

Abstract

- This talk will focus on recent and future improvements in DB2 for z/OS. The key focus points will be on continuous availability. This session will highlight some of the key changes.
- We expect DB2 UDB for z/OS Version 8 to make a fundamental change in many areas, reengineering much of DB2 and helping deliver more growth. There are key improvements that help improve scalability, improved ability to port applications, much better ability to deliver continuous availability. The ability to handle very large databases is changing in fundamental ways. Support for key vendor applications is compelling. The enhancements for 64 bit virtual storage will make management simpler, improve scalability and availability.



Agenda


- Availability Issues
 - ▶ Reduce planned & unplanned downtime
 - ▶ Manage to Service Level Agreement
 - ▶ Service Strategy
 - ▶ Checkpointing and long running Units of Recovery (URs)
 - ▶ CF Duplexing
 - ▶ Virtual Storage Management
 - ▶ Data Recovery
 - ▶ GDPS PPRC HyperSwap

- Achieving DB2 High Availability by DB2 release
 - ▶ DB2 Version 6
 - ▶ DB2 Version 7
 - ▶ DB2 Version 8
 - ▶ Futures



What can go wrong in Continuous Availability?

- Users
- Applications
- Operations
- Database Administration
- Systems Administration
- Software bugs
- Data Integrity
- Performance
- Locking
- Continuous Operations



Availability Issues - Reduce planned & unplanned downtime

- Planned downtime
 - ▶ Data Sharing allows hardware and software changes to be non-disruptive to applications
- Unplanned downtime
 - ▶ Data sharing eliminates a single point of failure for a CEC, a z/OS LPAR, a DB2 subsystem
- Fast, non-disruptive database operations (even without data sharing)
 - ▶ Online Copy & Reorg
 - ▶ Concurrent Copy – allows you to create a non-fuzzy copy while applications update the data
 - ▶ Geographically Dispersed Parallel Sysplex
 - ▶ Faster restarts
 - ▶ Online parameters
 - ▶ Online LOAD resume
 - ▶ ...



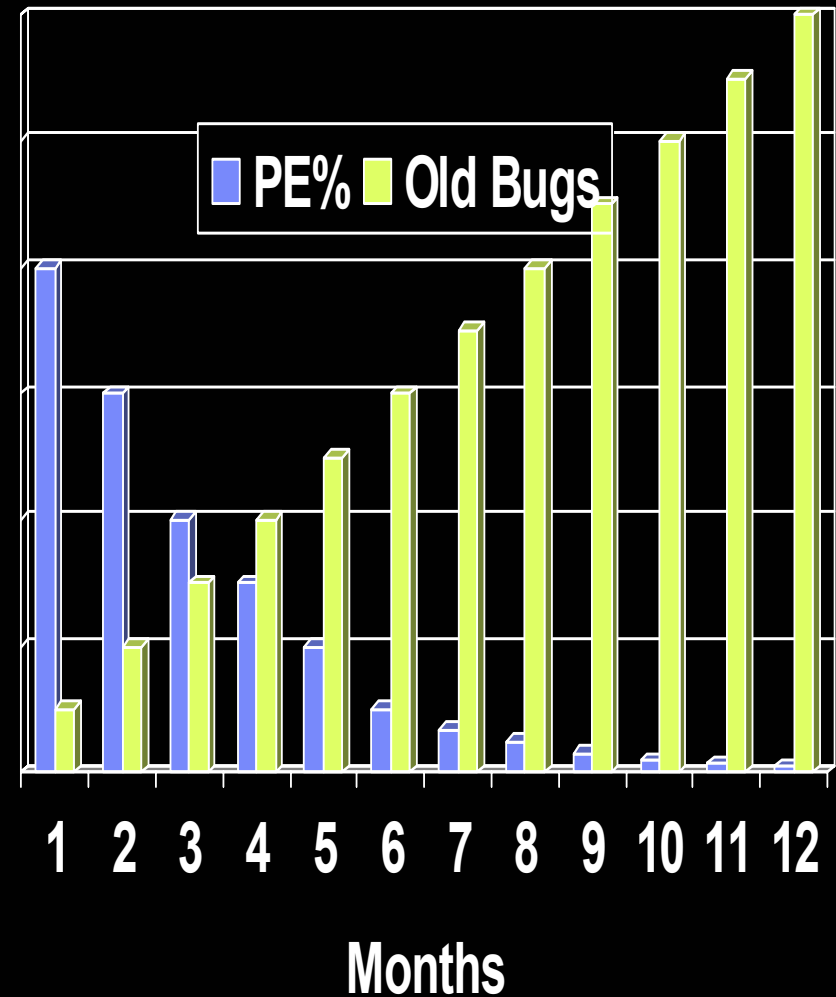
Availability Issues - Manage to Service Level Agreement

- Manage
 - ▶ Set availability criteria - balance needs & costs
 - ▶ Use criteria to determine procedures
- Practice operation and recovery
 - ▶ If not meeting criteria, then tune
 - ▶ Education for skills and time to do the job
- Stay reasonably current with DB2 releases & fixes
- Isolation and controlled replication
- Commit & Copy Frequency & Management



Availability Issues - Service Strategy

- Stay reasonably current with DB2 releases & fixes
- Balance for severity
 - ▶ Potential PTF in Error (PE)
 - ▶ Problems encountered instead of avoided
- Work load and windows for installing service
- Recommended process:
 - ▶ Consolidated Service Test levels
 - ▶ Upgrade service 3-4 times / year
 - ▶ Examine hipers and PEs weekly? monthly?
 - ▶ Stage through dev to test to prod



Best Practices for Restart

- Restart is the process of starting DB2 on an operating system logical partition (LPAR) after a normal stop, or an unexpected termination. The time required to restart DB2, after a termination, is dependent on a variety of factors:
 - ▶ How far back in the DB2 recovery log the last checkpoint was taken
 - ▶ The number of log records written from the last checkpoint for outstanding units of recovery (URs)
 - ▶ Whether objects are deferred for restart log apply processing
 - ▶ Whether postponed abort is used (limit backout)
 - ▶ How many datasets need to be opened for restart log apply
- Therefore...
 - ▶ Checkpoint frequently (2-5 minutes – the closer to two minutes, the better)
 - ▶ Don't have any long running (URs)
 - ▶ Consider deferred restart for selected objects or limiting the backout of long-running URs
 - ▶ Reduce the number of datasets open for write access
 - ▶ Use dual-logging and put each log/BSDS in different ICF catalog



Checkpointing and long running Units of Recovery (URs)

- To checkpoint based on logrecs or time?
 - ▶ Checkpointing based on number of log records written will give most consistent restart times
 - ▶ Consider checkpointing based on time interval when you have varying logging rates through time and restart time requirements are flexible
- Checkpoint every 2-5 minutes is a good general recommendation
- -SET LOG command can dynamically change checkpoint interval
 - ▶ LOGLOAD for # log records
 - ▶ CHKTIME for time interval (V7)
 - ▶ -DISPLAY LOG shows checkpoint frequency and status
- Last 100 checkpoints are recorded in the BSDS
 - ▶ Use DSNJU004 (Print Log Map) to display



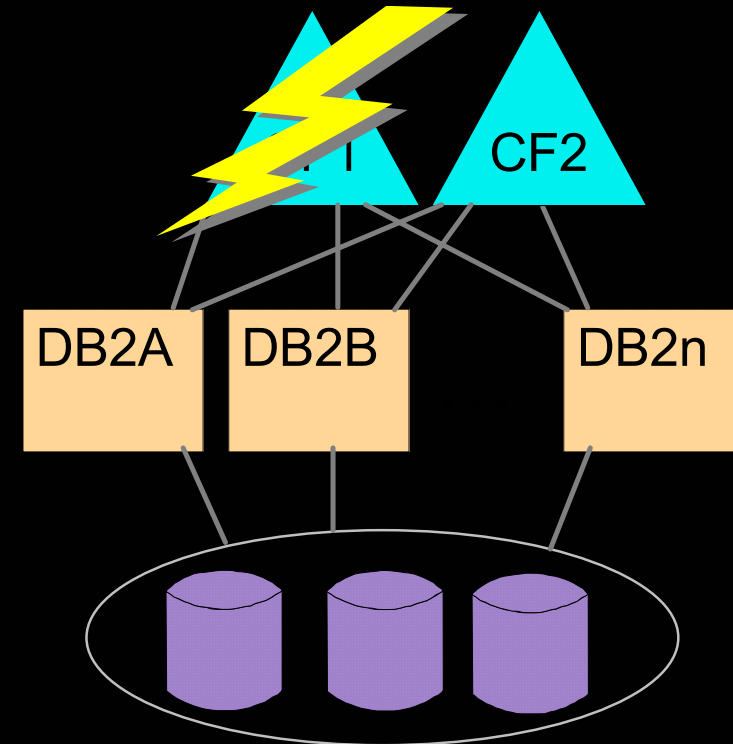
Checkpointing and long running Units of Recovery (URs)

- Identifying Long Running URs
 - ▶ Problems caused by long running URs:
 - Long restart times (use of Postponed Abort can mitigate this)
 - Long lock hold time can impact concurrency
 - Reduces effectiveness of lock avoidance checking
 - Prevents online Reorg and other utilities from running
 - ▶ Zparms control when DB2 issues long-running UR warnings:
 - UR CHECK FREQ , message DSNR035I (V5)
 - number of checkpoints that a UR has not committed
 - UR LOG WRITE CHECK, message DSNJ031I (V7)
 - number of log records that a UR has not committed
 - More granularity than UR CHECK FREQ
 - ▶ IFCID 0313 written, if active, when long-runner detected
 - Use IFCID 0313 to keep a history of problematic URs
 - ▶ Use UR CHECK FREQ for most accurate warning of restart impact
 - ▶ Use UR LOG WRITE CHECK to accurately identify applications that write a lot of log records
 - DB2 Accounting reports can also help identify these
 - V8: LONG-RUNNING READER THRESHOLD Zparm
 - Write IFCID 0313 when long-running reader detected



CF Duplexing

- GBP duplexing available since V5.
 - ▶ Recommendation: use it.
- Duplexing for Lock and SCA available in V7
 - ▶ Uses z/OS "system managed" CF duplexing which became GA in Feb '04
 - ▶ Useful to remove Internal Coupling Facility (ICF) as a single point of failure
 - ▶ Performance impact will depend mostly on intensity of CF lock requests for your application and distance between CFs
 - ▶ Pre-req's:
 - z/OS 1.2 or above plus maintenance
 - See CFDUPLEXING PSP Bucket
 - CFCC level 12 driver 3G (zSeries), or CFCC level 11 driver 26 (9672)
 - CF-to-CF links
 - ▶ <http://w3-1.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10272>



Virtual Storage Management

- Background
 - ▶ DBM1 "out of storage" is one of the leading causes of customer reported outages
 - Abend 04E/00E20003 or 04E/00E20016
 - ▶ Several drivers:
 - Bigger machines, higher workload volumes
 - Increasing use of dynamic SQL
 - New Java and Websphere workloads
 - Over-allocation of buffer pools, threads
 - Massive number of open datasets (with compression)
 - ▶ Largest consumers of DBM1 virtual include:
 - Buffer pools
 - EDM pool
 - Local dynamic statement cache
 - Thread storage
 - Compression dictionaries
 - ▶ DB2 has several enhancements to allow you to control virtual storage usage



Virtual Storage Management - some tuning tips

- If you are running short of DBM1 virtual, consider one or more of the following:
 - ▶ Monitor INFO APAR II10817 and apply recommended maintenance
 - ▶ Use data space buffer pools if running on a 64 bit machine
 - ▶ Use hiperpools if not on a 64 bit machine
 - ▶ Scale back usage of RELEASE(DEALLOCATE)
 - Version 8 removes need for REL(DEAL) to avoid XES lock contention
 - ▶ Set CTHREAD and MAXDBAT defensively
 - ▶ Set CONTSTOR=YES
 - ▶ Reduce number of open compressed datasets
 - ▶ Run IRLM PC=YES, reduce ECSA
 - With Version 8 "PC=YES" is forced
 - ▶ Consider making more use of data sharing to spread the work
 - ▶ Consider implementing Version 8 64bit virtual



Best Practices for Recovery

- Recovery is running RECOVER utility to recover one or more objects
 - ▶ The restore phase copies the data from the appropriate backup. If there are incremental backups, then they are also processed during this phase. The time for this phase is dependent upon:
 - Size of the objects (if flash copies cannot be used)
 - Media of the backups (tape, slow disk, fast disk, flashed copies)
 - ▶ The log apply phase positions the log to the point at which the backup was taken, and using a range index on the log (SYSLGRNX), scans the portions of the log when the objects were open for write access. The time for this phase is dependent upon:
 - The amount of log data that has to be scanned for applicable log records
 - The number of log records that need to be applied
 - Whether archive logs are needed, and whether archive logs reside on disk or tape
- Therefore...
 - ▶ Keep recent backups on fast disk (VTS considerations)
 - ▶ Take backups frequently enough to meet your recovery time objective
 - ▶ Keep 24-48 hours of log on disk
 - ▶ Use fast log apply



Data Recovery

- RECOVER utility time = restore time + log scan time + log apply time
- Restore time:
 - ▶ Number of pages, number of objects?
 - ▶ ICs on tape or DASD?
 - ▶ Degree of parallelism?
- Log scan time:
 - ▶ Image copy frequency
 - ▶ Archive logs needed to recover?
 - ▶ Log read from archive is not as efficient as from active (~25%)
 - ▶ Archive logs on tape or DASD?
 - ▶ Reads from DASD are faster
- Log apply time:
 - ▶ Update frequency and update patterns
 - ▶ Maximal fast log apply (4-10X difference)




Data Recovery - Recommendations for faster data recovery

- Reduce COPY cycle time
- For tape, take dual image copies to avoid image copy fallback
- Consider incremental COPY
 - ▶ Perform regular MERGECOPY of incremental COPYs
- COPY indexes
- Use DASD to write image copies and manage by DFSMS
- Keep at least 24-48 hours of recovery log on DASD
- Large, dual active logs
- Use VSAM striping for active logs
- Avoid access to archive log datasets
- Write archive log to DASD and manage by DFSMS
- IBM DB2 Archive Log Accelerator for z/OS can speed up arch log access
- Reorganize SYSLGRNX
- Large BP0 (at least 10,000 buffers) for recovering the cat/dir
- Use ESA Compression - be careful before V8 because of virtual storage constraint



GDPS PPRC HyperSwap

- Site failover
- DB2 data sharing group must still be recycled to recover CF structures
- New technology under development to avoid having to recycle DB2 members
- DASD failover
- New GDPS HyperSwap Manager announced 2/15/05
- Entry level GDPS offering
- Keep data available to end user apps during disk subsystem maintenance or failure
- DASD failure is transparent to DB2
- GDPS HyperSwap prerequisites
 - ▶ Disk subsystems that support PPRC level 3 (extended query)
 - ▶ IBM Tivoli System Automation for GDPS/PPRC HyperSwap Manager
- http://bvrghsa.ibm.com/projects/g/gdpsweb/hyperswap_manager.html



Near continuous availability for DASD failures or site failures

Achieving DB2 High Availability by DB2 release

- DB2 Version 6
- DB2 Version 7
- DB2 Version 8
- Futures




DB2 Version 6 Availability

- Minimize Unplanned down time by:
 - ▶ Faster, More Consistent Restart
 - Faster Log Processing
 - Ability to Start Even With Long Running UR
 - ▶ Faster Recovery for Data and Indexes
 - Faster Log Processing
 - Ability to Copy and Recover Index
- Minimize Planned down time by:
 - ▶ ALTER Partition Range and VARCHAR length
 - ▶ Set LOGLOAD parameter online
 - ▶ Faster, more parallel utilities
- Minimize Both Planned and Unplanned down time by:
 - ▶ Group Buffer Pool Duplexing



DB2 Version 6 Availability - Repartitioning

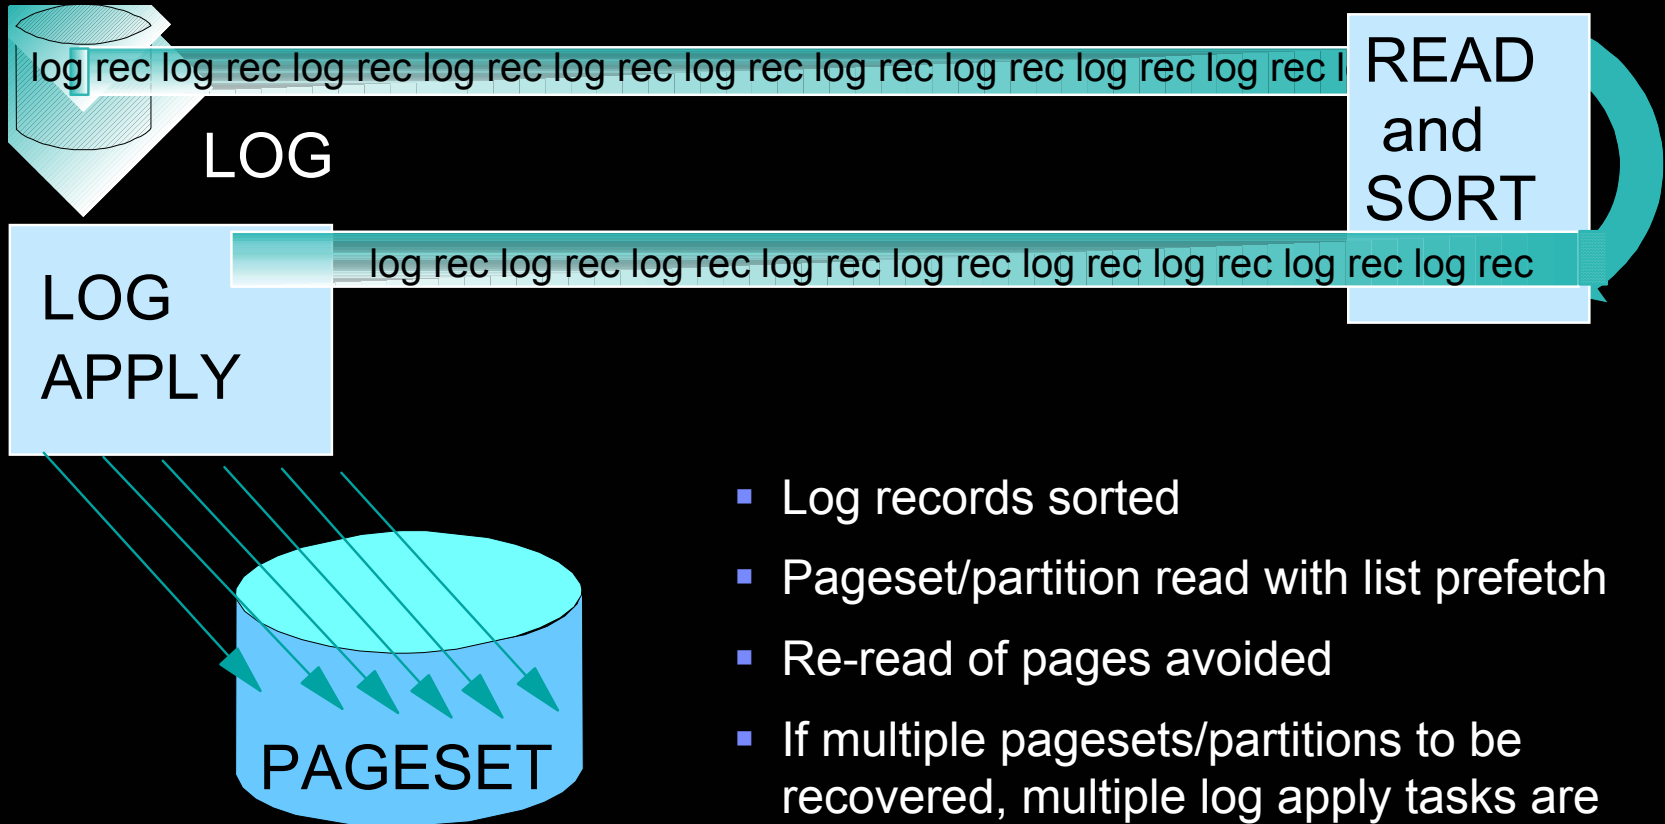
- Reduced down time and simplified operation
- ALTER INDEX enhanced to allow altering of limit key values on Partitioning Indexes i.e. "repartitioning"



```
ALTER INDEX IX1 ... PART 1 Values '2222'  
PART 6 Values '8500'
```

- Data Partitions affected are placed in REORP state
- (REORG Pending)
 - ▶ Partitions 1 and 2, and 6 and 7 in the example
 - ▶ REORG the affected partition range(s)
- Plans/Packages invalidated... rebind

DB2 Version 6 Availability - Fast Log Apply



- Log records sorted
- Pageset/partition read with list prefetch
- Re-read of pages avoided
- If multiple pagesets/partitions to be recovered, multiple log apply tasks are created

DB2 Version 6 Availability - Restart Processing option

- The problem:
 - ▶ Long running UR's which must be backed out at system restart
- A solution:
 - ▶ Allow some backouts to run AFTER restart is completed

- Two new install parameters:
 - ▶ LIMIT BACKOUT (LBACKOUT)
 - ▶ BACKOUT DURATION (BACKODUR)
 - Multiplier for LOGLOAD - max number of records to be read during Restart's backward log scan phase
 - Insures that short UR's are backed out at restart

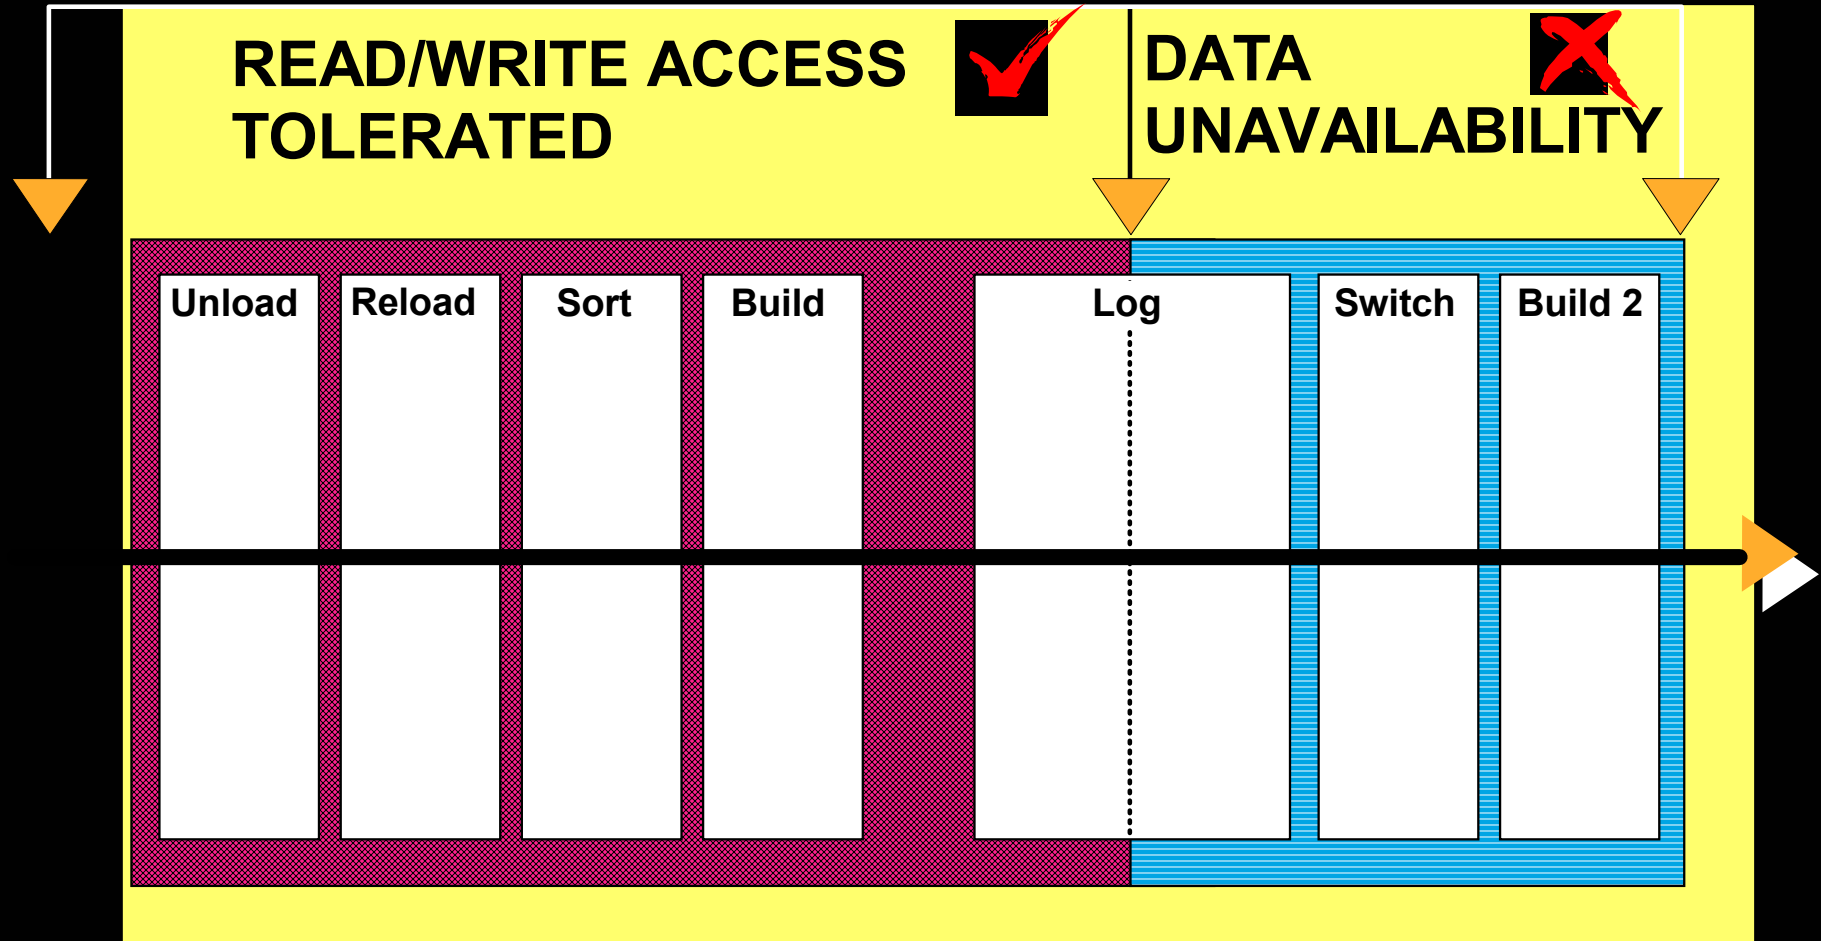


DB2 Version 7 Availability

- Online utility improvements
 - ▶ Online Reorg
 - Avoid data set rename
 - Parallel BUILD2 phase
 - ▶ Online Load resume
 - ▶ Timeout and retry options
 - Reduce impact on concurrently running applications
- Unload from image copy or shrlevel change
- Online subsystem parameters
- Consistent restart enhancements (ability to cancel long backouts)
- Faster utilities: Unload, Load parallelism
- Ability to cancel thread without rollback
- Real Time Statistics – help manage when copies should be run
- Improved Storage Information



DB2 Version 7 Availability – Online REORG




DB2 Version 7 Availability – Online REORG FASTSWITCH

- With V6:
 - ▶ Shadow copies, RENAME during SWITCH phase
 - ▶ Data sets deleted in UTILTERM phase
- With V7 (Driven by FASTSWITCH parameter on reorg utility):
 - ▶ DB2 will allow two naming conventions for dataset names ('I0001' or 'J0001')
 - ▶ Current instance node registered in catalog & directory
 - ▶ At UTILTERM J or I data sets are deleted
Difference on the ESS disks was about 92% faster or about 12 times faster.
 - ▶ Avoids costly data set renames during switch phase
 - ▶ Impacts stand-alone utilities, SMS, HSM and RACF definition, Customer procedures that specify dataset name)

FASTSWITCH	NO		YES		Delta %	12 X
REORG SHRLEVEL CHANGE	CPU (sec)	Elapsed (sec)	CPU (sec)	Elapsed (sec)	CPU	Elapsed
	1	349	1	28	0	-92

DB2 Version 7 Availability - BUILD2 parallelism

- BUILD2 phase of online REORG rebuilds non partitioning indexes
- V7: adds parallelism
- With 1 NPI improvement about 30% in CPU & elapsed time
- With 5 NPIs built by parallel tasks about 80% elapsed time
- V8 ability to eliminate BUILD2 phase with DPSI



DB2 Version 7 Availability - On-line LOAD RESUME

Syntax

```
LOAD DATA
RESUME YES
SHRLEVEL CHANGE
INTO TABLE TEST.CARS
( CARID      POSITION ( 1: 3) CHAR
, PRODUCER  POSITION ( 5:17) CHAR
, MODEL     POSITION (20:29) CHAR
...

```

```
001 MOTOR CORP.    STAR
002 ENGINE MASTER  SPEEDY
003 MOBILES, LTD.  CONFOR
...

```

- Security: LOAD (not INSERT) privilege
- Timeout of the LOAD: Like utility (not like SQL appl.)
- Uses LOAD PART parallelism
- Intelligent commit interval

```
INSERT INTO TEST.CARS
VALUES ('001'
, 'MOTOR CORP.'
, 'STAR' );
INSERT INTO TEST.CARS
VALUES ('002'
, 'ENGINE MASTER'
, 'SPEEDY' );
INSERT INTO TEST.CARS
VALUES ('003'
, 'MOBILES, LTD.'
, 'CONFOR' );
...

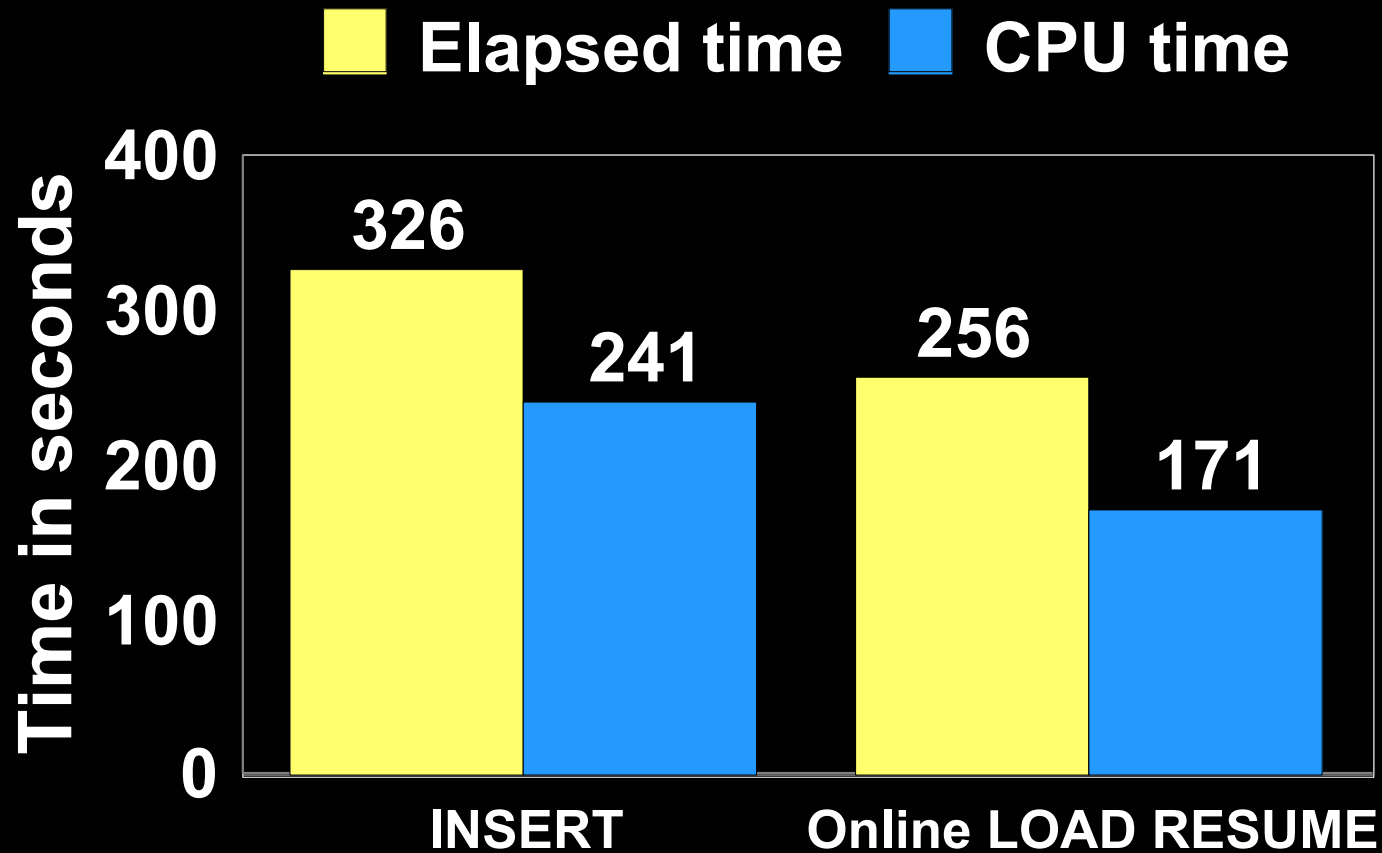
```

Processing

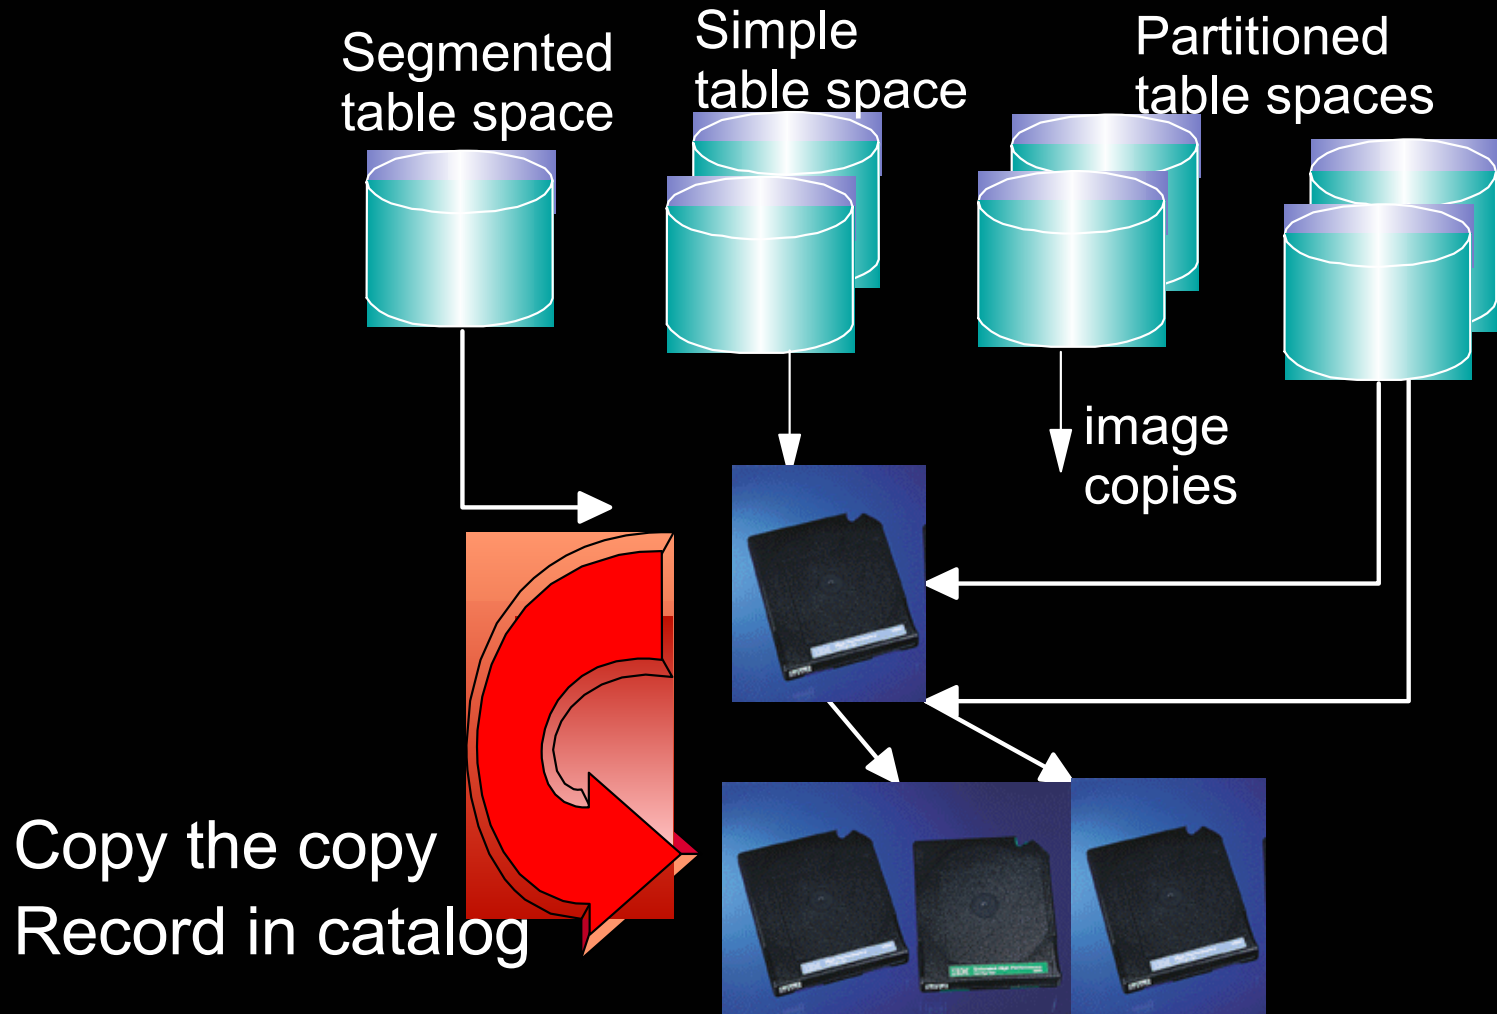
- **Serialization:** Claims (not drains)
- **Logging:** LOG YES only
- **Triggers:** Fire
- **RI:** Parent key must exist
- **Duplicate keys:** First accepted
- **Clustering:** Preserved
- **Free space:** Used (not provided)



DB2 Version 7 Availability - Online LOAD RESUME

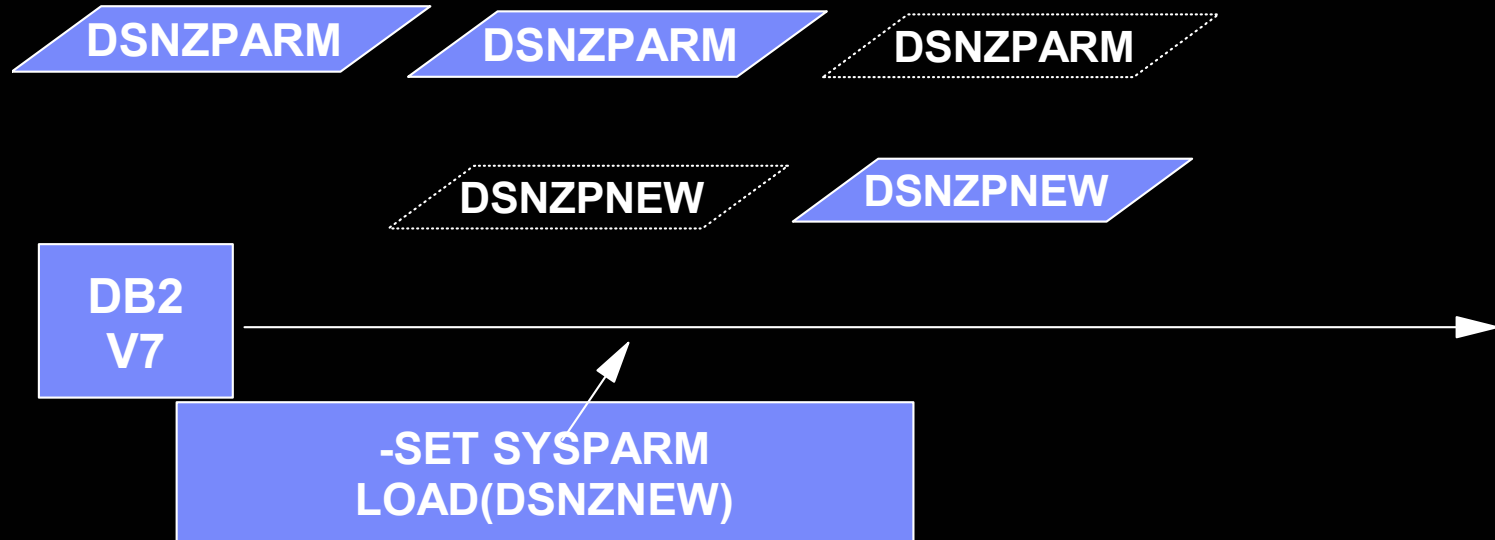


DB2 Version 7 Availability – Copy To Copy



DB2 Version 7 Availability - Online subsystem parameters

- Allow many parameters to be changed while DB2 is up
- Load module granularity (multiple DSNZPARM load modules)
- Can change over 60 values
- Restart of DB2 resets all values to startup DSNZPARM



DB2 Version 7 Availability - Adding workfiles

- Work files are created with the sequence:

~~-STOP DATABASE (DSNDB07)~~

CREATE TABLESPACE DSN4K01 IN DSNDB07

BUFFERPOOL BP0 CLOSE NO USING VCAT C710;

CREATE TABLESPACE DSN32K01 IN DSNDB07

BUFFERPOOL BP32K CLOSE NO USING VCAT C710;

~~-START DATABASE (DSNDB07)~~

- Now you do not need to STOP the DSNDB07 and production to add or remove a table space



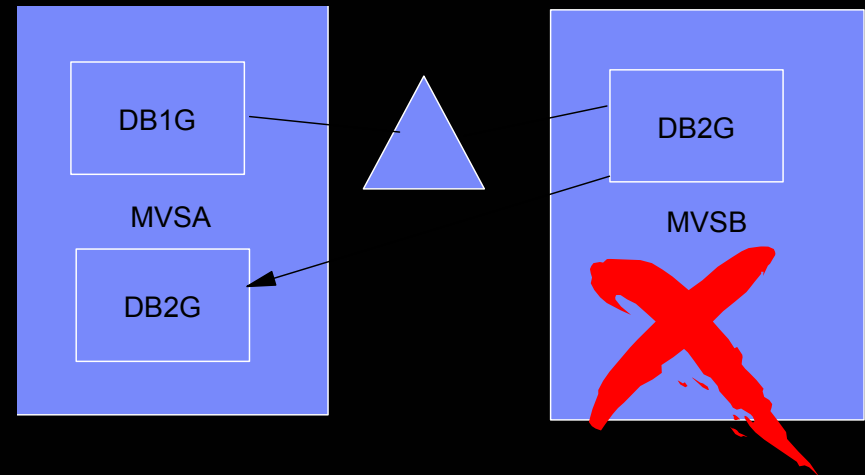
DB2 Version 7 Availability – Real-time Statistics

- Real-time statistics help DBA or Monitor Programs to identify objects that need database maintenance
 - ▶ Efficient use of DBA time
 - ▶ Improve system and application performance
 - ▶ Reduce batch window
- DSNACCOR easy way to identify DB2 objects that
 - ▶ Need to run REORG / RUNSTATS / COPY
 - ▶ Are in restricted access states (e.g. COPY pending, ...)
- Real-time statistics provides a foundation for DB2 to manage itself in the future
- See Appendix G & H in latest Administration Guide



DB2 Version 7 Availability - DB2 Restart Light details

- Use to recover from LPAR failures
- Quickly restart a failed DB2 member on another LPAR with minimal disruption to release retained locks
- Brings up DB2 with minimal storage footprint
- Terminates normally after retained locks are freed
- On restart:
 - ▶ No EDM and RID pools, LOB manager, RDS, RLF
 - ▶ Reduced number of service tasks
 - ▶ PRIMARY bufferpools only, No Hiperpools,
 - $VPSIZE = \min(VPSIZE, 2000)$
 - ▶ Short VDWQ's
 - ▶ Castout(NO) used for shutdown
 - ▶ If autostart IRLM, will override IRLM to PC=YES



DB2 Version 7 Availability - Consistent Restart Enhancements

- What's new:
 - ▶ CANCEL option on RECOVER POSTPONED
 - ▶ Rollback of ALL postponed abort URs are canceled
 - ▶ Cannot be used if CATALOG/DIRECTORY updated in UR
 - ▶ DSNV435I message indicates cancel has been scheduled
 - ▶ DSNV042I message indicates Rollback for PA URs has been canceled
- NOBACKOUT option on CANCEL (DDF) THREAD
 - ▶ Can be used to cancel a long running thread without backing out the changes
 - ▶ DSNV042I message indicates Rollback for the indicated thread has been canceled
 - ▶ Also applies to 'regular' IN-ABORT threads
- New REFP state which can be cleared by
 - ▶ recovery to point in time prior to exception
 - ▶ load replace
 - ▶ start db() space() access(force) BAD IDEA
- **WARNING: WHEN YOU CANCEL ROLLBACK YOU TAKE RESPONSIBILITY FOR DATA INTEGRITY!!!**



DB2 Version 8 Availability

- Schema evolution
 - ▶ Alter some data attributes with minimal disruption
 - ▶ Add partitions, rotate partitions, extend column
- Up to 4096 partitions
- Partitioned secondary indexes
- Additional online parameter changes
- 64 bit virtual storage
- System-Level Point in Time Recovery
- Utility Improvements
- Increase in Active and Archive Logs
- Additional Online Zparms
- Sliding Secondary Allocation Size
- Automatic LPL recovery (no more `–START DB() SPACENAM())`
 - ▶ If failure (MSGDSNB357I) , then `RECOVER LOGONLY`
- ...



DB2 Version 8 Availability – Schema Evolution

- Allows altering of many table, tablespace, and index attributes without impacting application availability
- Before you had to drop/recreate or stop the object
- Support in previous releases
 - ▶ V4: allow alter of CLOSE or ERASE attribute without invalidating dependent objects
 - ▶ V5: allow alter of PRIQTY and SECQTY without requiring stop
 - ▶ V5: allow expanding length of VARCHAR columns, even if indexed
 - Index versioning
 - ▶ V5: DDL drains instead of X lock on tablespace
 - ▶ V6: allow changing partition boundaries without dropping table
 - ▶ V7: allow dynamically add more workfile TS's
 - ▶ V7: allow alter LOCKSIZE for non-linked catalog objects



DB2 Version 8 Availability – Schema Evolution

- Changing Partition Definitions
 - ▶ Add Partition
 - ▶ Rotate Partition
 - Move first part to last
 - E.g. keep rolling 36 months
- ALTER INDEX : ADD COLUMN, PADDED or NOT, CLUSTER or NOT



DB2 Version 8 Availability – Schema Evolution

- Table or column Changes
 - ▶ Increase column within numeric data types
 - smallint, integer, decimal, float
 - No loss of precision allowed
 - ▶ Change to expand character data type
 - ▶ Change varchar to / from char



DB2 Version 8 Availability – Schema Evolution

- Unbundling Partitioned Table Attributes
 - ▶ Partition without an index
 - May be able to have one less index
 - ▶ Data Partitioned Secondary Index
 - ▶ Cluster on any index
 - May be able to have more efficient clustering
 - E.g. partition by date, cluster by account
- REORG automatically rebalance partitions



DB2 Version 8 Availability - Data Partitioned Secondary Index

- Physical partitions like table
- Benefits
 - ▶ Eliminate REORG BUILD2 phase
 - ▶ Improve index recovery
 - ▶ Facilitate parallel partition processing
 - ▶ Eliminate LOAD PART contention
 - ▶ Reduce Data Sharing overhead
- Some tradeoffs for index use
 - ▶ Need to check multiple partitions
 - ▶ Not for unique indexes



DB2 Version 8 Availability - Data Sharing Enhancements

- Batching of GBP writes and castouts
 - ▶ Improved data sharing performance, especially for batch updates
- Reduced global contention for tablespace L-locks
 - ▶ Improved data sharing performance, especially for OLTP
 - ▶ Recommendation for RELEASE(DEALLOCATE) can be rescinded
- Changed pages written to GBP at Phase1 instead of Phase2
 - ▶ Removes need for IMMEDIATE PH1 Zparm or Bind option
 - ▶ Equivalent performance for Ph1 vs. Ph2 writes
- More efficient index split processing for data sharing



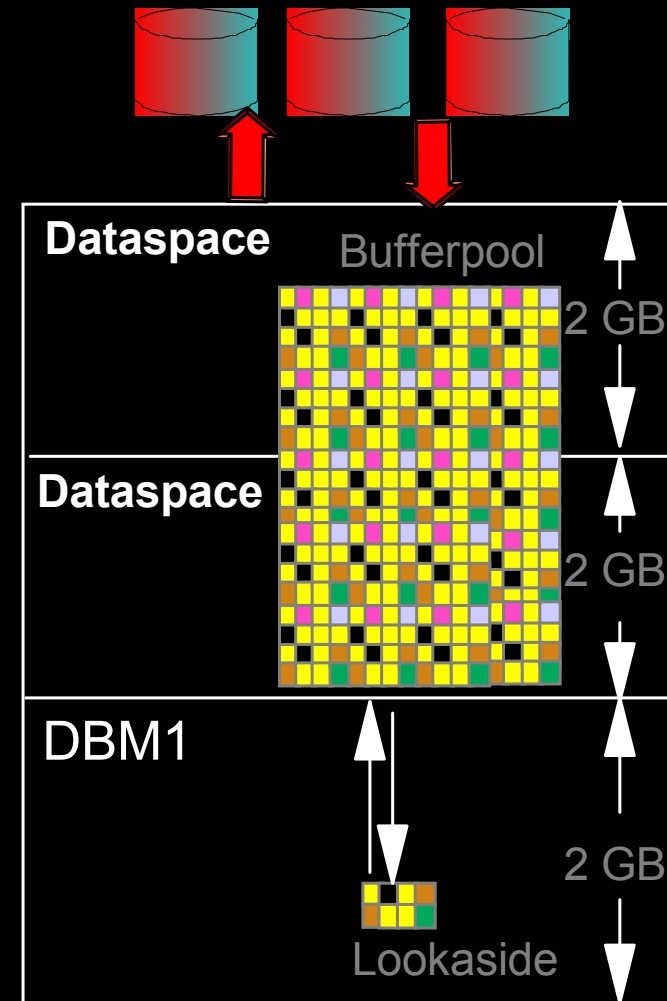
DB2 Version 8 Availability - Automatic Space Management

- Sliding secondary allocation quantity size
- Applies to DB2 managed pagesets only
- Tries to avoid VSAM maximum extent limit errors
- Can reach maximum dataset size before running out of extents
- Uses cylinder allocation
 - ▶ Default PRIQTY
 - 1 cylinder for non-LOB tablespaces and indexes
 - 10 cylinders for LOB tablespaces
 - ▶ Improved SQL SELECT and INSERT performance
 - 2x improvement relative to track allocation
- Can be used for
 - ▶ New pagesets: No need for PRIQTY/SECQTY values
 - ▶ Existing pagesets: SQL ALTER PRIQTY/SECQTY values to -1 plus schedule a REORG



DB2 Version 8 Availability - 64 bit evolution

- zSeries, z/OS, z/Architecture & large real storage
- Real storage for performance
 - ▶ Improve all versions
 - ▶ V6 data space advantages
 - zSeries & (z/OS or OS/390 V2R10 & 64 bit addressing)
- 64 bit virtual
 - ▶ See z/OS 64-bit Virtual Storage Roadmap
- Many incremental changes ...
- Avoid storage outages



DB2 Version 8 Availability - 64 bit evolution

- Why 64-bit?
 - ▶ Needed for DB2 to continue to scale on a single OS image
 - ▶ More concurrent threads
 - ▶ Bigger buffer pools
 - ▶ Large memory exploitation
- DBM1 and IRLM are 64bit address spaces in Version 8
- The following data areas are moved above the 2GB "bar"
 - ▶ Buffer pools, BM control blocks
 - ▶ Castout buffers
 - ▶ RID pool
 - ▶ EDM DBD cache (OBDs)
 - ▶ Global dynamic stmt cache
 - ▶ Sort pool
 - ▶ Trace tables (Global, Lock, BB)
 - ▶ Accounting blocks
 - ▶ Compression dictionaries
 - ▶ IRLM locks
- Max buffer pool size is lifted to 1 TB
 - ▶ New PGFIX(YES) option can give significant performance benefits
- Data space pools, hiper pools are eliminated - easier management



DB2 Version 8 Availability – System Level PIT Recovery

- Easier, more flexible, less disruptive, faster recovery
- Handle large numbers of table spaces & indexes
- Two new utilities introduced:
 - ▶ BACKUP SYSTEM: Fast volume-level backups
 - DB2 databases and logs
 - Data sharing group scope
 - z/OS V1R5, DFSMSHsm Fast Replicate, DFSMSdss, & FlashCopy required (FlashCopy Version 2 recommended)
 - ▶ RESTORE SYSTEM
 - To an arbitrary point-in-time
 - Handles creates, drops, LOG NO events



DB2 Version 8 Availability – Improved Data Integrity Checking

- LOBs
 - ▶ CHECK DATA to unmark invalid LOBs
 - Improved availability, PQ77366 (V7)
 - ▶ LOB space reuse tracking
 - Improved diagnostic logging for data corruption, PQ77368 (V7)
 - ▶ COPY CHECKPAGE enhanced for LOBs in PQ74793
- Do not set COPYP if CHECKPAGE detects an error
 - ▶ Implemented in V8
- DSN1LOGP CHECK(DATA) option
 - ▶ Check for regressed data page problems by analyzing the log
 - ▶ V8 NFM: enhanced to also handle indexes
- CHECK INDEX option to check integrity of non-leaf pgs (V8)
- Online Check Index capability
 - ▶ V8 APAR PQ92749



DB2 Version 8 Availability – Online Reorg Enhancements

- REBALANCE on REORG SHRLEVEL NONE|REF
- DISCARD on REORG SHRLEVEL CHANGE
- No BUILD2 phase if all secondary indexes are DPSIs
- REORG SHRLEVEL REFERENCE on all DB2 Catalog Tables
- PQ96628 avoid OLR drain failures



DB2 Version 8 Availability – Post-GA Enhancements

- Utilities performance enhancements
 - ▶ PQ86532 - Unload Performance
 - ▶ PQ86311 - RUNSTATS non-padded index performance improvement
 - ▶ PQ84160 - DB2 Cross Loader performance improvement
 - ▶ PQ89685 - DB2 UTILITIES Performance Enhancement: UT Serial Lock
 - ▶ PQ92911 - LOAD PART REPLACE parallelism performance
 - ▶ PQ90884 - RUNSTATS DSTATS Performance
 - ▶ PQ96628: Online REORG claim/drain deadlock avoidance
 - ▶ PQ92749 - Online CHECK INDEX
 - ▶ PQ90263 - Crossloader support for larger LOBs
 - ▶ Upcoming: LOAD/UNLOAD support for larger LOBs



DB2 Version 8 Availability – Post-GA Enhancements

- Locking and concurrency
 - ▶ PQ89070 - Lock Avoidance for Singleton Select with ISO(CS) and CURRENTDATA(YES)
 - ▶ PQ87381 - INSERT performance enhancement to help avoid need for REORG - not yet delivered
- Managability
 - ▶ PQ88983 - Allow Cancel Rollback after log data set access error
 - ▶ PQ90936 - DISPLAY ARCHIVE enhancement to show DASD archive log read and also HSM RECALL waits
 - ▶ PQ93548 - System Point-In-Time or Disaster Recovery - scan log for prior checkpoint when log truncation is prior to last checkpoint in the BSDS queue

DB2 Version 8 Availability

- RAS and monitoring enhancements
 - ▶ PQ83649 DISPLAY THREAD SERVICE(WAIT)
 - ▶ PQ88073 - new EXPLAIN STMTCACHE ALL option, results written to DSN_STATEMENT_CACHE_TABLE
 - ▶ PQ91101 - Improve statement cache and real storage statistics
 - ▶ PQ90547 - DDF/RRSAF Accounting Rollup Enhancements
 - ▶ PQ93009 - Add new criteria for DSNACCOR Stored Procedure to recommend RUNSTATS
 - ▶ PQ99524 - improve robustness of DB2 when stored procedures are cancelled
 - ▶ PQ98043 - compress stack storage when it gets large



DB2 Future Availability

- Online Schema enhancements
- Online 'LOAD REPLACE'
- Faster, more consistent DB2 restart
- WLM synergy for system robustness at high utilization
- Self tuning BPs
- 64-bit enhancements
- Online REORG BUILD2 avoidance and parallelism
- Online CHECK LOB/DATA
- REORG LOB SHRLEVEL(REF)
- Volume-level backup and recovery
- Archive log enhancements
- Dynamic change of early code
- Broken page recovery enhancements
- Outage Analysis items



Availability Summary

- Substantial improvements in New DB2 Versions
- Service Level Agreements differ substantially
- Availability is dependent upon your database, application & subsystem design & practices
- Much more information available
 - ▶ SG24-5486 DB2 for OS/390 and Continuous Availability
 - ▶ SG24-6129 DB2 V7 Performance Topics
 - ▶ Redpaper DB2 V7 Selected Performance Topics
 - ▶ DB2 for z/OS & OS/390
 - Web: Access or Download PDF files
 - ibm.com/software/db2zos +presentations.html Then click Library
 - ▶ Red Books International Technical Support Organization
 - ibm.com/software/db2zos Then click Support, then Redbooks then put book number or topic or DB2 into Search and click Go.
 - New Tools for DB2 for OS/390 & z/OS Presentation Guide, SG24-6139
 - DB2 for OS/390 & z/OS Version 7 Presentation Guide, SG24-6121
 - DB2 Version 6 Performance Topics, SG24-5351
 - Storage Management and DB2, SG24-5462
 - Version 6 Technical Update, SG24-6108



Disclaimers & Trademarks*

Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion.

You should not rely on such information to make business plans. Any discussion of OEM products is based upon information which has been publicly available and is subject to change.

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries: AIX, AS/400, DATABASE 2, DB2, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, SYSTEM/390, SQL/DS, VM/ESA, IBM, Lotus, NOTES. The following terms are trademarks or registered trademarks of the MICROSOFT Corporation in the United States and/or other countries: MICROSOFT, WINDOWS, ODBC

