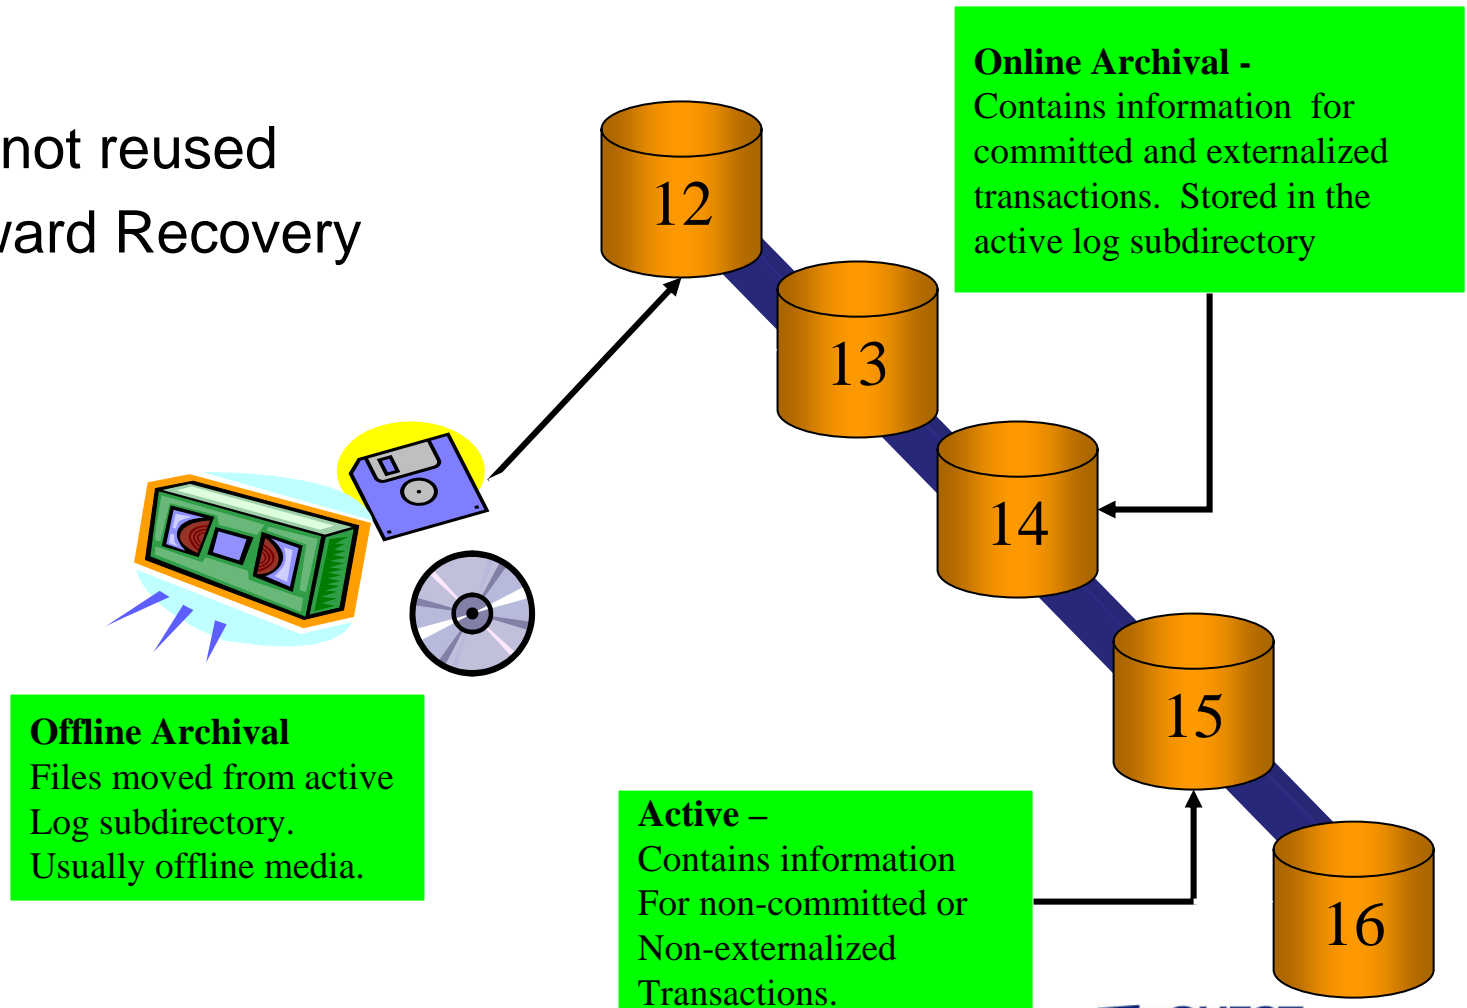
Circular Logging: DB2 LUW

- **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

- **RESTORE SYSTEM**
 - When BACKUP SYSTEM used
- **Crash**
 - DB2 restart
- **Roll-Forward**
 - IC plus log apply
 - LOGONLY
- **Point in Time**
 - IC only (TOCOPY)
 - TOLASTCOPY
 - TOLASTFULLCOPY
 - TOLOGPOINT
 - TORBA

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)

Common Utilities

z/OS

- DSNTIAUL/Fast Unload
- LOAD
- REORG
- RUNSTATS
- QUIESCE
- CHECK DATA

LUW

- EXPORT
- LOAD/IMPORT
- REORG
- RUNSTATS
- QUIESCE
- Set Integrity

Unloading Data

z/OS

- **DSNTIAUL**
 - IBM sample program
- **REORG UNLOAD PAUSE**
- **UNLOAD Utility**
 - 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
 - Enforce Constraints
- **ONLINE**
 - SHRLEVEL CHANGE

LUW

- **Load**
 - Locks single table in TS
 - Insert/Replace
 - RUNSTATS
 - Good for large amounts of data
 - READ ACCESS for rows not being loaded Load directly from a SQL query
- **ONLINE**
 - **Allow Read Access**
- **Import**
 - Insert process
 - Update
 - Replace
 - Good for small amounts of data

Reorganizing Data

z/OS

- **Tablespace**
 - Log Yes/No
 - Unload Pause
 - Shrlevel
- **Index**

LUW

- **Table**
- **Index**
- **REORGCHK**
 - Determines when Reorg is required
- **Automated Reorganization**
 - Automated REORGCHK

Statistics

z/OS

- RUNSTATS
 - Tablespace
 - Index
- Real Time Statistics
 - Stored Procedure
 - DNACCOX
 - V9
 - DSNACCOR
 - V8
 - Monitor for
 - REORG
 - RUNSTATS
 - Image Copy

LUW

- RUNSTATS
 - Table
- Statistics Profile
 - Customized parameters for individual tables
- Automated statistics collection
 - DB2 schedules RUNSTATS via maintenance policy

Quiesce

z/OS

- Establishes point of consistency in log
 - Flushes all changed pages to disk
 - Used to avoid RI violations

LUW

- Instance or Database level
 - Used for administering maintenance
 - All objects in “Quiesced Mode”
 - Only authorized users can access objects
 - Similar to ACCESS(MAINT) in z/OS
- UNQUIESCE
 - Removes Quiesce Mode

Checking Referential Integrity

z/OS

- **Check Data**
 - Removes “Check Pending” flag
- **Exception Tables**
 - Contain RI violations
- **Invocation**
 - Standalone
 - Load Utility “Enforce Constraints”

LUW

- **Set Integrity**
 - Removes “Set Integrity Pending” state
- **Exception Tables**
 - Contain RI Violations
- **Invocation**
 - Embedded in Application
 - Dynamic SQL

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

