



IBM Software Group

## Support de l'adressage 64 bits ...

[Catherine\\_chochoy@fr.ibm.com](mailto:Catherine_chochoy@fr.ibm.com)

**DB2** Information Management Software

© 2005 IBM Corporation



IBM Software Group

## Storage management before and after DB2 for z/OS V8

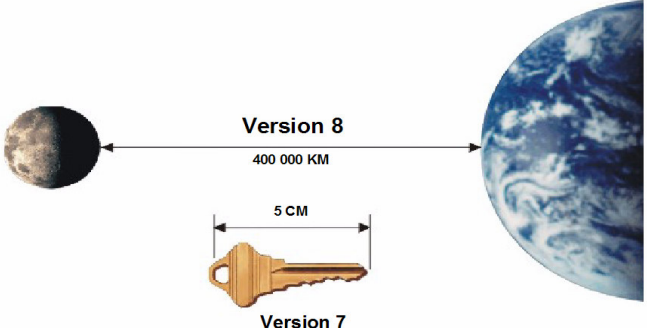
sources: Florence Dubois, redbook sg246079

**DB2** Information Management Software

© 2005 IBM Corporation

IBM Software Group | DB2 Information Management Software IBM

## How much bigger is 8 billion times?



© 2004 IBM Corporation

© 2005 IBM Corporation

IBM Software Group | DB2 Information Management Software IBM

## Before DB2 V8 – What is the problem?

- Background
  - Each address space has an addressing range of 2GB based on 31 bit addressing
    - Maximum of 16MB available "below the line"
    - Maximum of 2032MB available "above the line"
  - Practical maximum available to DB2 below and above the line much less
- GETMAIN processing by DB2
  - Requests may be conditional or unconditional
  - "Short on Storage" condition can occur for both
  - DB2 recovery routines may be able to clean up
  - Individual DB2 threads may abend with 04E/RC=00E200xx when
    - Insufficient storage available e.g., 00E20003 & 00E20016
  - Eventually DB2 subsystem may abend with abend S878 or S80A when critical task and no toleration of error

© 2005 IBM Corporation

## What are the drivers?

- Workload growth
- Subsystems/applications accessing DB2 recycled less frequently
- Growth in thread related storage over successive DB2 releases
- Slow downs and CTHREAD / MAXDBAT throttle set too high
- New workload types: ERP & CRM e.g., use of local DSC
- Long running persistent threads
- Dynamic SQL
- Widespread use of Compression plus wide Partitioning
- Size of ECSA to support WebSphere, IRLM PC=NO, etc
- Over allocation of ECSA and other extended common areas
- Query parallelism with high degree
- Heavy concurrent SQL prepare activity without DSC turned on
- Over allocation of bufferpools, EDM Pool, etc

## VSTOR allocation for a typical DB2 V7 installation

Usage	Typical Value
Virtual Buffer Pools	40 to 800MB
EDM Pool	20 to 400MB
User Thread Storage	0.1 to 2MB per active user thread
System Thread Storage	40KB per system thread
Compression Dictionary	64KB per open compressed data set
Local Dynamic Statement Cache	0 to 300MB
Other Areas	200 to 400MB

Recommended White Paper - John Campbell / Mary Petras  
 "DB2 UDB Server for OS/390 Virtual Storage Management"

## What is allocated in DBM1?

- Storage allocated for both subsystem and threads
  - Majority of storage allocated « above the line »
- |   |  |
|---|--|
| <ul style="list-style-type: none"> <li>➤ DB2 Code</li> <li>➤ Virtual (Buffer) Pools</li> <li>➤ Controls blocks for VP and HP buffers</li> <li>➤ Lookaside buffers for data space BP</li> <li>➤ EDM Pool including Global Dynamic Statement Cache</li> <li>➤ RID Pool</li> <li>➤ Compression dictionaries for open datasets</li> </ul> | <ul style="list-style-type: none"> <li>➤ <u>RDS OP Pool</u></li> <li>➤ <u>Agent Local Storage</u></li> <li>➤ <u>Stack Storage</u></li> <li>➤ <u>Internal Trace Pool used by Buffer Mgr. and Data Mgr.</u></li> <li>➤ <u>Local Dynamic Statement Cache</u></li> <li>➤ Storage required by various service tasks</li> <li>➤ Fast log apply buffers</li> <li>➤ Real Time Statistics blocks</li> </ul> |
|---|--|

## How much VSTOR is available to DBM1?

- RMF Virtual Storage Private Area Report
  - Interval data collected in SMF Type 78-2
  - RMF Monitor I session option: *VSTOR(D,xxxxDBM1)*
  - RMF Post Processor option: *REPORTS(VSTOR(D,xxxxDBM1))*
- Calculate amount of storage available above the line  
 REGION ASSIGNED – (MAX LSQA/SWA/229/230 PAGES ALLOCATED  
 + MAX USER REGION PAGES ALLOCATED)
- How much is enough?
  - Greater than 500MB spare is AOK (**GREEN**)
  - Between 200-500MB spare is boundary condition (**AMBER**)
  - Less than 200MB action is required (**RED**)

IBM Software Group | DB2 Information Management Software

Adobe Acrobat - [DB2 for z/OS Virtual Storage Management - Before and After zSeries.pdf]

VIRTUAL STORAGE ACTIVITY PAGE 3  
 OS/390 SYSTEM ID DED1 DATE 07/01/1999 INTERVAL 14.59.596  
 REL. 02.06.00 RPT VERSION 2.6.0 TIME 10.15.00 CYCLE 1.000 SECONDS

JOB NAME - DBCDML PRIVATE AREA SUMMARY REGION REQUESTED OK  
 STEP NAME - DBCDML REGION ASSIGNED (BELOW 16M) 12.0M  
 PROGRAM NAME - SDBTARCP REGION ASSIGNED (ABOVE 16M) 1862M  
 NUMBER OF SAMPLES - 00

PRIVATE STORAGE MAP

BELOW 16M				EXTENDED (ABOVE 16M)			
ADDRESS	LOGA/SMA	SIZE	STATUS	ADDRESS	LOGA/SMA	SIZE	STATUS
BFFFFFF	229/230	1280K	BOTTOM OF ALLOCATED AREA	7FFFFFFF	229/230	1329M	2CF88000
AC0000	10.15.01	10.6M	UNUSED		10.15.01	512M	7FFFFFFF
C00000	10.15.01	10.6M	UNUSED		10.25.31	512M	CF880000
240000	10.15.01	10.6M	UNUSED		10.25.31	512M	CF880000
5000	10.15.01	132K	UNUSED		10.25.31	512M	CF880000
1000	10.15.01	16K	UNUSED		10.25.31	512M	CF880000

-----  
 MIN MAX AVG MIN MAX AVG  
 LOGA/SMA/229/230 0K 10.15.01 0K 0K \*\*\*\*\* 10.23.32 1564M 10.20.31 1768M  
 FREE PAGES (BYTES) 0K 10.15.01 0K 0K \*\*\*\*\* 10.23.32 1564M 10.20.31 1768M  
 LARGEST FREE BLOCK 0K 10.15.01 0K 0K \*\*\*\*\* 10.23.32 1564M 10.20.31 1768M  
 PAGES ALLOCATED 936K 10.15.01 936K 10.15.01 936K 1251M 10.21.22 1273M 10.16.22 1263M  
 USER REGION 936K 10.15.01 936K 10.15.01 936K 1251M 10.21.22 1273M 10.16.22 1263M  
 FREE PAGES (BYTES) 10.9M 10.15.01 10.9M 10.15.01 10.9M \*\*\*\*\* 10.17.51 2024M 10.23.52 \*\*\*\*\*  
 LARGEST FREE BLOCK 10.6M 10.15.01 10.6M 10.15.01 10.6M 1840M 10.25.31 1840M 10.15.01 1840M  
 1M GETMAIN LIMIT 10.6M 10.15.01 10.6M 10.15.01 10.6M 1840M 10.25.31 1840M 10.15.01 1840M  
 PAGES ALLOCATED 132K 10.15.01 132K 10.15.01 132K 21.7M 10.15.01 21.7M 10.25.31 21.7M  
 (IN BYTES)

1862-(1273+21.7)=567.3MB available

9 © 2005 IBM Corporation

IBM Software Group | DB2 Information Management Software

## How much VSTOR is used in DBM1?

- DB2PM Statistics Report|Trace Layout Long
  - ZPARM SMFSTAT=(...,6) to generate IFCID 225
  - ZPARM STATIME=5 (mins)
  - ZPARM SYNCVAL=0
- Develop and set virtual storage budget
  - Determine how much non-thread related storage is required
  - Develop how much storage is used per active thread
  - Plan on keeping at least 200 MB spare for tuning, growth, recovery, etc.
  - Determine how many active threads can be supported
- Set CTHREAD and MAXDBAT defensively for robustness to protect system

10 © 2005 IBM Corporation

## 64-bit REAL storage support

- Prerequisites: zSeries, OS/390 V2 R10 or z/OS, ESAME 64-bit LPAR, DB2 V6+
- Data spaces provided a good short term solution by exploiting 64-bit REAL
  - Buffer pools and statement caching in data spaces
    - Frees up space for other work in the DBM1 address space
    - Performance penalty when not 100% backed by real storage
- Advantages of data spaces over hiperpools
  - Read and write cache with direct I/O to data space
  - Byte addressability
  - Very large buffer pool sizes
    - 32GB for 4K page size
    - 256GB for 32K page size
- Excellent performance experienced with z900 and large processor storage
  - Performance dependent upon being in 64-bit REAL mode

## DB2 V8 – 64-bit VIRTUAL storage support

- DB2 Version 8 is 64-bit exclusive
  - Buffer pools always allocated above 2GB bar
  - Eliminates need for hiperpools and data space pools
  - "Data access" modules enhanced to access 64-bit addressable buffers "in place"
  - NO internal data movement as per data space buffer pools today
- Sizing and placement
  - Buffer pool max size is 1TB
    - The actual maximum = the REAL storage available
  - Total buffer pool max size is 1TB
  - Page manipulation blocks (PMBs) moved above the bar
  - Castout buffers (data sharing) above the bar

## Limit changes


- Max number of read, write, castout engines increased to 600
- z/OS V1.2 provides a new MEMLIMIT JCL keyword
  - Controls how much VSTOR above 2GB bar is available in address space
  - DB2 sets MEMLIMIT value to 4TB (minimum) ensuring sufficient memory for operation
- VPSIZE default values:

BP0	20000	(Minimum 2000 (from 56))
BP8K0	1000	
BP16K0	500	
BP32K	250	


## What else is moving above the 2-GB bar?

- RID pool
- Sort pool
- Compression dictionaries
- EDM pool - DBDs and OBDs

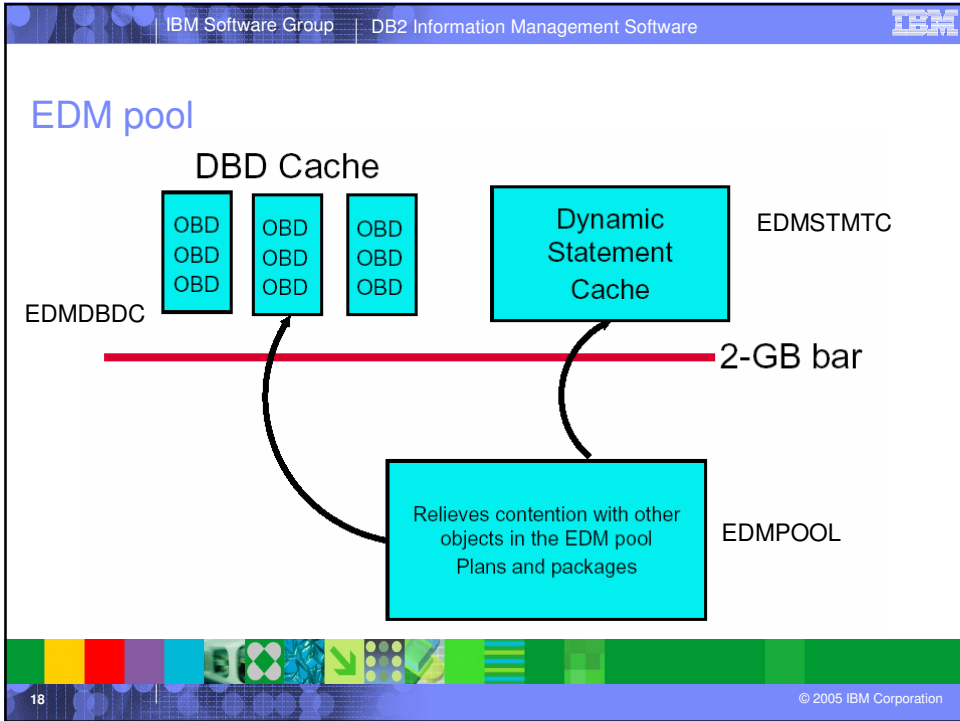
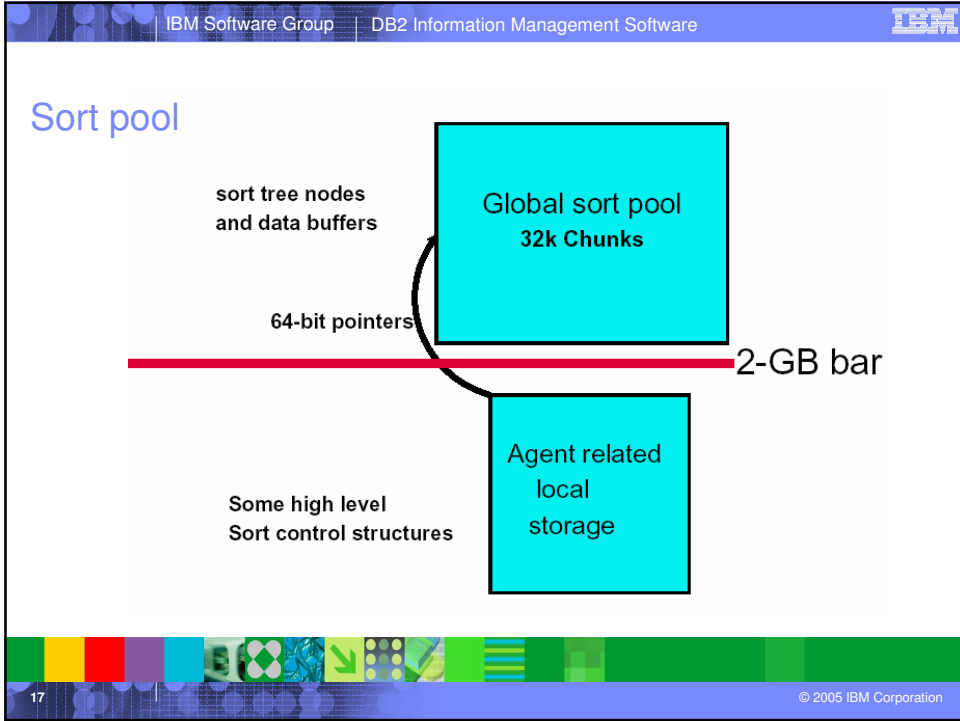
## RID pool



## Compression dictionaries







## How much is 64-bit virtual support worth for you? ...

- At first sight, re-engineering of DBM1 to exploit 64-bit should make a very significant difference in terms of providing significant VSCR
  - More concurrent threads
  - Higher transaction throughput
- Now consider V6 or V7 installation achieving significant VSCR
  - Maximum use of Dataspace bufferpools (400-800MB)
  - Using dataspace extension to EDM Pool for GDSC (80-160MB)
- What is the 'net' benefit in V8
  - Dataspace Lookaside Pool (80-120MB) is eliminated
  - Buffer Manager control blocks (70-120MB) going above 2GB bar
  - Other items going above the 2GB bar (compression dictionaries, certain EDM/RID pool components, thread sort pool, etc.)

## How much is 64-bit virtual support worth for you? ...

- Most of the thread storage stays below the 2GB bar
  - Agent Local, Stack Storage
  - Local Dynamic Statement Cache (0-400MB)
  - Expect some regression
- How many additional threads can be supported will depend on 'Thread Footprint' which will vary by workload depending on
  - Duration of thread
  - SQL workload
  - RELEASE parameter setting on plan/package bind
  - Effectiveness of thread storage contraction (CONTSTOR=YES)
- Net benefit might only work out at 300-500MB for some installations
- Still need to monitor and track VSTOR usage with IFCID 225

## DB2 PE Statistics Report V8 - Extract

DBM1 AND MVS STORAGE BELOW 2 GB		QUANTITY
TOTAL DBM1 STORAGE BELOW 2 GB	(MB)	81.70
TOTAL GETMAINED STORAGE	(MB)	34.06
VIRTUAL BUFFER POOLS	(MB)	N/A
VIRTUAL POOL CONTROL BLOCKS	(MB)	N/A
EDM POOL	(MB)	32.12
COMPRESSION DICTIONARY	(MB)	N/A
CASTOUT BUFFERS	(MB)	N/A
DATA SPACE LOOKASIDE BUFFER	(MB)	N/A
HIPERPOOL CONTROL BLOCKS	(MB)	N/A
DATA SPACE BP CONTROL BLOCKS	(MB)	N/A
TOTAL VARIABLE STORAGE	(MB)	33.74
TOTAL AGENT LOCAL STORAGE	(MB)	29.23
TOTAL AGENT SYSTEM STORAGE	(MB)	20.36
NUMBER OF PREFETCH ENGINES		4.00
NUMBER OF DEFERRED WRITE ENGINES		207.00
NUMBER OF CASTOUT ENGINES		0.00
NUMBER OF GBP WRITE ENGINES		0.00
NUMBER OF P-LOCK/NOTIFY EXIT ENGINES		0.00
TOTAL AGENT NON-SYSTEM STORAGE	(MB)	8.86
TOTAL NUMBER OF ACTIVE USER THREADS		14.35
RDS OP POOL	(MB)	0.00
RID POOL	(MB)	0.00
PIPE MANAGER SUB POOL	(MB)	0.00
LOCAL DYNAMIC STMT CACHE CNTL BLKS	(MB)	0.99

## DB2 PE Statistics Report V8 – Extract ...

...		
THREAD COPIES OF CACHED SQL STATEMENTS		0.00
BUFFER & DATA MANAGER TRACE TBL	(MB)	0.00
TOTAL FIXED STORAGE	(MB)	0.21
TOTAL GETMAINED STACK STORAGE	(MB)	13.69
STORAGE CUSHION	(MB)	18.16
24 BIT LOW PRIVATE	(MB)	0.14
24 BIT HIGH PRIVATE	(MB)	0.48
31 BIT EXTENDED LOW PRIVATE	(MB)	33.57
31 BIT EXTENDED HIGH PRIVATE	(MB)	100.61
EXTENDED REGION SIZE (MAX)	(MB)	1610.00
EXTENDED CSA SIZE	(MB)	280.86
DBM1 STORAGE ABOVE 2 GB		QUANTITY
FIXED STORAGE	(MB)	2.54
GETMAINED STORAGE	(MB)	270.57
COMPRESSION DICTIONARY	(MB)	0.00
CACHED DYNAMIC SQL STATEMENTS (MAX)	(MB)	100.38
DBD CACHE (MAX)	(MB)	100.38
VARIABLE STORAGE	(MB)	67.29
VIRTUAL BUFFER POOLS	(MB)	0.00
VIRTUAL POOL CONTROL BLOCKS	(MB)	0.00
CASTOUT BUFFERS	(MB)	0.00

## Real Storage Use

- Important subsystems such as DB2 should not be paging to auxiliary storage
- RMF Monitor II Address Space Data (ASD) Report
  - Real storage usage by DB2 DBM1 address space
- RMF Monitor III Option 3 Storage Frames (STORF) resource
  - Number of REAL+AUX frames used by DB2 DBM1 address space
- UIC value highlighted in the ASD report is a good indicator if you are running out of real frames
  - If this value starts to dive it is an indication that you are starting to run out of real storage
    - > 200 (GREEN)
    - 100 - 200 (AMBER)
    - < 100 (RED)

## RMF Mon II ASD Report

RMF - ASD Address Space State Data													Line 1 of 44							
													CPU=	2 / 2	UIC=2540	PR=	0	System= NB01		Total
JOBNAME	SRVCLASS	P	L	LS	PR	CS	ESF	CS	TAR	X	PIN	ES	TX	SWAP	WSM					
						F			TAR	WSS	M	RT	SC	RV	RV					
*MASTER*	SYSTEM	1	NS	FF	4649					0	----			0	0					
PCAUTH	VEL30STC	1	NS	F7	101					0	X	----		0	0					
RASP	SYSTEM	1	NS	FF	155					0	X	----		0	0					
TRACE	VEL30STC	1	NS	F7	67					0	X	----		0	0					
DUMPSRV	SYSTEM	1	NS	FF	73					0	----			0	0					
XCFAS	SYSTEM	1	NS	FF	1220					0	X	----		0	0					
GRS	SYSTEM	1	NS	FF	1440					0	X	----		0	0					
SMSPDSE	SYSTEM	1	NS	FF	1776				432	X	----			0	0					
CONSOLE	SYSTEM	1	NS	FF	381					0	X	----		0	0					
WLM	SYSTEM	1	NS	FF	898					0	X	----		0	0					
ANTMAIN	SYSTEM	1	NS	FF	2009					0	X	----	1	998						
ANTAS000	VEL30STC	1	NS	F7	2293					0	X	----	1	998						
OMVS	SYSTEM	1	NS	FF	23.2K				20.5K	X	----			0	0					
IEFSCHAS	SYSTEM	1	NS	FF	33					0	X	----	1	0						
JESXCF	SYSTEM	1	NS	FF	86					0	X	----		0	998					
ALLOCAS	SYSTEM	1	NS	FF	1628					0	X	----	1	0						

IBM Software Group | DB2 Information Management Software

# RMF III Option 3 STORFRMF

Jobname	C	Class	Cr	TOTAL	ACTV	IDLE	WSET	FIXED	DIV	AUX SLOTS	PGIN RATE	ES RATE
CORIDBM1	S	VEL50STC		352K	352K	0	352K	295K	37	0	0	
OMVS	S	SYSTEM		23235	23235	0	23235	1347	41	12660	0	
VLF	S	VEL50STC		17945	17945	0	17945	183	0	22512	0	
RMF	S	SYSSTC		10042	10042	0	10042	101	0	801	0	
*MASTER*	S	SYSTEM		4649	4649	0	4649	4229	0	812	0	
CORIDIST	S	VEL50STC		4531	4531	0	4531	122	0	0	0	
JAVARMI	S	VEL30STC		4292	4292	0	4292	82	0	0	0	
RMFGAT	S	SYSSTC		3539	3539	0	3539	75	0	665	0	
CORIMSTR	S	VEL60STC		2362	2362	0	2362	143	0	0	0	
ANTAS000	S	VEL30STC		2293	2293	0	2293	93	0	6455	0	
SMF	S	SYSTEM		2218	2218	0	2218	57	0	81	0	
ANTMAIN	S	SYSTEM		2009	2009	0	2009	123	0	6834	0	
SMSPDSE	S	SYSTEM		1776	1776	0	1776	311	0	708	0	
ALLOCAS	S	SYSTEM		1628	1628	0	1628	38	0	511	0	
CATALOG	S	SYSTEM		1521	1521	0	1521	706	0	668	0	
SMS	S	SYSSTC		1502	1502	0	1502	57	0	92	0	

Command ==> Scroll ==> CSR

25 © 2005 IBM Corporation

IBM Software Group | DB2 Information Management Software

# DB2 64-bit Summary

Legend:  
■ 64-bit AS  
■ 31-bit AS

2 GB

Components and their AS types:

- User pgm/st proc: 31-bit AS
- MSTR: 31-bit AS
- DBM1: 64-bit AS (includes Buffer pool, RID pool, Sort pool, BP ctl blks, Compr dict, Castout buf, DBD, DSC)
- IRLM: 64-bit AS (includes Locks, IRLM code)
- DDF: 31-bit AS
- DB2 code: 31-bit AS

+ materialized LOBs, accounting blocks, lock and buffer traces etc...

© 2004 IBM Corporation

26 © 2005 IBM Corporation

## Immediate Benefits



### Simplified buffer pool monitoring and tuning

- Only ONE type of buffer pool
- Buffer pool size limits increased
  - Eases the worry over monitoring once the system stabilizes

### DBM1 VSTOR constraint relief

- Increase CTHREAD - ECSA allocation may need increasing if CTHREAD raised
- Expect a single DB2 subsystem to run larger workloads
  - Could defer going to data sharing
  - Consolidating members in a group

### Expected increases in workload

- Maximum number of prefetch, deferred write, and castout engines is now 600 in order to decrease "engine not available" conditions

### IFCID's 217 and 225 reflect DBM1 VSTOR usage above and below 2GB bar

© 2004 IBM Corporation

## Key messages

- V8 64-bit support will not eliminate VSC below the 2GB bar in DBM1
- Will provide valuable additional relief, but will vary by installation
- Will be able to exploit all available processor storage on latest processor models (currently 256GB, current DB2 limit of 1TB)
  - Some **additional number of active threads and DBATs**
  - **XXL bufferpools** to eliminate I/O and speed up remainder
  - Increase exploitation of ESA **Compression**
  - **Larger thread Sort Pool**
- Must have **sufficient real storage** to fully back increased usage
- Must continue to plan for, **monitor and tune** VSTOR usage below 2GB bar
- **Additional exploitation of 64-bit virtual to move thread related storage above 2GB bar under consideration for future release**