Answer: ‘It Depends’ What’s Your DB2 Question?

Session Number TDZ-1888

Christopher J. Crone
IBM – DB2 for z/OS Development
Jeff Lane
HCSC

Information On Demand 2010
The Premier Forum for Information & Analytics

Gain Insight. Optimize Results.

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, 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 ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

Agenda

- Objective
- Object Usage Options & Features
- Zparms
- The 3 R’s (Reorgs, Runstats, and Rebinds)
- Best Practices
- Summary

Objective

- DB2 has many options and DBAs often ask what is the best choice? Many times the answer is … It Depends!

- We will examine various topics and discuss pros and cons of choices surrounding the issue. Some answers could be no-brainer, direct, and straight-forward. On the other hand, some answers may be less obvious, thought provoking, or even controversial … It Depends!

- This presentation is intended to be Q&A type format to promote audience interaction, state facts, identify best practices, and expose 'gotchas' to help DBAs make better, informed decisions when faced with so many choices.

- What will you get out of this presentation? … It Depends!

Object Usage Options and Features

SMS or User Defined?

Answer – It Depends

- The Roger Miller Answer
- V10 – Catalog is SMS Managed managed
  - DSNTIJSS – part of V10 migration

What page size should I use for my tables?

Answer – It Depends

- Smaller page sizes tend to be better for Random I/O
- Larger page sizes tend to be better to Sequential processing
- Other things to consider
  - Locking
    - If you are using Page Level Locking, changing the page size can affect concurrency.
    - Compression increases the number of rows per page and can also affect concurrency
  - Spacemap pages
    - A spacemap can cover more rows if a larger page size is used. This may be useful for INSERT intensive workloads
  - Page size affects amount of data written to the CF.
    - If you update a row on a 4K page, 4K is written to the CF
    - If you update the same row on a 32K page, 32K is written to the CF

What page size should I use for my indexes?

Answer – It Depends

- Larger pages offer the potential to reduce the number of index leaf page splits, which are especially painful for GBP-dependent indexes in a data sharing environment
  - An Index split requires two synchronous I/Os
  - An Index page collapse requires five synchronous I/Os
  - Reduce index tree lotch and p-lock contention
- Potential to reduce the number of index levels
  - Reduce the number of getpages for index traversal
  - Reduce CPU resource consumption
- Large index page size may aggravate index bufferpool hit ratio for random access
- Larger Index pages (at least 8K) are required to be used with Index Compression

Large Index Page Size Examples

Rows In Table	1,000,000,000								
Key Length	**4**	**8**	**16**	**32**	**64**	**128**	**256**	**512**	**1024**
Page Size									
4096									
Entries/Leaf	336	252	168	100	56	29	15	7	3
Leafs	2,976,191	3,968,254	5,952,381	10,000,000	17,857,143	34,482,759	66,666,667	142,857,143	333,333,334
Non-Leaf fanout	331	242	158	93	51	26	13	7	3
Index Levels	**4**	**4**	**5**	**5**	**6**	**7**	**9**	**11**	**19**
8192									
Entries/Leaf	677	508	338	203	112	59	30	15	7
Leafs	1,477,105	1,968,504	2,958,580	4,926,109	8,928,572	16,949,153	33,333,334	66,666,667	142,857,143
Non-Leaf fanout	666	488	318	187	103	54	27	14	7
Index Levels	**4**	**4**	**4**	**4**	**5**	**6**	**7**	**8**	**11**
16,384									
Entries/Leaf	1360	1020	680	408	226	120	61	31	15
Leafs	735,295	980,393	1,470,589	2,450,981	4,424,779	8,333,334	16,393,443	32,258,065	66,666,667
Non-Leaf fanout	1,336	980	639	376	207	108	55	28	14
Index Levels	**3**	**4**	**4**	**4**	**4**	**5**	**6**	**7**	**8**
32,768									
Entries/Leaf	2725	2044	1362	817	454	240	123	62	31
Leafs	366,973	489,237	734,215	1,223,991	2,202,644	4,166,667	8,130,082	16,129,033	32,258,065
Non-Leaf fanout	2,676	1,963	1,280	755	414	218	111	56	28
Index Levels	**3**	**3**	**3**	**4**	**4**	**4**	**5**	**6**	**7**

Should I use Tablespace Compression?

Answer – It Depends

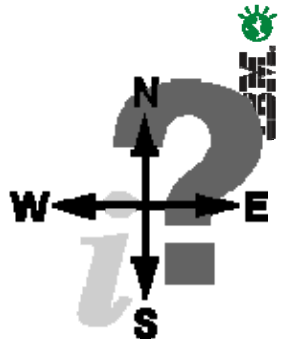

- What size is the average row?
- How many rows/page on average?
- What does DSN1COMP say the compression ratio would be?
- What size (kb/mb/gb/tb) is your table (tablespace)?
- What is the table used for?
  - Read Intensive?
  - Write Intensive?
- How is the table populated?
  - V10 helps here by creating the compression dictionary on the fly
- How long is the data kept?

Should I use Index Compression?

Answer – It Depends

- Is the Index used in a Warehouse or OLTP environment?
- Indexes are compressed on Disk
- Indexes are not compressed in the bufferpool
- Compression is on a page level
- Compressed indexes are stored on a 4K page
  - Extra CPU required on INSERT/UPDATE to ensure a key will fit on the page when compressed
    - This CPU is charged to the transaction
  - Pages are Compressed when written to disk. Pages are Uncompressed when loaded into bufferpool
    - The CPU is charged to the address space SRB time if the READ/WRITE was asynchronous
    - The CPU is charged to the transaction if the READ/WRITE is synchronous

http://www.redbooks.ibm.com/redpapers/pdfs/redp4345.pdf

Audience Questions?

ZPARMS

What value should I use for CHKFREQ?

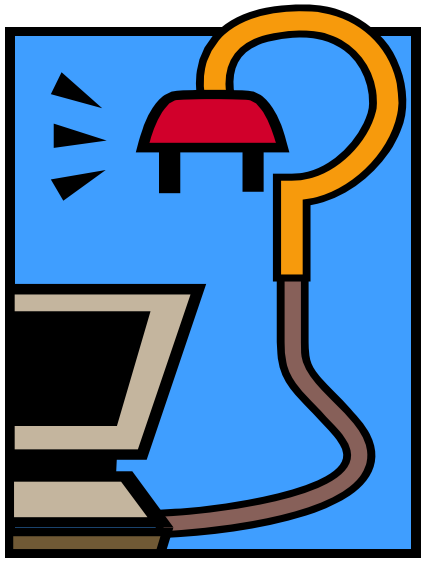

Answer – It Depends

- In V8 and V9, a customer has to choose between Checkpointing based on the number of log records or the time.
- Checkpointing is the process where DB2 flushes written pages from the bufferpools
  - Setting CHKFREQ to 200,000 would ensure a checkpoint occurs every 200,000 log records.
  - Setting CHKFREQ to 5 would ensure a checkpoint would occur every 5 minutes
- Setting a checkpoint either based on minutes or log records is a compromise
- DB2 V10 adds CHKTYPE, CHKLOGR, and CHKMINS
  - You can now specify both max duration and a max number of log records/checkpoints

What value should I set for MINSTOR/CONTSTOR/EDMBFIT

Answer – It Depends

- MINSTOR/CONTSTOR/EDMBFIT all are intended to provide similar functions
  - Trade CPU for storage
  - Intended to help minimize storage footprint and ensure system stability on systems that are close to the 2G DBM1 Private Storage limit
  - There are other options that can also affect the storage footprint of systems and the effectiveness of MINSTOR and CONTSTOR
    - Bind option RELEASE(COMMIT) or RELEASE(DEALLOC)

MINSTOR/CONTSTOR

- MINSTOR
  - Changes Agent Local (AGL) pool from First Fit to Best Fit. Only effective if you have fragmentation issues
  - No easy way to determine if you have fragmentation (a dump needs to be analyzed)
  - Online changeable, but you can really only measure effectiveness after an IPL
- CONTSTOR
  - Trades CPU for storage
  - Does not compress AGL, stack storage, or local statement cache
  - Ineffective for RELEASE(DEALLOC) threads
  - Fine tuning of this zparm can be done via SPRMCTH (number of commits) and SPRMSTH (threshold before contraction is applied)

EDMBFIT – It Depends – but not very often

- EDMBFIT is intended to switch the algorithm for allocating an object in the EDMPOOL from First Fit to Best Fit.
  - EDMBFIT = NO is the Recommended Best Practice
    - There is no need for this ZPARM for this with larger EDM Pools
    - This ZPARM is removed from DB2 V10
  - EDMBFIT = YES should only be used for very very small EDM Pools
  - Starting with V8 and Continuing with V10 the EMPOOL has been moving ATB – with each release more and more storage has been moved above and thus there is very little reason to constrict the EDMPOOL(s) to the point that EDMBFIT makes any sense

How should I set MAXDBAT and CTHREAD

Answer – It Depends

- Availability is a major reason to code these defensively
- Is the system part of a Data Sharing group?
  - How many members?
  - Do you have any workload affinity?
- What does your workload look like?
  - Do you have peaks and valleys?
- Is your thread footprint consistent?
  - Large Batch thread footprint?
  - RELEASE COMMIT or DEALLOCATE?
- Are you using Inactive DBATs?
- With JAVA (JDBC) connections: be aware of autocommit, and cursor with hold options. Make sure to close your database connection.
- V10 offers relief in this area

Should I use ASCII, EBCDIC, or Unicode?

Answer – It Depends

- Do you need to encode multiple languages?
  - Are the languages compatible?
    - English and Spanish = yes
    - English and Chinese = no
- What Programming language are you using?
  - COBOL and PL/I support UTF-16
  - Java is UTF-16 natively, but supports UTF-8 without penalty
- Is your application local or remote?
  - What is the encoding of your application?
- Does your application rely on messaging?
  - IMS Connect, MQ, CTG, other????
- What display technology are you using?
  - Green Screen, Thin Client, Thick Client?

The 3 R’s – Reorgs, Runstats, Rebinds

When should I run REORG?

Answer – It Depends

- Do I have to Reorg?
  - Tables?
  - Indexes?
- When do I have to Reorg? (no choice)
- What is the best practice to Reorg?
- How often should we Reorg?

When should I run RUNSTATS?

Answer – It Depends

- When do I have to run Runstats? (no choice)
- What is the best practice to run Runstats?
  - How often should I run Runstats?
  - What to Runstat and where – tbls, indx, lob, catalog…?
  - What Runstats parms to use?
  - How does Runstats affect different processing/workloads?

When should I rebind?

Answer – It Depends

- When you have to Rebind? (no choice)
  - When a package is invalidated
- When is it a best practice to Rebind?
  - After Runstats
  - At least once per release
- Do I need to rebind packages?
  - Even in V10?
- What about Plan Stability – Should I use it?
- If access paths changed, is it good or bad?

Audience Questions?

BEST PRACTICES

Should I use Release(Commit) or Release(Deallocate)?

Answer – It Depends

- Increases EDM pool storage
- Increase in thread footprint
  - Possible that many seldomly used statements are cached
- DBD concurrency may be impacted
- Can lower CPU usage (trading storage for CPU)
  - V10 makes this decision easier

Should we add another index to our table?

Answer – It Depends

- How many indexes are on the table now?
- What is the (SQL) activity on the table – is the cost (cpu, dasd, …) to maintain the new index worth the benefit it provides?
- Analysis may be required to determine what (SQL) is going to use the index vs what (SQL) should use the index?
- Other considerations
  - Bufferpool impact?
  - Datasharing? Potentially increased CF overhead
  - Logging
  - Utility impacts
    - REORG, RUNSTATS, RECOVER will be affected
    - COPY may also be affected

Should we use Row Level Locking or Page Locking?

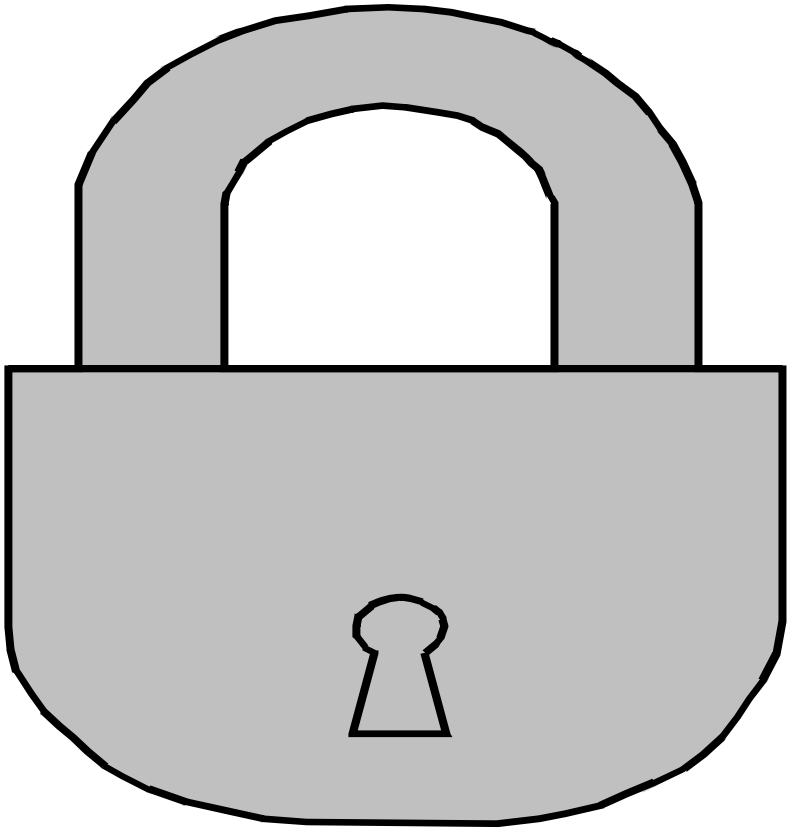

Answer – It Depends

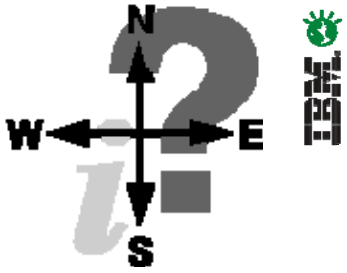

- Row level locking may increase the number of locks and cpu.
- Row level locking can improve concurrency.
- Data Sharing overhead with row level locking - worth the cost?
- Do we have to rebind if we switch from page to row level locking?
- Stick with page locking until it's too painful

How should we configure DSNDB07 (sortwork) tablespaces?

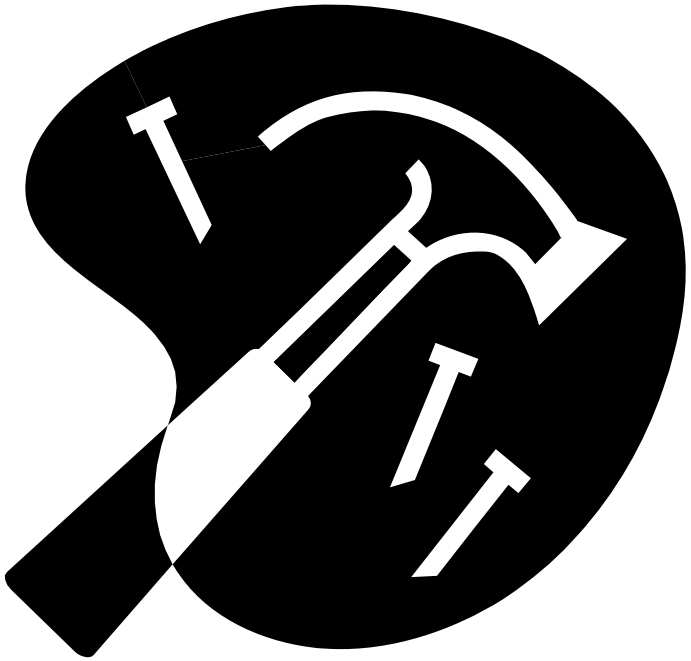

Answer – It Depends

- What version are you on V8, V9, V10?
- Do you have DGTTs?
- Many small tablespaces vs few large tablespaces?
- 4K or 32K page sizes?
  - V9, V10 – if record length > 100, preference is for 32K
  - Monitor using IFCID 2 (stats)
- How have you set extents (zero or non-zero SECQTY)?
- What's your Segsize(V9, V10)?
- What's your WFDBSEP setting(V9, V10)?
  - Yes – Forces hard separation based on SECQTY
  - No – Tries to separate based on SECQTY

Do I care how my DB2 Logs are configured?

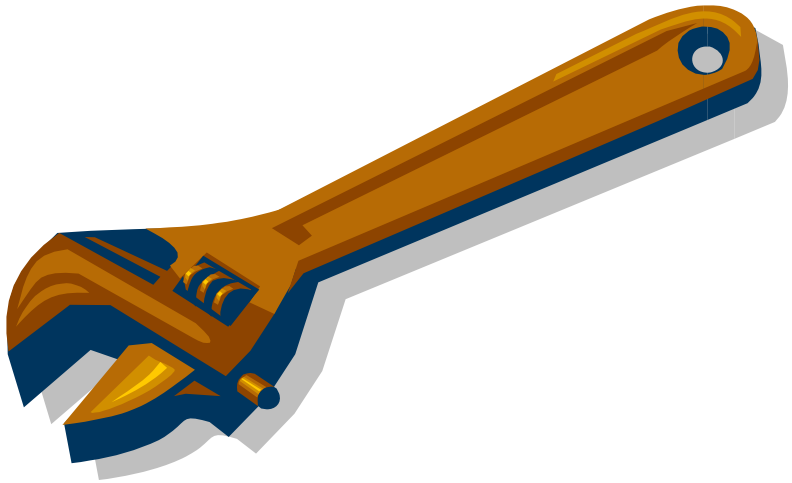

Answer – It Depends – what's your Job?

- DB2 System Programmers?
  - Backup & Recovery Strategy
    - System backup & recovery requirements – local, DR, …
  - Log Configuration
    - Number, size, placement, striping, …
    - Archive Strategy
      - VTS (Tape), DASD
      - Retention
  - Performance
- Application DBA?
  - Object backup/recovery requirements – local, DR, …
  - Performance
- Application Developers?
  - Long running URs
- Storage Admins?
  - Storage capacity
  - Performance

What is an optimal DB2 Bufferpool Configuration?

Answer – It Depends

- What are you trying to optimize?
  - Performance?
  - Servicability?
  - Scalability?
  - People Cost?
  - ????
- Some things to think about
  - One bufferpool vs many?
  - Tablespace and index bufferpools?
  - Random vs sequential bufferpools?
  - Required bufferpools?
    - Catalog bufferpools

Do we really need SDSNEXIT concatenated before SDSNLOAD?

Answer – It Depends

- Recommended in install guide
  - Issues tend to revolve around DSNHDECP
- Do you run one system or multiple systems?
- DB2 V8 requirement? Different in V9?
- Future SDSNEXIT usage?
  - DB2 V10 allows specification of DSNHDECP name (similar to DSNZPARM name previously allowed).

What's the best practice for INSERT workloads?

Answer – It Depends

- Table type?
  - Segmented
  - Classic Partitioned table space (Member Cluster)
  - UTS (PBG/PBR) (Member Cluster in V10)?
- Random or Sequential INSERT?
- Percent Free/Freepage?
- Data Sharing?
- Locking – Row or Page Level locking?
- Append?
- Compress(yes/no)?
- Page Size/Index Page Size?
- MAXROWS n?

See Sessions 3443 and 3444 – Optimizing Insert Performance

Summary

- We covered a variety of topics from object usage to best practices. DB2 provides many options and in many situations the best choice isn't always clear.

- We wanted to give you some additional insight to these topics to help you choose the best answer for your situation.

- Did we answer your DB2 questions? The answer very well could be … It Depends!

Thank You!
Your Feedback is Important to Us

- Please complete the survey for this session by:
  - Accessing the SmartSite on your smart phone or computer at: iodsmartsite.com
    - Surveys / My Session Evaluations
  - Visiting any onsite event kiosk
    - Surveys / My Session Evaluations
  - Each completed survey increases your chance to win an Apple iPod Touch with daily drawling sponsored by Alliance Tech

Session: 1888
Answer: ‘It Depends’ … What’s your DB2 question?

Chris Crone

IBM
cjc@us.ibm.com

Jeff Lane

HCSC
Jeffrey_Lane@bcbsil.com