

No More Log RBA Resets!

DB2 11's Extended RBA and LRSN Support

Dave Levish

IBM DB2 for z/OS Development

August 8, 2014

Session 15769



#SHAREorg



Copyright (c) 2014 by SHARE Inc.  Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>

© 2014 IBM Corporation

Acknowledgements and Disclaimers

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. 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. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

All performance data presented may or may not have been collected in a controlled environment and may not be representative of the actual performance seen in a specific environment.

© **Copyright IBM Corporation 2014. All rights reserved.**

• **U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.**

IBM, the IBM logo, ibm.com, DB2® for z/OS®, IBM InfoSphere® Data Replication, and IBM InfoSphere Change Data Capture for z/OS® are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Remote Recovery Data Facility® (RRDF®) is a trademark of E-Net Corp.

Other company, product, or service names may be trademarks or service marks of others.

Agenda

- What's the situation?
- What's the solution?
- The Message

What's the problem? (round 1)

- **Running out of Log RBA...** 6 bytes x'FFFFFFFFFFFFFF'
 - (Continuously ascending **R**elative **B**yte **A**ddress for log records)
 - In DB2® Version 1 days (30 years ago), expected to occur in sometime around 2030!
- **Why?**
 - *Much* faster DASD
 - Replication (captures entire “before” row) and compression
 - Consolidation of DB2 data sharing members enabled by VSCR
 - Large, fast, higher n-way processors provide more throughput
 - heavy logging rates
- **What happens?**
 - **DB2 stops** until the log is reinitialized – ugly procedure


RBA Today

- DB2's Relative Byte Address (RBA) for logging is 6 bytes
 - 256TB of log record addressing per DB2 subsystem
 - 256TB of log record addressing per DB2 group member
- Some customers wrap the log in 2 months, others just approaching limit
 - Warning message DSNJ032I issued at x'F00000000000' (repeated every log switch)
 - Action needed message DSNJ033E at 'FFFF00000000' and DB2 shuts down

RBA Today

- Manual recovery actions are needed to reset RBA to zeros
 - Data Sharing:
 - Shut down affected member and start a new one in its place
 - Reset log RBA when member logs are no longer needed
 - Operational issues, such as automation, routing, and others
 - Must establish new recovery base
 - Non Data Sharing:
 - Reset PGLOGRBA field in every page to zero
 - Special form of COPY utility – extended outage (often multi-day)
 - http://pic.dhe.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.admin/src/tpc/db2z_subsystemendrba.htm
 - Procedure in b/g material

History Lesson – Why are we here?

- Don Haderle – the father of DB2 – thought 6 bytes was enough for 50 years
- Based on
 - 3350 DASD (1.2 MB/sec max)
 - High availability 
 - **12 hours / day**
 - **5 days / week**
- When DB2 V1 GA'ed in 1983, 3380 DASD (3.0 MB/sec) was the norm
- 2013 – FICON Express can write almost 1 GB/sec

LRSN Today – aka the problem, round 2

- The data sharing Log Record Sequence Number (LRSN) is based on the high order 6 bytes of the 8-byte time-of-day clock (plus “delta”)
 - hits end of range in **2042 w/o delta**
- No current circumvention for end of LRSN
- LRSN advances whether DB2 is running or not

- 6-byte LRSN value has precision to only 16 microseconds
 - Can cause LRSN uniqueness ‘spin’, burning extra CPU
 - Changes in DB2 9 and 10 address most LRSN spin situations
 - Spin still happens, mainly with row update/delete, and sometimes multi row insert

But wait, there's more: The LRSN Delta

- *Some* data sharing groups have a non-zero LRSN “delta” which gets added to the TOD clock
 - Causes the LRSN to be exhausted before 2042
 - LRSN “delta” value may be set when data sharing is enabled or re-enabled
 - Cause: data sharing enabled when the highest non shared RBA written is > current store clock upper 6 bytes
 - Print Log Map (DSNJU004) shows the LRSN delta value
 - STCK TO LRSN DELTA 000000000000
- Some non data sharing customers have enabled data sharing to circumvent RBA nearing end-of-range
 - This often causes a non-zero LRSN delta since high written RBA > LRSN

Solution

Solution

- Enlarge RBA to 10 bytes
- RBA addressing capacity of 1 yottabyte (2^{80})
- What's a yottabyte?
 - 1,208,925,819,614,629,174,706,176 bytes
- **DB2 for z/OS 11**

Larger RBA and LRSN

- **Expand the RBA and LRSN to 10 bytes**
 - RBA addressing capacity of 1 yottabyte (2^{80})
 - X'00000000112233445566' – add positions to high order RBA to left
 - LRSN extended on left by 1 byte, on the right by 3 bytes
 - x'00LLLLLLLLLLLLLLLL000000'
 - >30,000 years and 16M times more precision
- **Why 10 bytes??**
 - 8 bytes is not enough to solve LRSN issues and insufficient for both RBA and LRSN in the longer term
 - One byte on lower end of LRSN means DB2 could spin on current processors
 - 8-byte RBA -- could exhaust the range in < week with technology now in IBM labs

LRSN progression

- Numeric progression after end of 6-byte LRSN:

– 6 Byte end of range	00	FFFFFFFFFFFFFF	000000
– Next value	01	000000000000	000000
– The next value is	01	000000000000	000001
– Eventually...	01	000000000000	FFFFFF
– And next value is	01	0000000000001	000000
– ... 142 years later	02	000000000000	000000

Solution in DB2 V11 NFM

- In NFM, DB2 uses 6-byte values from CM until you take action to convert
- Two conversion tasks (can be done in either order):
 1. Convert BSDSes to new format to enable logging with larger RBA and LRSN values
 2. Convert DB2 objects to new page format via REORG in piecemeal fashion over days/weeks/months
- These tasks are optional in DB2 V11 (but recommended)
 - If you don't need larger RBAs/LRSNs now, no need to convert
 - Performance can be better if you convert BSDSes
 - Avoid internal conversion overhead on log write
 - Eliminate LRSN spin (Lab measurements are 0-6% improvement; heavy batch can be many percent more)

DB2 V11 – Keep this in mind...

DB2 11 is fundamentally a 10-byte system. 6-byte RBA and LRSN values are **always converted before processing.**

(Only log manager knows for sure)

Elegant solution designed for user flexibility / availability

- Many internal changes – affects the bones of DB2
 - Utilities - most have no external changes
 - Catalog and directory columns
 - Trace records: numerous IFCIDs are changed for extended RBA/LRSN
 - Most messages show 10 byte values **in all modes**
 - **New page format to handle 10-byte PGLOGRBA equivalent**
 - Utilities, SQL, DDL can handle either page format
 - Definition: “Basic” means 6-byte RBA/LRSN. “Extended” is 10-bytes
- ➔ Check with your vendor for specific tools you use: they must be able to handle/tolerate as appropriate:
- Basic and Extended RBA/LRSN (sometimes in CM!)
 - Both page formats

Modes of DB2 11 implementation

- **Any mode:** The Unit of Recovery ID (URID) extends to 10 bytes
 - -DISPLAY THREAD output, restart status table messages
 - Padded in CM with zeros
- **CM** – 6-byte RBA/LRSN and pages in current or “basic” format
 - New catalog objects for DB2 11 are created in CM in basic page format, but with extended RBA/LRSN columns, where applicable
 - RECOVER utility: TORBA, TOLOGPOINT, RESTOREBEFORE keywords extended to accept either 6- or 10-byte RBA/LRSN input
 - Standalone log utilities accept larger RBA/LRSN values
 - REPORT RECOVERY output accommodates extended RBA/LRSN values
 - Work files are in extended format

Modes of DB2 11 implementation

- **ENFM**

- DB2 10 RBA/LRSN related catalog columns are extended to 10 bytes
- You can convert any catalog objects extended format before NFM
 - They are being REORGed or created anyway
- Converting catalog objects during ENFM migration:
 - 2 ZPARMs can be set in CM but are ignored in that mode
 - *OBJECT_CREATE_FORMAT=EXTENDED*
 - » *Newly created catalog objects are created in extended format*
 - *UTILITY_OBJECT_CONVERSION=EXTENDED*
 - » *All REORGed catalog objects become extended format*
- Note: ENFM migration is shorter than DB2 10. Only 6 REORGs and 6 new table spaces
- ESP customer: “we can do production migrations during the week with business transactions running”

Modes of DB2 11 implementation

- **NFM**
 - Now you can convert BSDSes for extended RBA/LRSN
 - Begin to convert objects any time you REORG
 - Almost all partitioned objects can be converted independently (i.e. by part(s))
 - New job DSNTIJCV converts catalog/directory basic → extended page format
 - Optional, if you are not close to end of the log
 - Converts any not yet converted; over time all are in extended page format

Modes of DB2 11 implementation

- **NFM**

- SCA conversion done during NFM migration
 - Happens automatically; no manual action required
 - Recommended action: *IF* system managed duplexing is active for SCA, disable before NFM migration. **Note:** duplexing is not recommended for SCA due to possible performance impact.
 - Conversion rebuilds may happen during group restart after CM
- Conversion **MUST** have occurred before either the RBA or LRSN reach all x'FF'
 - Objects become “read only” since updates cannot be made in basic format

Procedure to Convert to Larger LRSN

- New warning messages for LRSN reaching end of range
 - DSNJ034I - WARNING: one year or less
 - DSNJ035E – Critical: 2 weeks of LRSN end (note “E” for “action”)
- Convert the BSDSes of all data sharing members to activate the extended log record format (DSNTIJCB job while DB2 is down)
 - Roll through one data sharing member at a time for continuous availability
 - Typically runs in under 5 seconds
 - Practice / test before actual conversion
 - Conversion is to a new copy, not in-place


Procedure to Convert to Larger LRSN

- Convert table spaces and indexes to the new page format
 - When LRSN advances past the 6-byte limit (x'FFFFFFFFFFFF') objects in the basic format become read only but still opened as RW
- Can be done in either order, but usually makes sense to convert BSDSes first
- Procedure for non sharing DB2 subsystems similar – in b/g material

Converting the Page Format


- An object can be converted to / from extended format by REORG or REBUILD utilities
 - ZPARM default or utility keyword
- A partitioned TS or index can be converted one part at a time
 - Except PBG hashed TS which must be converted at the TS level (all parts)
- CLONE tables cannot be converted
 - DROP CLONE
 - Convert base table, then ALTER ADD CLONE
 - Re add all the cloned authorizations
- SYSINDEXPART and SYSTABLEPART indicate the page format

Current 4K page format (Basic)



- Pghead contains PGLOGRBA (6 bytes stored) since DB2 V1
 - RBA/LRSN of last update to each page
- Sync bits are contained in PGTAIL (2 bytes)
- Last byte is PGEND (parity)
 - “E” - x’11000101’ or
 - “N” - x’11010101’
 - Underlined bit called PGBigTailApplies
 - Zeros mean off (does not apply)
- PGBigTailApplies = 0 → DB2 uses 6-byte RBA/LRSNs in pghead area

4K page with Extended RBA/LRSN



All page sizes have same format
 Maximum row length unchanged
 Fewer rows may be stored/page

- PGBigTail (20 bytes)
 1. Unused – 6 bytes
 2. Reserved – 2 bytes
 3. PGBigRBA – 10 bytes – Extended RBA/LRSN stored
 4. PGTAIL – 2 bytes
 - byte 1 reserved
 - Byte 2 – PGEND either
 x'01000010' or
 x'01010010'

Parity is bits 2-4 (100 or 101) and PGBigTailApplies is bit 7 -
 If one, it is ON, (Extended format PGBigRBA)
 If zero, it doesn't (Basic format)

New ZPARMs

- OBJECT_CREATE_FORMAT
 - Parameter states page format new objects use, basic or extended
 - Default is BASIC (6-byte RBA/LRSN format)
 - EXTENDED only takes effect in NFM (10-byte format)
 - – If you are going to convert, set it in CM and catalog objects touched in ENFM will be converted for you
- UTILITY_OBJECT_CONVERSION (ignored during CM)
 - Should page conversion be done by default by REORG, LOAD, REBUILD?
 - BASIC: the output object will be basic format (default)
 - EXTENDED: the output object will be extended format (only NFM)
 - NONE: the output object remains the same format
 - NOBASIC: the output object must be extended format (no exceptions!)
 - RBALRSN_CONVERSION utility parm overrides (not NOBASIC)
 - Specify the format of the object on completion of the utility
 - *Can't specify in CM*
 - – Specify EXTENDED for catalog objects with altered columns during ENFM to convert

UTILITY_OBJECT_CONVERSION=NOBASIC

- After objects are converted to extended mode, installation does not want them converted to basic by anyone
- Sets UTILITY_OBJECT_CONVERSION=NOBASIC
- Keeps clueless DBAs from shooting themselves in the foot!
- The only ZPARM value that is not overridden by utilities execution control statement RBALRSN_CONVERSION
 - The rest do
 - Valid only in NFM (setting ignored prior)
- After RBA limit reached, customer absolutely does not want user to revert!
- **Note:** RECOVER restores object to basic if the recovery base is basic

Main RBA/LRSN Sensitive Utilities

- RECOVER
 - REORG / REBUILD INDEX / LOAD REPLACE
- ➔ All utilities support both basic and extended page formats
- DSN1COPY – When source and target have differing page formats, copied data format and catalog will disagree
 - REPAIR – new option to correct catalog/OBD for format of actual data sets after DSN1COPY
 - LISTDEF – adds Filtering on Basic/Extended page format
 - Utilities display/messages are in extended format regardless of actual object format **in all modes**

DSN1COPY from DB2 11 to another subsystem

From DB2 11 Version	To Version	Basic Page format	Extended page format
NFM	DB2 11 non-CM	Yes	Yes
NFM	DB2 10	Yes	No
CM	DB2 10	Yes	No
CM *	DB2 10	Yes	No

- Reason for conversion back and forth is porting to another DB2
- Table shows the page format for a DB2 11 object that must be copied to a DB2 subsystem that is on another version of DB2
- A REORG is necessary if an object is already in extended page format
 - DB2 10 and earlier understand only 6-byte basic page formats
 - **Note:** DB2 9 end of service was June 27, 2014

RECOVER

- RECOVER accepts longer values for TORBA, TOLOGPOINT, RESTOREBEFORE
 - Six-byte RBAs may be specified as ten-byte operands with the six-byte value right justified in ten-bytes, i.e. x'00000000RRRRRRRRRRRRRR'
 - Six-byte LRSNs may be specified as ten-byte operands in byte locations two through seven, i.e. x'00LLLLLLLLLLLLLLLL000000'
 - Corrects any mismatch between recovered data and catalog
 - i.e. current = Extended, but TORBA = Basic

Utilities in General

- Use DSNZPARM UTILITY_OBJECT_CONVERSION globally
 - Unless overriding Utility statement parms (not for NOBASIC)
 - RBALRSN_CONVERSION
- Mapping table handles extended RBA/LRSN
 - Automatically created!
- Any SHRLEVEL can convert
- Index and underlying object unrelated
 - Index can be left in different page format from underlying objects
 - Example: 50 parts extended – 500 parts in basic – NPI can be either
- Summary of table on next slide

Utilities and Page Format Conversion

Utility	Partition	Index Partitions	Non Partitioned Index
Reorg Tablespace	Y	Y	Y
Reorg PART	Y	Y	Y
Reorg INDEX		Y	Y
Reorg INDEX PART		Y	N
Rebuild INDEX		Y	Y
Rebuild INDEX PART		Y	N
Load Replace Tablespace	Y	Y	Y
Load Replace PART	Y	Y	N

- Tablespace reorg required for only:
 - Hashed tables (hash algorithm to spread data)
 - XML Versioned T/S – to expand START_TS and END_TS from 8 > 10 bytes - Must occur before or with conversion to extended page format then Part Reorg OK

Recovery Considerations

Example for table space 'TS1':

T1 - image copy

T2 - quiesce

T3 - convert to extended page format

T4 - quiesce

T5 - image copy

T6 - quiesce

Q: Can I recover TS1 back to time T2?

A: Yes. The catalog will be updated to indicate "basic" format

Q: Can I then roll forward to time T4?

A: No. RECOVER does not allow log apply across a REORG

Considerations for Replication products

- Replication products must be updated for DB2 11
- IFCID 306 products will not work in DB2 11 CM
 - IFCID 306 always returns 10-byte log format in all modes!
 - IBM InfoSphere® Data Replication
 - IBM InfoSphere Change Data Capture for z/OS (aka CDC)
 - Any others
- Same for products that have user-written log capture exit
 - (CSECT DSNJW117 in load module DSNJL004)
 - Interface changed significantly, may prevent restart
- In DB2 11 CM any coexisting DB2 10 members see log records in old format
 - Workaround for IFCID 306 based products if customer does not yet have the updated version installed

End of RBA Timeline

Messages are not issued once BSDS is converted

F00000000000	FFF80000000	FFFF00000000	FFFFFFFFFFFF
DSNJ032I	Soft limit -904 00C2026D	DSNJ033E Critical	Hard limit -904 00C2026E
Every log switch and Restart	No SQL updates for basic objects	DB2 stops 00D10251	Convert BSDS and objects NOW!

Begin Converting
To Extended

DB2 in MAINT mode
Convert BSDS

End of LRSN Timeline

Messages are not issued once BSDS is converted

1 Year left	2 months left	2 weeks left	FFFFFFFFFFFF
DSNJ034I	Soft limit -904 00C2026D	DSNJ034I Critical DSNJ035E	Hard limit -904 00C2026E
Every log switch and Restart	No SQL updates for basic objects	DB2 stops 00D10251	Convert BSDS and objects NOW!

Begin Converting
To Extended

DB2 in MAINT mode
Convert BSDS

Messages and Limits

Messages	RBA (D/S or non-D/S)	LRSN D/S only	Reason Code
DSNJ032I	X'F00000000000'	n/a	n/a
DSNJ033E	X'FFFF00000000'	n/a	00D10251
DSNJ034I	n/a	1 year or less	n/a
DSNJ035E	n/a	2 weeks	00D10251
Advisory limit ('soft')	X'FFF800000000'	n/a	-904 00C2026D
Advisory limit DS ('soft')		2 months	-904 00C2026D

- Messages appear at DB2 restart and every log switch – not just a speed bump!
- Soft limit is like “Bridge Out” sign before road ends:
 - Rollback and abort allowed
 - SQL updates not allowed
 - DDL ok if DROP/CREATE for extended format tablespace –

utilities okay too

Notes

- DB2 will issue the messages at times other than log switch. One could construct a scenario on a query dominated system where the message may never be issued before hitting the hard limit if the active log data set was large enough and little logging occurred.
- When hard limit reached (end of RBA/LRSN), any objects that have not been converted to extended format are read only, because it is impossible to update correctly the PGLOGRBA. The object must be converted with REORG or REBUILD to enable read / write access.
- DSNJ034I WARNING is issued at every log switch and DB2 restart. It is started as closely relative to the prior events as possible (ENDLRSN – 1 year). DB2 deletes the message and repeats.
- DSNJ035E is issued along with DSNJ034I CRITICAL at 2 weeks. When DSNJ035E is issued, DB2 terminates unless it is running in MAINT mode or restart light (not desired but allowed). If you crash and hit brick wall you can start DB2 light..
- But by this time the soft limit will have been reached. IBM hopes no one ever sees DSNJ034I CRITICAL or DSNJ035E
- The DSNJ032 / 33 / 34 / 35 messages are not issued after the BSDS is converted. If database objects remain in the old format, updates will fail with -904 as the end of the 6-byte range nears. “Soft” limit” implemented to ensure that the installation has time to resolve any outstanding updates to basic format objects.

Benefits of extended RBA/LRSN

- Provides permanent relief for operational issues/outages caused by hitting the end of RBA or LRSN
 - Penalty in CM for basic format and 6-byte logs (but better than V10)
 - – Converting logs to 10-byte and tables to extended format may yield a few percent CPU improvement (your mileage may vary)
 - Eliminate LRSN spin
 - Some batch updates showed > 20% improvement
- Complete Flexibility
 - Conversion of BSDS and DB2 objects can be done in any order
 - Conversion of DB2 objects performed non-disruptively with REORG or LOAD
 - Can be done either way – Basic to Extended or Extended to Basic
 - No extra effort – can be done over time


The Message

- **DB2 11 is a 10-byte RBA/LRSN system**
- If you think you are being “conservative” or “safe” by remaining in 6-byte RBA/LRSN mode ...
- Think again
 - Every 6-byte RBA/LRSN that is touched is converted to 10-bytes before processing, and may need to be converted back to 6-bytes.
- Conversion is user friendly and flexible
 - “There is no “can’t do” in this function”
 - *ESP customer*

Resources

- **Technote:** The expanded RBA and LRSN in DB2 11 for z/OS
 - <http://www-01.ibm.com/support/docview.wss?uid=swg21665964&ce=ism0062&ct=swg&cmp=ibmsocial&cm=h&cr=im&ccy=us>
 - Or search on “IBM DB2 technote RBA”
- **DB2 11 manuals:**
 - What’s New:
 - Expanded RBA and LRSN log records
 - Utility RBA and LRSN handling
 - Administration Guide:
 - When RBA or LRSN limits are reached
- **Redbook:** DB2 11 for z/OS technical overview
 - The Extended RBA and LRSN

Thank You!




#SHAREorg



Copyright (c) 2014 by SHARE Inc.  Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>

Background slides



#SHAREorg



Copyright (c) 2014 by SHARE Inc.  Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>

Active Log Data Set Background

- Support for active log data sets up to 4GB was introduced in DB2 V6
 - The previous maximum was 2GB
 - Total active log space: 124GB (31 active logs * 4GB)
 - Up to 1000 archive logs, 4TB
- DB2 V8 NFM and later allows up to 93 active log data sets
 - Maximum active log space: 372GB
 - Up to 10000 archive logs, 40TB
- In DB2 11, active log formatter (DSNJLOGF) can preformat 10 log data sets per execution
 - One in DB2 10

Non Data Sharing Procedure to Convert to Large RBA

- Convert the BSDS to activate the extended log record format
 - New job DSNTIJCB that invokes the stand alone utility, DSNJCNVT
 - Typically runs in under 5 seconds
- Convert table spaces and indexes to the new page format
 - Once the RBA value has advanced past the 6-byte limit (x'FFFFFFFFFFFF') then any object still in the old format is forced to R/O access by DB2
 - No way to update the RBA now

BSDS Conversion

- Run new DSNJCNVT utility to convert BSDS after creating new ones.
 - Need new output BSDS because CI size is increased
 - Rename old BSDS with IDCAMS ALTER NEWNAME (SYSUT1/2)
 - Create new BSDS with larger CI size (SYSUT 3/4)
 - Typically takes a few seconds to execute
- SDSNSAMP job is DSNTIJCB

Log Limits and behavior

- Advisory logging occurs at RBA “FFF80000000’x or LRSN about 2 months before LRSN exhausted. Updates for rollback work
 - No SQL updates allowed
- Actual limit occurs when RBA/LRSN no longer fit in 6 bytes
 - LOAD REPLACE, REBUILD, REORG works for conversion into EXTENDED format only – ignores DSNZPARMs and utility settings
 - RECOVER of object in BASIC format sets REORP/RBLDP
 - RESTORE SYSTEM allowed for both limits for non-D/S. For D/S allowed only if SLB created after BSDS conversion or recover from tape..
- MERGECOPY – Basic format not allowed
- Objects created after hard limits reached are always extended format, regardless of DSNZPARM settings

DB2 10 Procedure to reset the RBA for non sharing DB2..1/3

To reset the log RBA value in a non-data sharing environment by using the COPY utility:

- Alter all of the indexes so that they have the COPY YES attribute by issuing the ALTER INDEX COPY YES statement. Make sure that you alter user-created DB2 catalog indexes and commit the changes after every thirty ALTER statements. You do not need to alter DB2 directory indexes, because by default, these indexes already have the COPY YES attribute.
- Tip: Optionally, you do not need to alter the indexes for the DSNDB06.SYSCOPY catalog table space, because later in this procedure you reset the log RBA value for all of these indexes.
- Issue the -STOP DB2 command to quiesce the subsystem that is approaching the end of the log RBA range.
- Restart DB2 in ACCESS(MAINT) mode.
- Issue the -DISPLAY THREAD command. Ensure that there are no INDOUBT or POSTPONED ABORT units of recovery.
- Issue the -DISPLAY UTILITY command. Ensure there are no active or stopped utilities.
- Issue the -DISPLAY DATABASE(*) SPACENAM(*) RESTRICT command. Ensure that all restricted states are removed.

DB2 10 Procedure to reset the RBA for non sharing DB2..2/3

- Quiesce the DB2 subsystem again by issuing the -STOP DB2 command.
- Use IDCAMS to delete and redefine the table spaces SYSUTILX, SYSCOPY, and SYSLGRNX and their corresponding indexes. Then, re-initialize these page sets. Complete the following steps to enable the COPY utility to reset the log RBA values in data pages and index pages as they are copied:
- Edit member DSN6SPRC of the prefix.SDSNMACS library and locate the entry SPRMRRBA.
- Change the SPRMRRBA setting to '1' and save the change Run the first two steps of your customized copy of job DSNTIJUZ to rebuild your DB2 subsystem parameter module (DSNZPxxx)
- Cold start this subsystem back to the RBA value of 0 (zero). This step removes all log data from the BSDS. This step requires utility DSNJU003 with the following options: CRESTART CREATE,STARTRBA=0,ENDRBA=0
- Start the DB2 subsystem in ACCESS(MAINT) mode.
- Rebuild any user-created indexes that are on the SYSIBM.SYSCOPY table by using the REBUILD INDEX utility.
- Take new, full image copies of all table spaces and indexes by using the COPY utility with the SHRLEVEL REFERENCE option to automatically reset the log RBA values.

Procedure ...3/3

- First, copy the catalog and directory table spaces and indexes, including any user-created DB2 catalog indexes.
- Copy the table spaces and indexes in user-created databases. You can do this step and the next step in parallel.
- Copy the tables spaces and indexes in default database DSNDB04. Note: Do not copy the table spaces in workfile database DSNDB07. DB2 automatically resets the log RBA values in these table spaces when they are used.
- Verify that the log RBA values were reset:
- Run a query against the tables SYSIBM.SYSCOPY, SYSIBM.SYSTABLEPART, and SYSIBM.SYSINDEXPART to verify that all objects were copied.
- Use the DSN1PRNT utility with the FORMAT option to print several pages from some of the objects so that you can verify that the PGLOGRBA field in the pages are reset to zero. The COPY utility updates the PGLOGRBA field and other RBA fields in header pages (page zero) so these fields will contain non-zero values.
- Stop DB2, and disable the reset RBA function in the COPY utility by following the instructions in step 9 and setting SPRMRRBA to '0'.
- Restart DB2 for normal access.
- Alter the DB2 catalog indexes and user-created indexes to have the COPY NO attribute by issuing the ALTER INDEX COPY NO statement. Commit the changes after every thirty ALTER statements. However, you should issue these ALTER statements over several days, because during this process SYSCOPY and SYSLGRNX records are deleted and contention might occur.