



SWG BetaWorks

DB2 9 for z/OS

Technical Education Series

“Data Sharing”

BetaWorks

Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR
- ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

Agenda

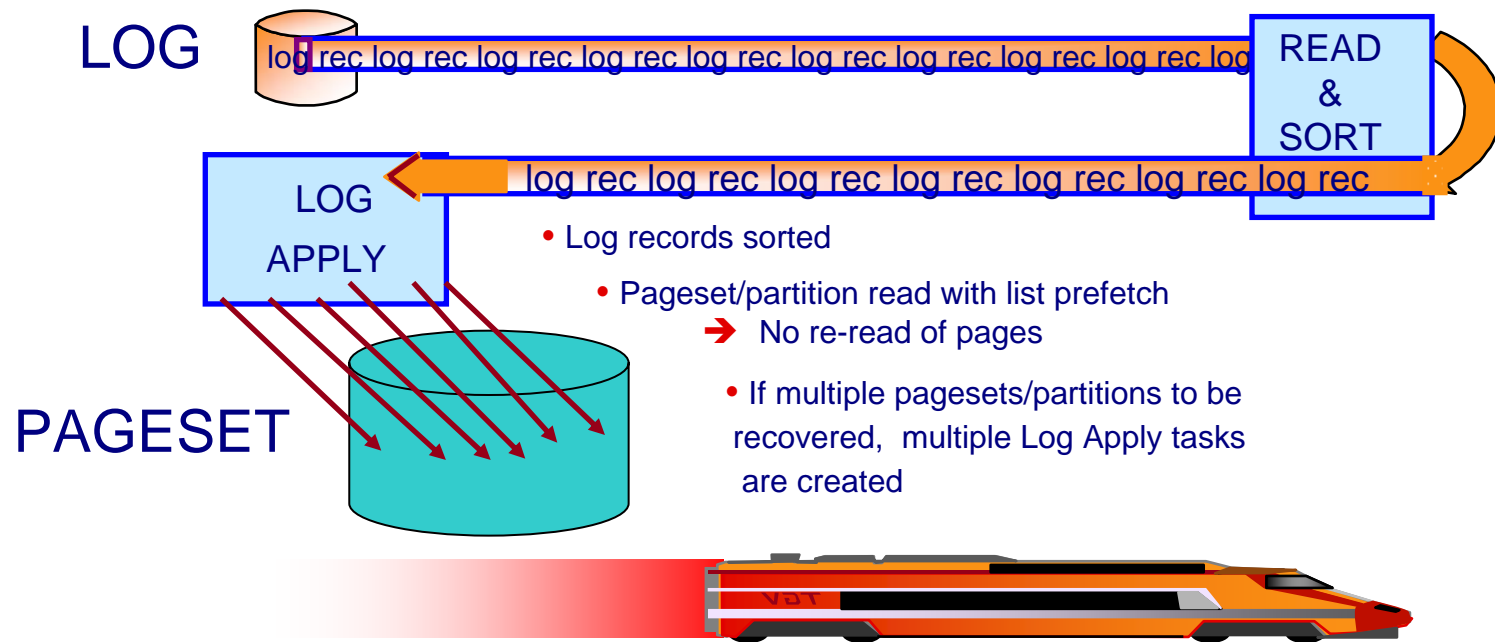
- Initiate object open as early as possible
- Simplify special open processing
- Deferred SYSLGRNX updates until after end restart
- Initiate Auto GRECP recovery at end of restart
- Allow table level retain locks to support postponed abort URs.
- Data Sharing logging improvement
- Further additional enhancements



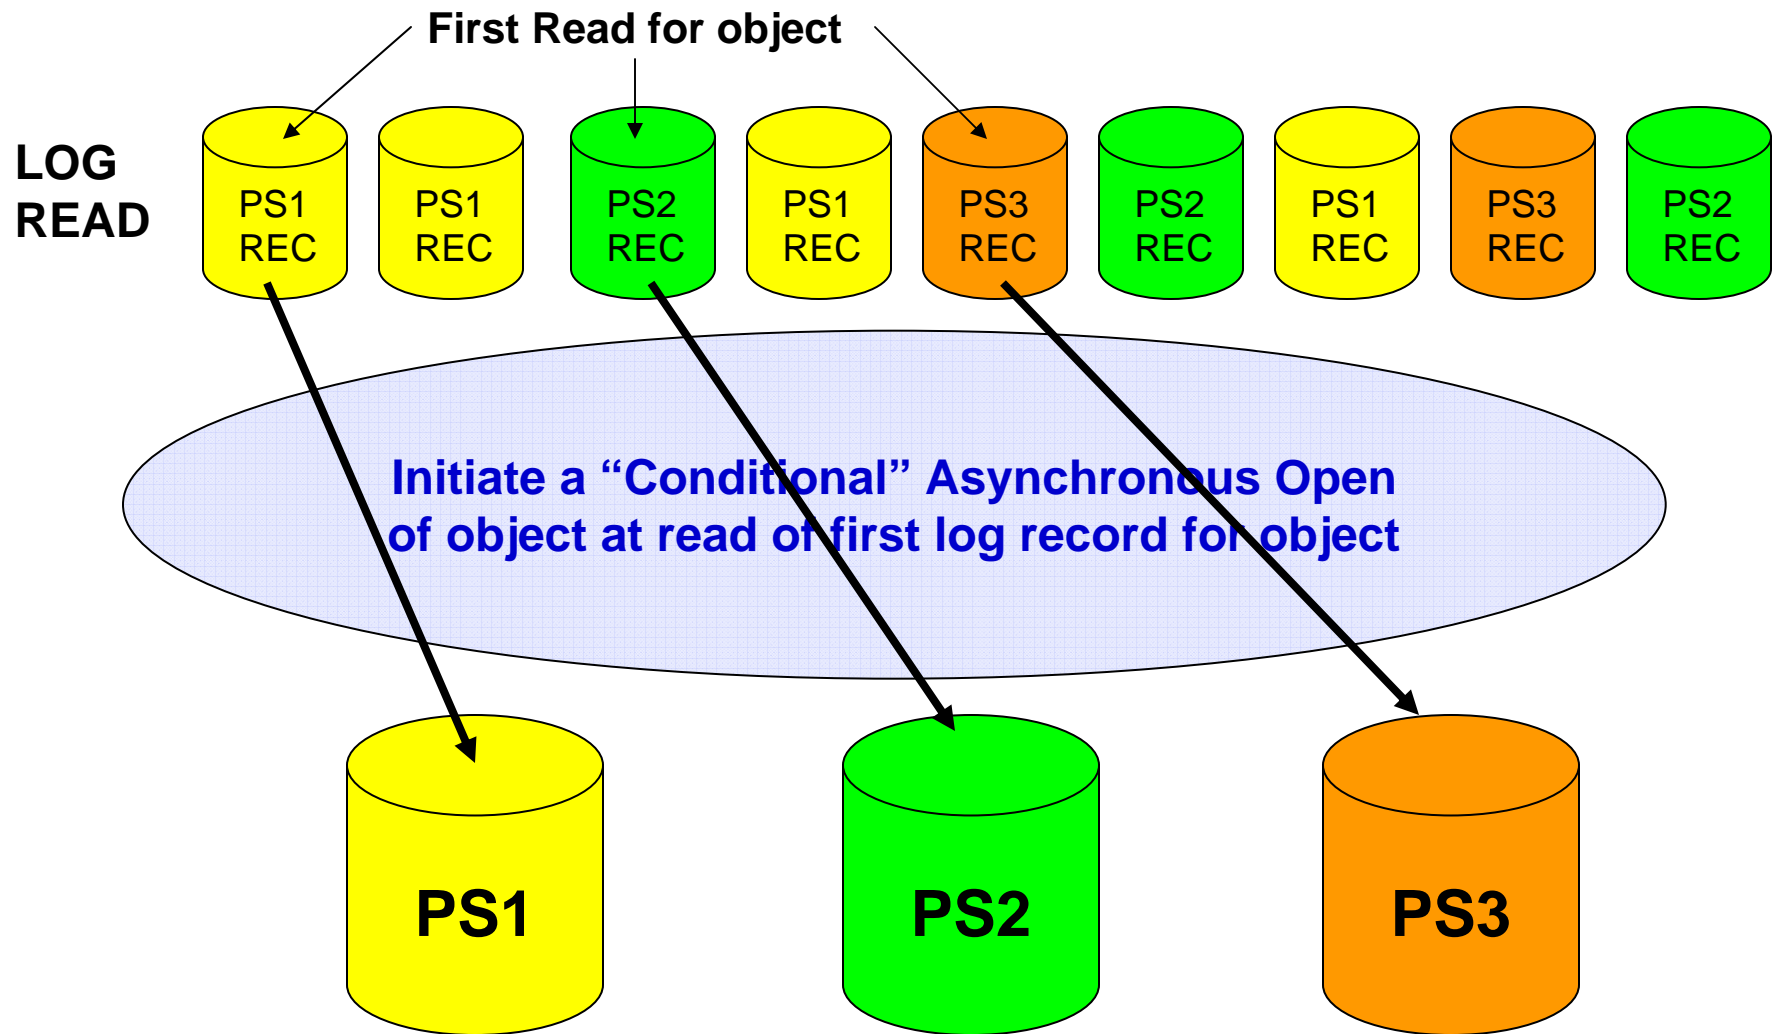
**Improving performance,
usability and availability
during DB2 restart.**

Initiate Object OPEN in FLR phase

- **Fast Log Apply** is always enabled during restart.
- When we are reading the logs, we are **not** applying log records.
 - If we know we have log records to apply to a particular dataset, then lets take advantage and open the dataset earlier - before we want the pages of that dataset.



Earlier Pageset Opens



Improved restart Log Apply time

- Initiate **asynchronous open** during Forward Log Recovery (FLR) phase
 - When the first log record of the object is read

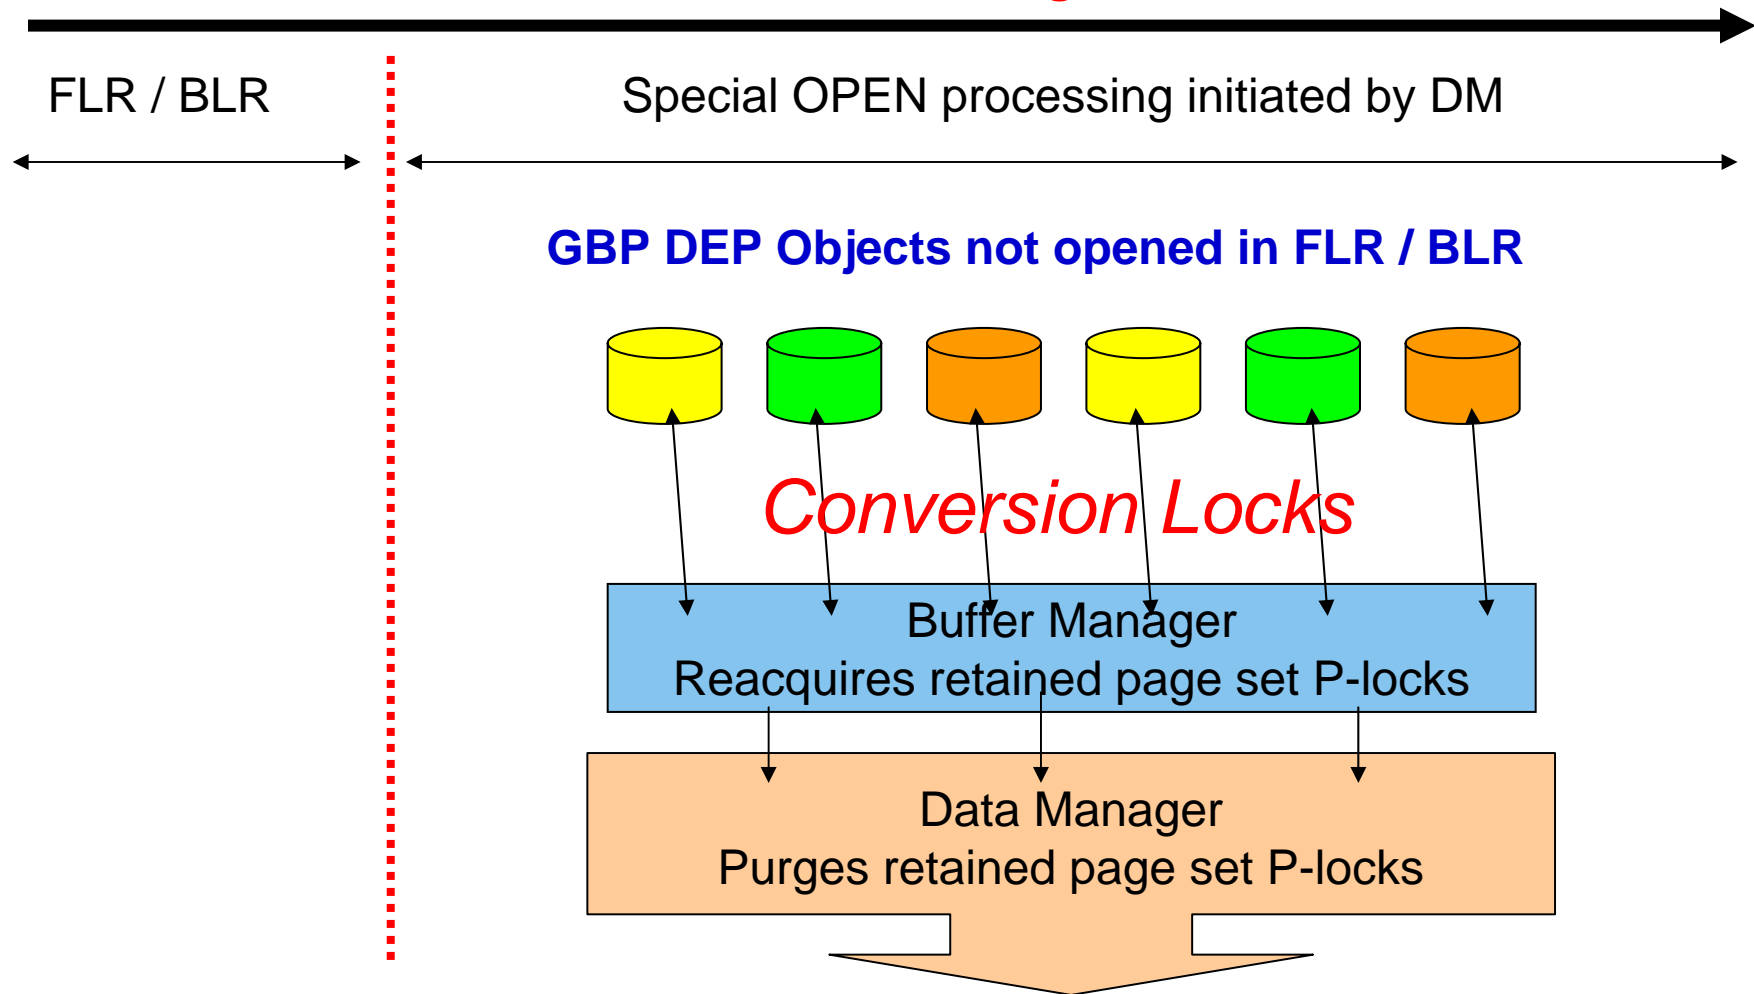
- **Conditional OPEN** request
 - If Log apply task tries to initiate a **synchronous open** request at the same time as the **asynchronous** open request is scheduled
 - **Synchronous** open **MUST** always succeed

- This is a performance enhancement, especially where we have 1000's of objects involved in the restart.
 - Remember there are limits of 20 to 40 parallel dataset opens with VSAM.

- **Important exceptions**
 - DB2 restarted with DEFER ALL option
 - DB2 restarted in System Level PITR mode
 - Tracker mode

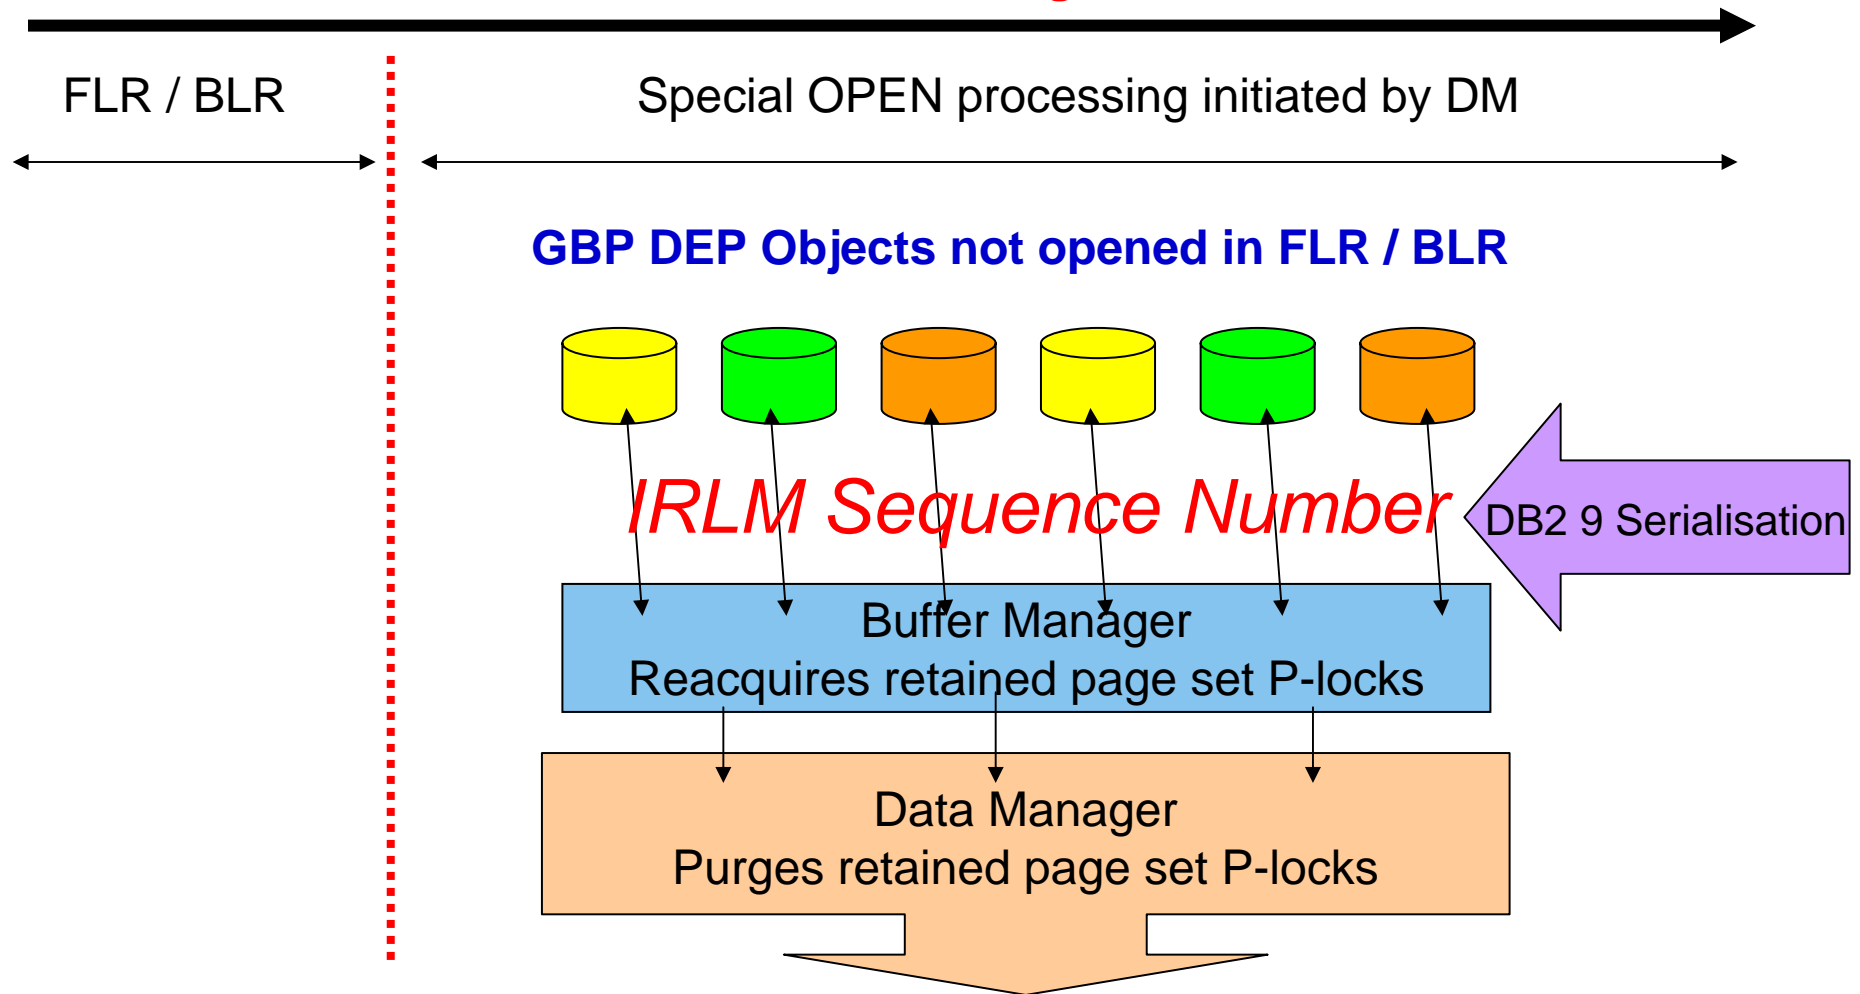
Special OPEN processing – V8

RESTART Processing in V8



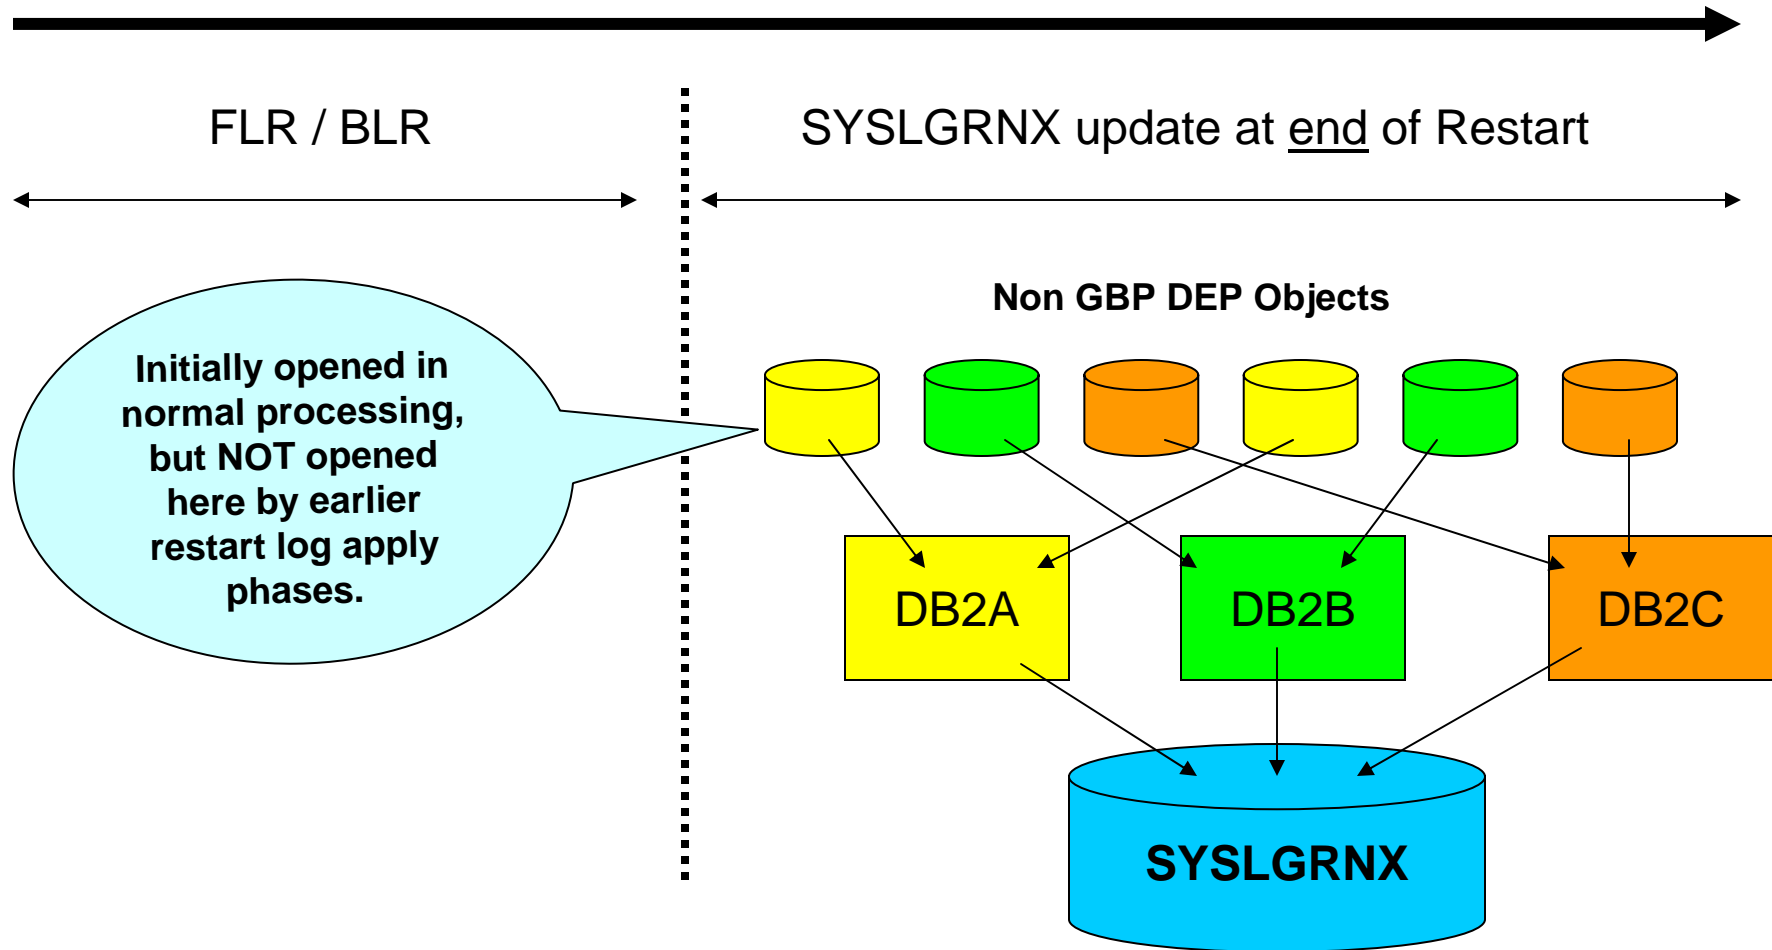
Simplifying special OPEN processing – DB2 9

RESTART Processing in DB2 9



SYSLGRNX update during restart – V8

RESTART processing in V8



Deferred SYSLGRNX update until after end of restart – DB2 9

- In V8, SYSLGRNX updates can queue behind other members making their updates
 - Increases the RESTART time

- In DB2 9 we will defer any SYSLGRNX update until the **pseudo close time**
 - Pseudo close triggered by first checkpoint after restart
 - Buffer Manager will close the syslog range.

- Benefits both data sharing and non data sharing systems

- Important exceptions
 - DB2 restarted with DEFER ALL option
 - DB2 restarted in System Level PITR mode
 - Tracker mode

Initiate Auto GRECP recovery after end restart

- Automatic Recovery of GBP DEP objects first introduced in V5, when:
 - Connectivity of ALL members to GBP is lost
 - GBP structure failure occurred

- But ONLY for failures occurring while DB2 is up and running
 - NOT for recovery during Group or Normal RESTART

- **The DB2 9 enhancement** will initiate auto GRECP recovery after DB2 has gone through a group or normal restart which involves GBP failure
 - Triggers after each member completes group restart

- Will save users writing automation / REXX to pull objects in GRECP
 - Spread –STA DB around members based upon castout ownership
 - 50 parallel commands – 20 objects (max) per command

Initiating recovery of GRECP objects after end restart

- The GBP **AUTOREC** setting determines auto GRECP initiation after restart

- Each member initiates recovery of GRECP objects which it knows about
 - Based upon the GBP dependent state (R/W Interest) of an object when DB2 came down

- Each member tries to acquire a conditional drain for each GRECP object
 - If successful, GRECP recovery of the object continues with log apply
 - Otherwise object remains in GRECP and DB2 issues DSN1005I
 - A possible reason could be that another member is still restarting and holding retain locks.
 - In this case that member will initiate GRECP recovery when it reaches its own end of restart

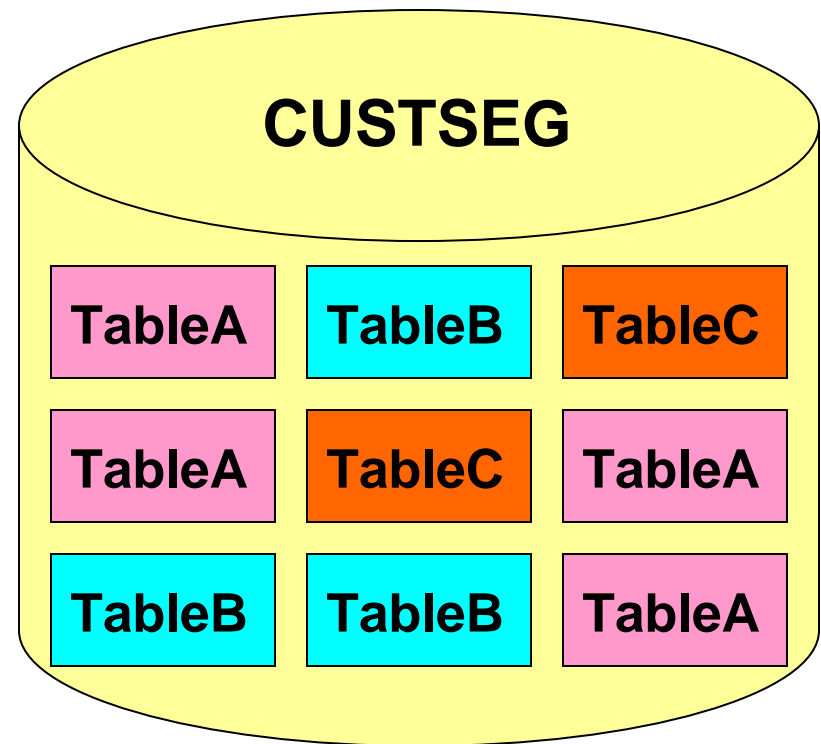
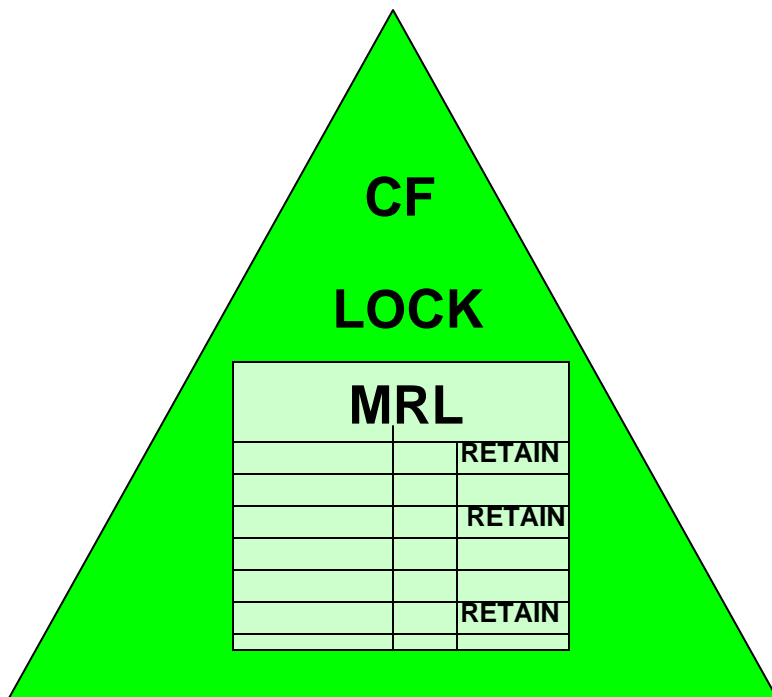
Additional information

- Disaster Recovery (DR) always involves a group restart
 - Hence auto GRECP recovery can be initiated at end of restart for DR members, if AUTOREC set

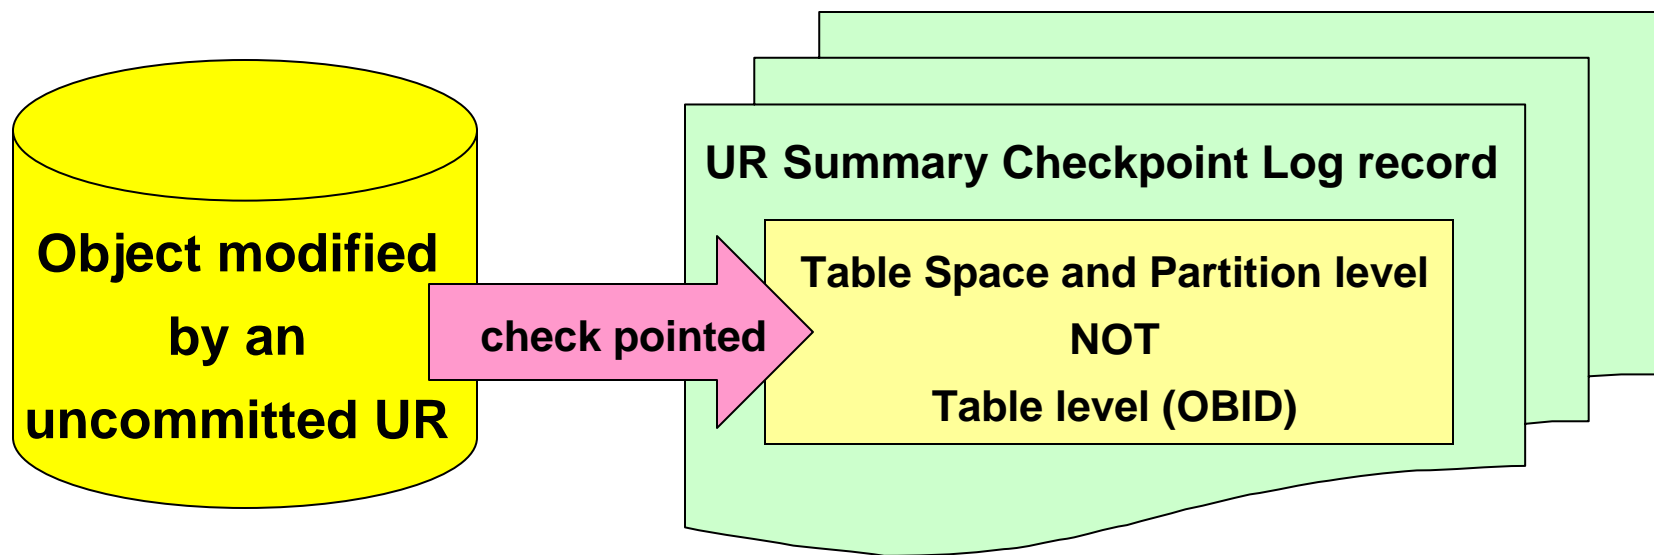
- Important exceptions
 - DB2 restarted with DEFER ALL option
 - DB2 restarted in System Level PITR mode
 - Tracker mode

Enhanced RETAIN Lock support

Allow table level retain lock support for the segmented table spaces, in data sharing environment, for Postponed Abort (PA) Unit of Recovery (UR)



Current V8 support



Information is used to:

- Determine how far back in the log to go in order to completely back out the UR.
- Build the resource name to maintain the retain lock for the postponed abort objects.

The DB2 9 enhancement

- Since this information is not kept at the table level, Data Manager could not maintain locking and the data change activity associated with postponed abort UR at the table level.
- We are adding table level information (OBID) into the UR summary checkpoint log record
 - If the UR goes postponed, end of restart processing will know which segmented table's retain lock to purge and which segmented table's retain lock not to purge.
- During the recovery of postponed abort UR, when the data changes for the table is completely backed out, the retain lock associated with that table will be purged.
 - Thus applications will have access to the table, even though the UR is still postponed.
- **Important** - since Advisory Restart Pending State (AREST) is at the table space level, the segmented table space will remain in AREST until the entire segmented table space is completely backed out.
 - Today, all the retain locks of the segmented tables are held until all the postponed abort URs are completely resolved.

LRSN Increment Wait Reduction

- In data sharing mode prior to DB2 9 for z/OS, the LRSN value of each log record on a given member was unique.
 - DB2 will re-drive pulling of STCK for LRSN, if two successive log records for the same member have the same value.

- The LRSN value increments every 16 microseconds.

- On modern processors, this can cause significant delays for workloads that have a high rate of log record creation.

- In DB2 9 for z/OS NFM and ENFM, the wait for LRSN increment is eliminated for most log writes.
 - In DB2 9, LRSN values only have to be unique within a given page
 - Compensating for any DB2 9 CPU performance regression.

-ACCESS DATABASE command

- Removes the GBP-dependent status for a table space, index space, or partition.
 - The MODE keyword specifies the desired state.
 - OPEN – Forces physical open of pageset
 - NGBPDEP – Converts pageset to non group buffer pool dependent

-ACCESS DB(dbname) SPACE(spacename) MODE(NGBPDEP)

- Useful when running work against a TS from one member only
- Drains all readers and writers on all other members
- Converts object to non GBP DEP, including castout of any changed pages
- Accesses from other members immediately sets page set to GBP DEP again

Enhanced “Unable to Obtain P-Lock” Messages

- DB2 modifies reason code 00C20255 (unable to obtain a P-Lock) and adds new codes to improve diagnosis of this event.
- RC00C20255 new description: “DB2 is unable to obtain a physical lock (P-lock) because another member holding the lock has not responded to the P-lock request.”
 - Old description: “DB2 is unable to obtain a physical lock (P-lock).”
- New RC00C2026A: “DB2 is unable to obtain a physical lock (P-lock) because another member holding the lock has been in DB2 shutdown for an extended amount of time.”
- New RC00C2026B: “DB2 is unable to obtain a physical lock (P-lock) because another member holding the lock has had logging suspended for an extended amount of time, due to either the active logs being full, or to an outstanding SET LOG SUSPEND command.”
- New RC00C026C: “DB2 is unable to obtain a physical lock (P-lock) because another member holding the lock encountered an abend in its P-lock exit.”
- This enhancement is only available in New Function Mode.

“Thank You for listening”

If you have any questions on this DB2 9 for z/OS session, then please send them to the BetaWorks team at:

Ian_Cook@uk.ibm.com
FLETCHPL@uk.ibm.com

