




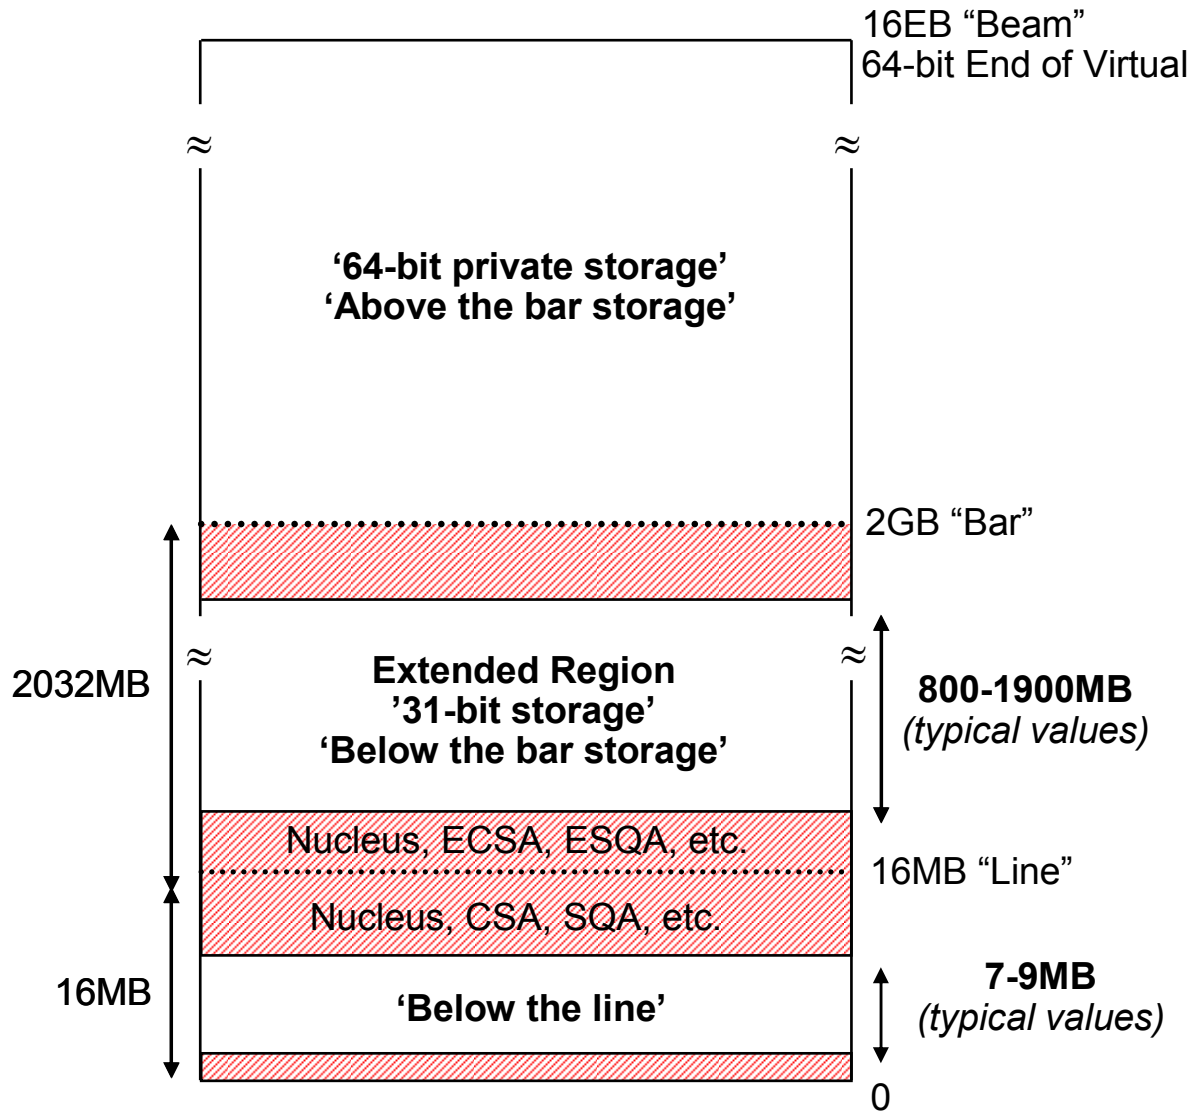
The Evolution of DB2 (Virtual and Real) Memory Monitoring

Florence Dubois
DB2 for z/OS Development
fldubois@uk.ibm.com





Virtual Storage Management – Terminology





31-bit Virtual Storage Constraint

800-1900MB
(typical values)




**Extended Region
'31-bit storage'**

(1) All Storage Except the Storage Cushion is consumed!
Full System Contraction starts to occur

▶ Potentially Degrading Performance

(2) Only 'must complete' processing succeeds
Individual DB2 threads (allied, DBAT) may abend with
04E/RC=00E200xx (e.g. 00E20003 & 00E20016)


(3) Ultimately DB2 subsystem may abend with S878 or
S80A due to non-DB2 subsystem component (e.g.
DFP) issuing unconditional MVS getmain





Tracking DB2 Storage


- IFC Records
 - IFCID 225 – Storage Summary
 - Enabled via Statistics Trace Class 1
 - Snapshot value as each DB2 Stats interval comes due
 - V8/V9 = interval defined by ZPARM STATIME – recommendation is 1 minute
 - V10 = written at fixed 1 minute intervals
 - SMF record type
 - V8 = Type 102
 - V9/V10 = Type 100 Subtype 4
 - IFCID 217 – Storage Detail Record
 - Storage usage at thread level
 - Effectively a dump SM=1 report but in IFC form
 - Available through Global Trace Class 10





Tracking DB2 Storage ...

- IFC Records ...
 - First class support provided by OMEGAMON XE for DB2 PM/PE
 - Statistics Trace | Report
 - Includes FILE and LOAD data base table support as well as upgrade (ALTER TABLE) of already installed table DB2PM_STAT_GENERAL
 - Record Trace Report
 - SPREADSHEETDD subcommand option
 - Generate a .csv file that can easily be loaded into a spreadsheet
 - REXX Tools (MEMU2, MEMUSAGE)
 - Available for download from DB2 for z/OS Exchange
 - <http://www.ibm.com/developerworks/software/exchange/db2zos>
 - » Click on 'View and download examples'
 - » The file is tagged with 'memu2'






Tracking DB2 Storage ...

- V8 APAR PK20800 8/07
 - DB2 command: -DISPLAY THREAD(*) SERVICE(STORAGE)
 - DSNV492I message that can be used by DB2 service for diagnostics

```
V91A      N  *      0 003.RCRSC 02 SYSOPR              0067      0
V490-SUSPENDED 07213-09:59:18.02 DSNRCRSC +00000230 01.51
V492-LONG 252 K VLONG 40 K 64 1028 K
```

- Includes Agent Local Non-System Storage usage
- Does not include Getmained Stack Storage usage
- The key values are the LONG storage pool and the VLONG storage pool values (252KB + 40KB = 292KB in the above example)
 - Reflect virtual storage consumption below the 2GB bar
 - May be used to identify poorly behaved applications or DB2 code issues






Tracking DB2 Storage ...

- DB2 Internal Monitor (V9 CM)
 - Automatically issues console messages (DSNV508I) when DBM1 31-bit virtual storage crosses (increasing or decreasing) thresholds
 - Threshold recently augmented in PM38435 to account for the storage cushion

```
DSNV508I -SE20 DSNVMON - DB2 DBM1 BELOW-THE-BAR
STORAGE NOTIFICATION
      77% CONSUMED
      76% CONSUMED BY DB2
      352M AVAILABLE OUT OF REGION SIZE 1553M
      WITH A 274M STORAGE CUSHION
```

- Identifies the agents that consume the most storage
- Can get status at any time using `-DIS THREAD(*) TYPE(SYSTEM)`

```
NAME      ST A   REQ ID          AUTHID   PLAN      ASID   TOKEN
VA1A      N  *    0 002.VMON 01   SYSOPR           002A    0
V507-ACTIVE MONITOR, INTERVALS=8216, STG=77%, BOOSTS=0, HEALTH=100
      REGION=1553M, AVAIL=352M, CUSHION=274M
```





MVS Storage Overview

- EXTENDED REGION SIZE (MAX) – QW0225RG
 - Total theoretical amount DB2 has access to
- 31 BIT EXTENDED LOW PRIVATE – QW0225EL
 - DB2 uses a small amount of Low private (bottom up storage)
 - DB2 code itself / reservation for pageset storage
- 31 BIT EXTENDED HIGH PRIVATE – QW0225EH
 - DB2 mostly uses subpool 229 Key 7 (top down storage)
- Other products also use address space storage
 - Dataset opens / DFP
 - SMF


DBM1 AND MVS STORAGE BELOW 2 GB	CONTINUED	QUANTITY
24 BIT LOW PRIVATE	(MB)	0.14
24 BIT HIGH PRIVATE	(MB)	0.47
31 BIT EXTENDED LOW PRIVATE	(MB)	130.86
31 BIT EXTENDED HIGH PRIVATE	(MB)	859.40
EXTENDED REGION SIZE (MAX)	(MB)	1296.00
EXTENDED CSA SIZE	(MB)	632.48





MVS Storage Overview ...

- ECSA – QW0225EC
 - Common storage area across all address spaces for a given LPAR
 - Large ECSA size would be 1GB with typical sizes being 300-500MB
 - Affects maximum available Extended Region
 - Biggest factor
 - Some customers due to the needs of other products have huge ECSA requirement leading to very small extended region size
 - Extensive use of ECSA by IMS across dependent regions
 - Mostly buffer pools, control blocks, data are in ECSA
 - Sizes are at user choice – For best performance they tend to be large
 - Not exploiting VSCR features of recent IMS releases
 - Generous over allocation for safety of ECSA and other common areas
 - Common LPAR image for Sysplex (best practice)





DB2 DBM1 Address Space Storage

- 31-bit / 24-bit DB2 storage
 - Getmained
 - Variable
 - Fixed Storage
 - Stack storage
- Non-DB2 getmained
 - SMF
 - Dataset / pageset


DBM1 AND MVS STORAGE BELOW 2 GB		QUANTITY
-----		-----
TOTAL DBM1 STORAGE BELOW 2 GB	(MB)	1092.89
TOTAL GETMAINED STORAGE	(MB)	143.62
EDM POOL	(MB)	97.66
TOTAL VARIABLE STORAGE	(MB)	846.56
TOTAL AGENT LOCAL STORAGE	(MB)	732.12
TOTAL AGENT SYSTEM STORAGE	(MB)	173.48
NUMBER OF PREFETCH ENGINES		412.00
NUMBER OF DEFERRED WRITE ENGINES		300.00
NUMBER OF CASTOUT ENGINES		300.00
NUMBER OF GBP WRITE ENGINES		300.00
NUMBER OF P-LOCK/NOTIFY EXIT ENGINES		461.00
TOTAL AGENT NON-SYSTEM STORAGE	(MB)	558.64
TOTAL NUMBER OF ACTIVE USER THREADS		343.50
NUMBER OF ALLIED THREADS		8.95
NUMBER OF ACTIVE DBATS		322.60
NUMBER OF POOLED DBATS		11.95
RID POOL	(MB)	1.29
PIPE MANAGER SUB POOL	(MB)	0.00
LOCAL DYNAMIC STMT CACHE CNTL BLKS	(MB)	0.99
THREAD COPIES OF CACHED SQL STMTS	(MB)	86.64
IN USE STORAGE	(MB)	44.78
STATEMENTS COUNT		5719.68
HWM FOR ALLOCATED STATEMENTS	(MB)	50.49
STATEMENT COUNT AT HWM		6240.00
DATE AT HWM		09/22/11
TIME AT HWM		02:18:52.46
BUFFER MANAGER STORAGE CNTL BLKS	(MB)	19.43
TOTAL FIXED STORAGE	(MB)	2.93
TOTAL GETMAINED STACK STORAGE	(MB)	99.78
TOTAL STACK STORAGE IN USE	(MB)	81.94
STORAGE CUSHION	(MB)	175.84





DB2 DBM1 Address Space Storage ...


- Getmained - QW0225GM
 - V7 – EDM, Compression Dictionaries, Bufferpools and miscellaneous
 - V8 – Compression Dictionaries and Bufferpools are moved above the bar
 - V9 – Part of the EDM pool (100% of SKCT/SKPT and ~30% of CT/PT) is moved above the bar
- Variable Storage - QW0225VR
 - Most interesting from a tuning perspective
 - Variable length blocks
 - Thread pools (AGL)
 - Used by both System and User
 - Local Dynamic Statement Cache
 - V9 – Approx. ~50% of LDSC is moved above the 2GB bar
- Fixed Storage - QW0225FX
 - High performance storage
 - Fixed length blocks
 - Not usually so interesting from a tuning perspective
 - Small change in the great scheme of things
- Stack Storage - QW0225GS
 - Save areas
 - Working program variables
 - Small amounts of high speed storage allocations
 - Cached in the DB2 address space to allow greater performance
 - Compressed only at full system contraction





Non-DB2 31-bit Storage in DBM1 Address Space

- Not tracked by DB2
- Non-DB2 storage is high private storage
 - **TOTAL DBM1 STORAGE** = TOTAL GETMAINED STORAGE QW0225GM + TOTAL GETMAINED STACK STORAGE QW0225GS + TOTAL FIXED STORAGE QW0225FX + TOTAL VARIABLE STORAGE QW0225VR
 - **NON-DB2 STORAGE** = MVS 31 BIT EXTENDED HIGH PRIVATE QW0225EH – TOTAL DB2 DBM1 STORAGE
- Used usually by MVS functions such as SMF
- Option DETAIL in SMFPRMxx can cause storage to creep and become very large
 - The big hit to DB2 in this area is the DDNAME tracking: allocation does not realise that we have closed off a page set and reallocated it again
 - SMF Type 30 subtype 4 and 5 will track all the DDNAMES
 - Most environments do not need SMF Type 30 subtype 4 and 5
 - Recommend NODETAIL





DB2 DBM1 Address Space Storage ...


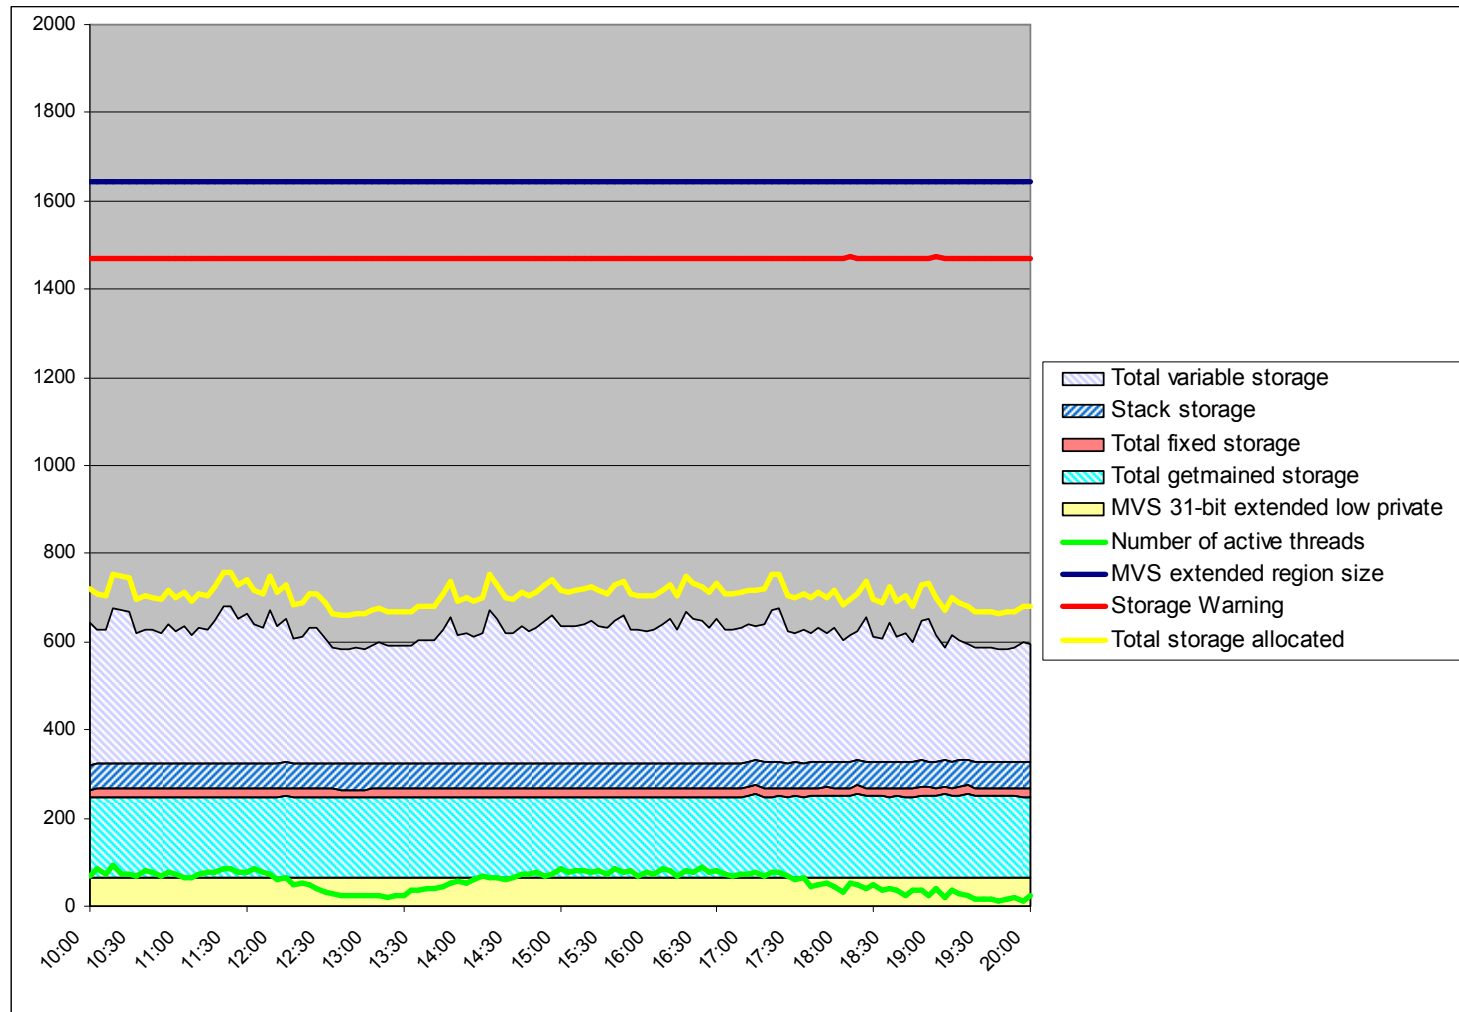
- 64-bit DB2 storage
 - Getmained
 - Fixed
 - Variable
 - Compression Dictionaries
 - DBD Pool
 - Dynamic Statement Cache
 - RDS Pool Above (V9)
 - Skeleton Pool (V9)
 - Buffer Pools
 - Buffer Control Blocks
 - Castout Buffers
 - 64-bit Shared Private Storage (V9)

DBM1 STORAGE ABOVE 2 GB		QUANTITY
GETMAINED STORAGE	(MB)	4585.91
FIXED STORAGE	(MB)	33.85
VARIABLE STORAGE	(MB)	1387.48
COMPRESSION DICTIONARY	(MB)	276.27
IN USE EDM DBD POOL	(MB)	90.43
IN USE EDM STATEMENT POOL	(MB)	330.10
IN USE EDM RDS POOL	(MB)	0.00
IN USE EDM SKELETON POOL	(MB)	1.67
STAR JOIN MEMORY POOL	(MB)	N/A
STORAGE MANAGER CONTROL BLOCKS	(MB)	N/A
VIRTUAL BUFFER POOLS	(MB)	29542.97
VIRTUAL POOL CONTROL BLOCKS	(MB)	928.85
CASTOUT BUFFERS	(MB)	37.50
SHARED GETMAINED STORAGE	(MB)	3.58
SHARED FIXED STORAGE	(MB)	17.27
SHARED VARIABLE STORAGE	(MB)	2533.17



Storage Monitoring Using IFC225



- Graphing IFC 225 data basic elements to study evolutionary trend





Storage Monitoring Using IFC225 ...


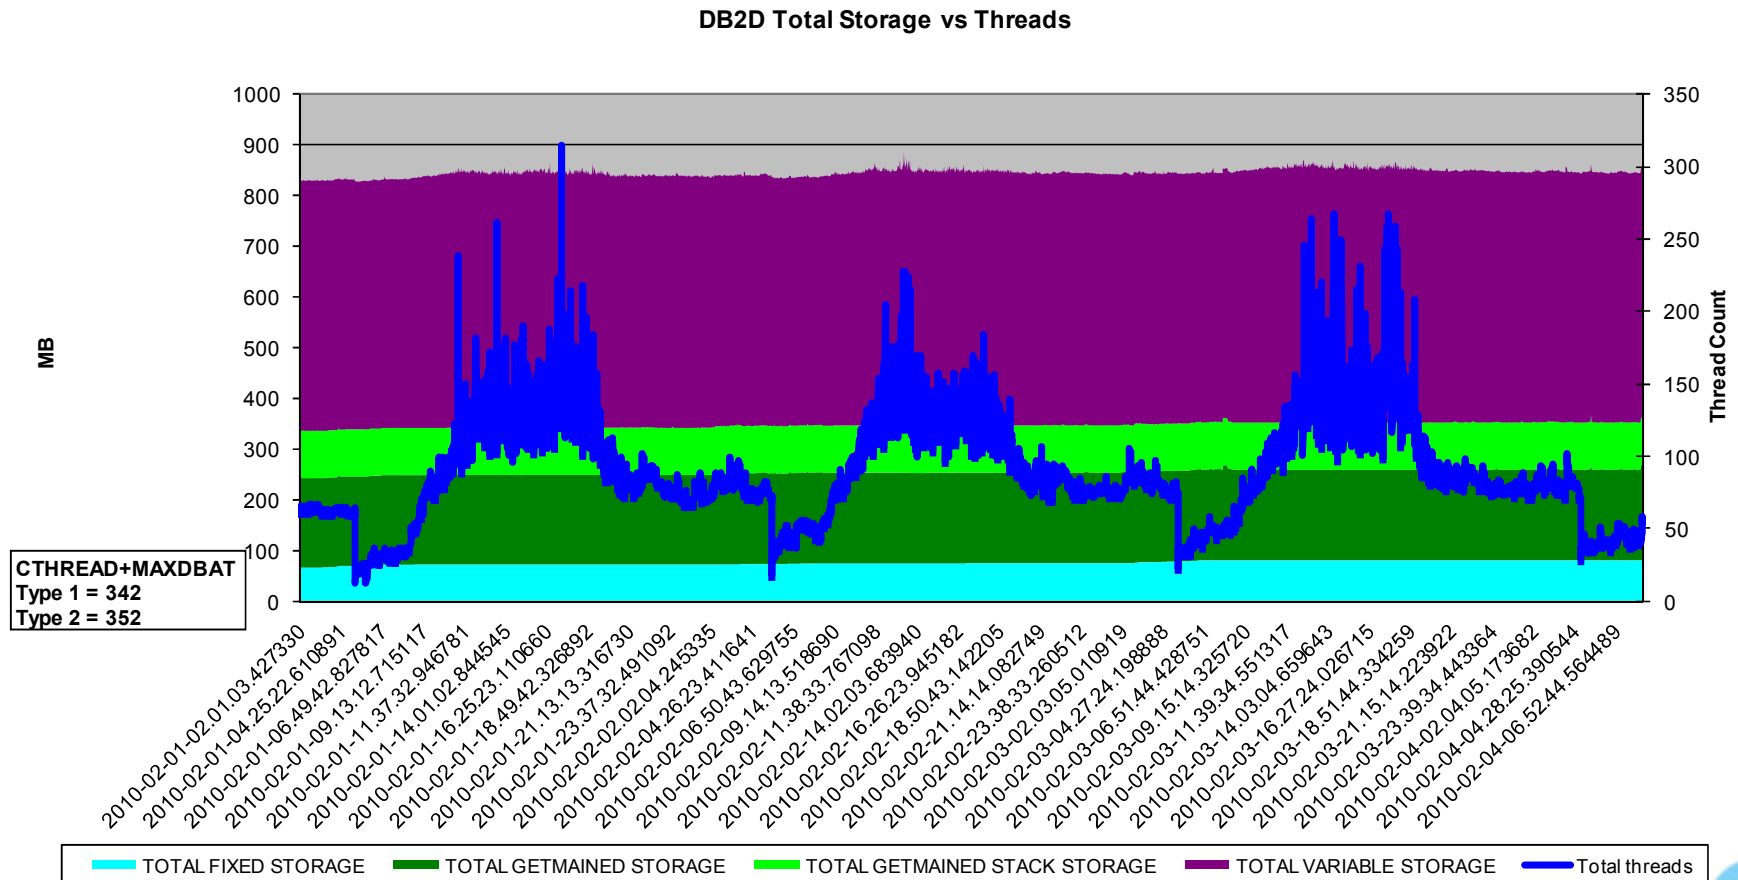
- 7 days data (Mon-Sun) - Leaky subsystem?





Storage Monitoring Using IFC225 ...


- What happened next Mon-Wed
 - This DB2 took a full week to “warm up”





Lessons To Be Learned


- Wrong data time can lead to erroneous conclusion
 - Full week showed a possible leak
 - 10 days showed DB2 storage usage to be stable
- Ideal data is from DB2 startup to DB2 shutdown
 - If not possible, then get as much as you can
- Do not think you know the point where the maximum usage is – you may be way off the mark
- More data will lead you to a better result





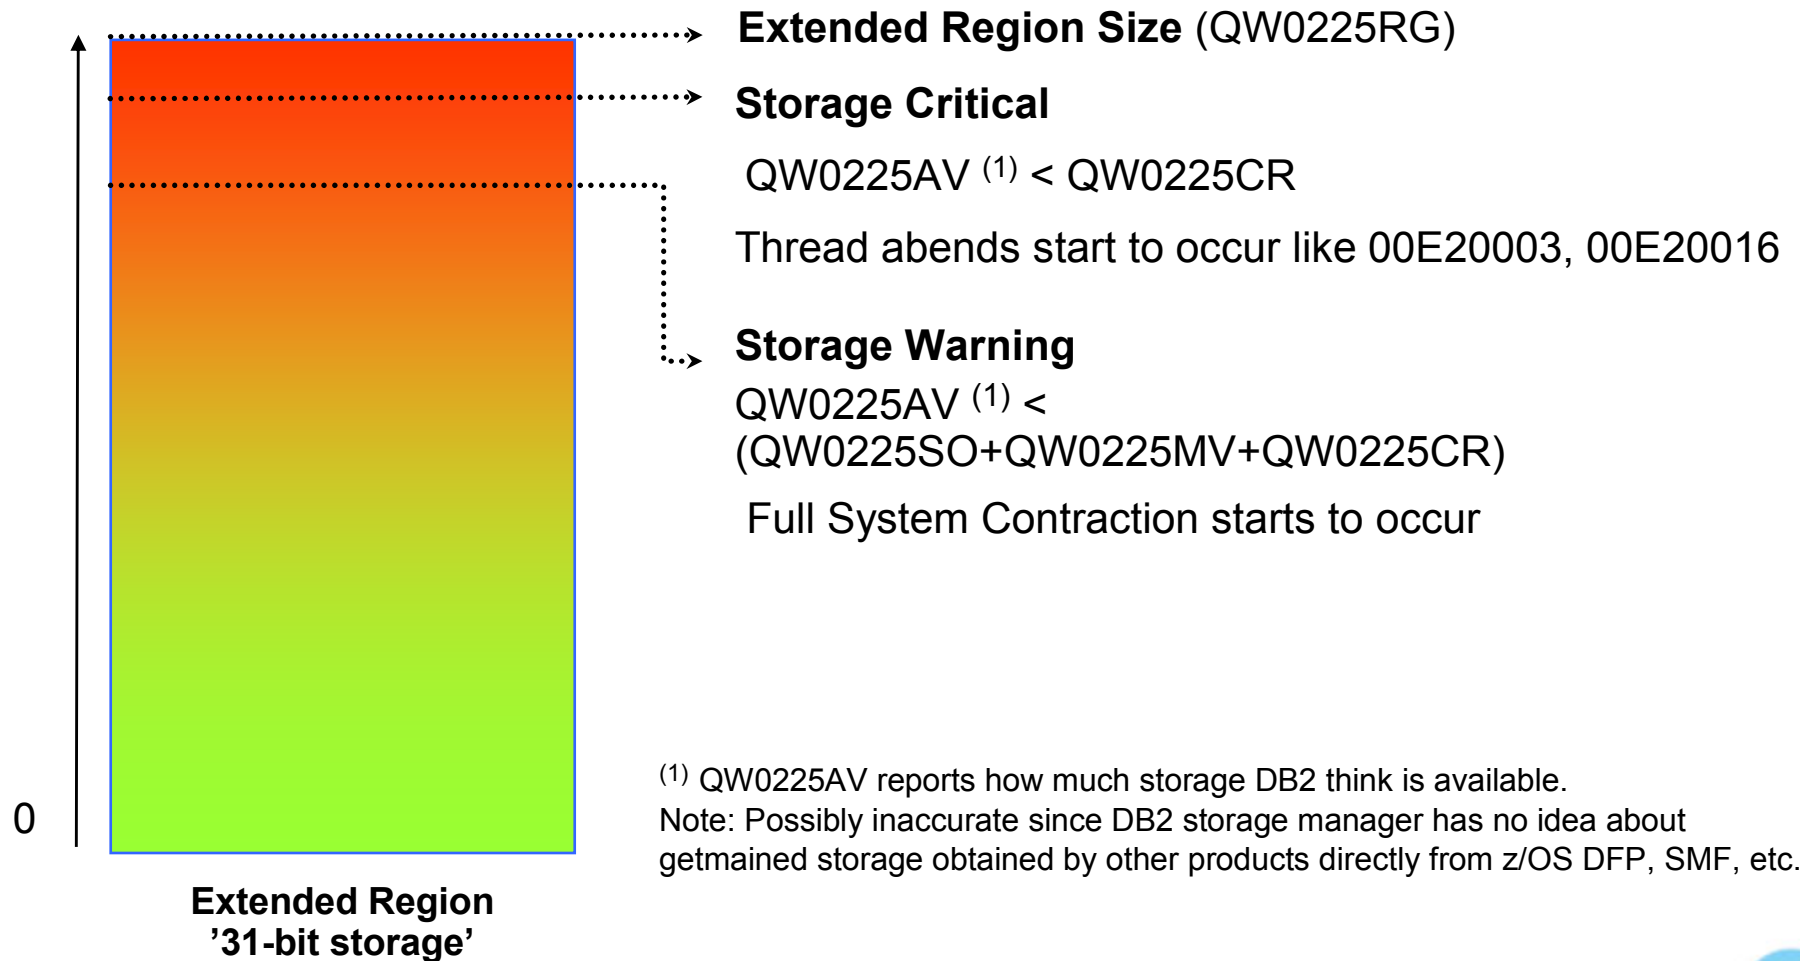
Storage Overuse: DB2 Storage Contraction

- When 'running low' on extended virtual, DB2 begins system contraction process
 - Attempts to freemain any available segments of storage
- 3 critical numbers for contraction
 - Storage reserved for must complete (e.g. ABORT, COMMIT)
 - $QW0225CR = (CTHREAD + MAXDBAT + 1) * 64K + 25M$
 - Fixed, real value
 - Storage reserved for open/close of datasets
 - $QW0225MV = (DSMAX * 1300) + 40K$
 - Virtual number and no guarantee
 - Warning to contract
 - $QW0225SO = \text{Max}(5\% \text{ of Extended Region Size, } QW0225CR - 25M)$
 - Storage Cushion = $QW0225CR + QW0225MV + QW0225SO$





Storage Overuse: DB2 Storage Contraction



(1) QW0225AV reports how much storage DB2 think is available.
Note: Possibly inaccurate since DB2 storage manager has no idea about getmained storage obtained by other products directly from z/OS DFP, SMF, etc.




Storage Overuse: DB2 Storage Contraction

****WARNING** DO NOT SPECIFY CTHREAD + MAXDBAT TOO HIGH IN DB2 V8 OR THE CUSHION WILL BE VERY LARGE**

- Examples:


	Case 1	Case 2	Case 3
CTHREAD	2000	400	400
MAXDBAT	2000	2000	150
DSMAX	15000	15000	15000
MVS extended region size (MB)	1700	1700	1700
Storage reserved for must complete (MB)	275	178	63
Storage reserved for datasets (MB)	19	19	19
Warning to contract (MB)	225	128	85
Storage Cushion (MB)	519	325	166





CONTSTOR


- Thread storage contraction turned on by zparm CONTSTOR=YES
 - Online changeable with immediate effect
- Maximum of 1 compress every 5 commits so very cheap to implement
- Ineffective for long-running persistent threads with use of RELEASE(DEALLOC)
- Compresses out part of Agent Local Non-System storage
 - Only compresses LONG storage (as per SERVICE(STORAGE))
 - Does not compress
 - Agent Local System, Getmained Stack Storage
 - Local Dynamic Statement Cache
- Controlled by two hidden zparms
 - SPRMSTH @ 1048576 and SPRMCTH @ 10
- Triggers
 - No. of Commits > SPRMCTH, or
 - Agent Local Non-System > SPRMSTH and No. of Commits > 5





MINSTOR


- Best fit algorithm for thread storage turned on by zparm MINSTOR=YES
 - Online changeable, may not have an effect due to already-cached pools
 - Restart recommended if this parm changed
- Changes the storage management of the user AGL POOL to “Best fit” rather than “First fit”
 - In order to find the best fit piece of storage, CPU cycles are used to scan and maintain ordered storage
 - In a POOL with low fragmentation, MINSTOR may not have a great effect but will cost CPU
- Only enable if fragmentation is a big issue
 - Only the SM=4 option of the DB2 Dump Formatter and a dump will really give you the definitive answer





Protecting the System


- Plan on a 'Basic' storage cushion = Storage cushion + 100MB
 - Avoid hitting short-on-storage and driving Full System Contraction
 - Provide some headroom
 - For tuning, some growth, Fast Log Apply, abnormal operating conditions
- Estimate the maximum number of threads that can be supported
 - Assuming the storage is proportional to the amount of threads, it is possible to predict a theoretical max. number of concurrent threads
 - It may be possible to run with more threads than the formula dictates, but there is the danger that large threads may come in and cause out of storage conditions
- Set zparm CTHREAD and MAXDBAT to realistic values that can be supported
 - CTHREAD and MAXDBAT are the brakes on the DB2 subsystem
 - Theoretical maximum: $CTHREAD + MAXDBAT = 2000$
 - Practical maximum is much less (typical range 300-850)
 - Avoid over committing resources
 - Deny service and queue work outside the system to keep system alive





Estimating Maximum Number of Threads

- Collect IFCID 225 since the start of DB2
 - Month end processing
 - Weekly processing
 - Utilities processing
 - Try to use a full application mix cycle
- Focus on time periods with
 - Increasing number of allied threads + active DBATs
 - Increasing use of getmained stack storage
 - Increasing use of AGL non-system
- Adjust the formula based on workload variations
- Protect the system by always using a pessimistic approach
 - Optimistic may mean a DB2 outage
- Recalculate on a regular basis as new workloads and/or parameters are changed
 - Continuous process






Estimating Maximum Number of Threads ...

- 'Basic' storage cushion (BC)
 - $(BC) = QW0225CR + QW0225MV + QW0225SO + 100MB$
- Calculate Max non-DB2 storage (ND)
 - $(ND) = \text{MAX}(\text{MVS 31 BIT EXTENDED HIGH PRIVATE } QW0225EH - \text{TOTAL GETMAINED STORAGE } QW0225GM - \text{TOTAL GETMAINED STACK STORAGE } QW0225GS - \text{TOTAL FIXED STORAGE } QW0225FX - \text{TOTAL VARIABLE STORAGE } QW0225VR)$
- Max. allowable storage (AS)
 - $(AS) = QW0225RG - (BC) - (ND)$
- Max. allowable storage for thread use (TS)
 - $(TS) = (AS) - (\text{MAX}(\text{TOTAL AGENT SYSTEM STORAGE } QW0225AS) + \text{MAX}(\text{TOTAL FIXED STORAGE } QW0225FX) + \text{MAX}(\text{TOTAL GETMAINED STORAGE } QW0225GM) + \text{MAX}(\text{MVS 31 BIT EXTENDED LOW PRIVATE } QW0225EL))$
- Average thread footprint (TF)
 - $(TF) = (\text{TOTAL VARIABLE STORAGE } QW0225VR - \text{TOTAL AGENT SYSTEM STORAGE } QW0225AS + \text{TOTAL GETMAINED STACK STORAGE } QW0225GS) / (\text{Allied threads } QW0225AT + \text{DBATs } QDSTCNAT)$
- Max threads supported = $(TS) / (TF)$


Remember to use the MAX impact value across all available data e.g. maximum system storage





Options to Provide Some Relief in V8/V9


- Balance EDM pool size with CTHREAD/MAXDBAT
 - Race condition as more threads come into play
 - Queuing or resource unavailable “-904” or storage abends
 - EDM pool tuning ROTs
 - FAILS DUE TO POOL FULL = 0
 - Very serious condition
 - % NON-STEALABLE PAGES IN USE (CTs/PTs) < 50% (V8)
 - % NON-STEALABLE PAGES IN USE (CTs/PTs) < 75% (V9)





Options to Provide Some Relief in V8/V9 ...


- Reduce use of BIND option RELEASE(DEALLOCATE)
 - Overuse with persistent threads can create a virtual storage issue
 - Accumulated storage can be left around until deallocation
 - Ineffective thread and full system storage contraction
 - Also drive up demand for EDM Pool resources
 - Best reserved for
 - High volume and/or performance-sensitive at reasonable volume OLTP plans/packages
 - Long-running batch programs that take frequent intermediate commits





Options to Provide Some Relief in V8/V9 ...

- Reduce size of the Local Dynamic Statement Cache (LDSC)
 - Enabled by BIND option KEEP DYNAMIC(YES)
 - Prepared statements are kept in thread storage across commit so that prepares can be avoided
 - Least Recently Prepared statements are thrown away from the LDSC at commit based on zparm MAXKEEPD
 - May see increase in number of “short prepares” and “full prepares” with associated increase in CPU resource consumption
 - Increase size of the Global Dynamic Statement Cache above the 2GB bar (zparm EDMSTMTC) to compensate by trying to reduce the number of “full prepares” and offset the increase in CPU resource consumption
 - Recommendation
 - Reduce zparm MAXKEEPD incrementally, and
 - Increase size of Global Dynamic Statement Cache above the 2GB bar






Options to Provide Some Relief in V8/V9 ...

- Invalidate statements from the LDSC and release the associated storage ahead of commit
 - MAXKEEPD is only enforced at commit
 - APAR PK21861 introduced new zparm CACHEDYN_FREELOCAL
 - Allow statements to be purged at end of section
 - CACHEDYN_FREELOCAL settings

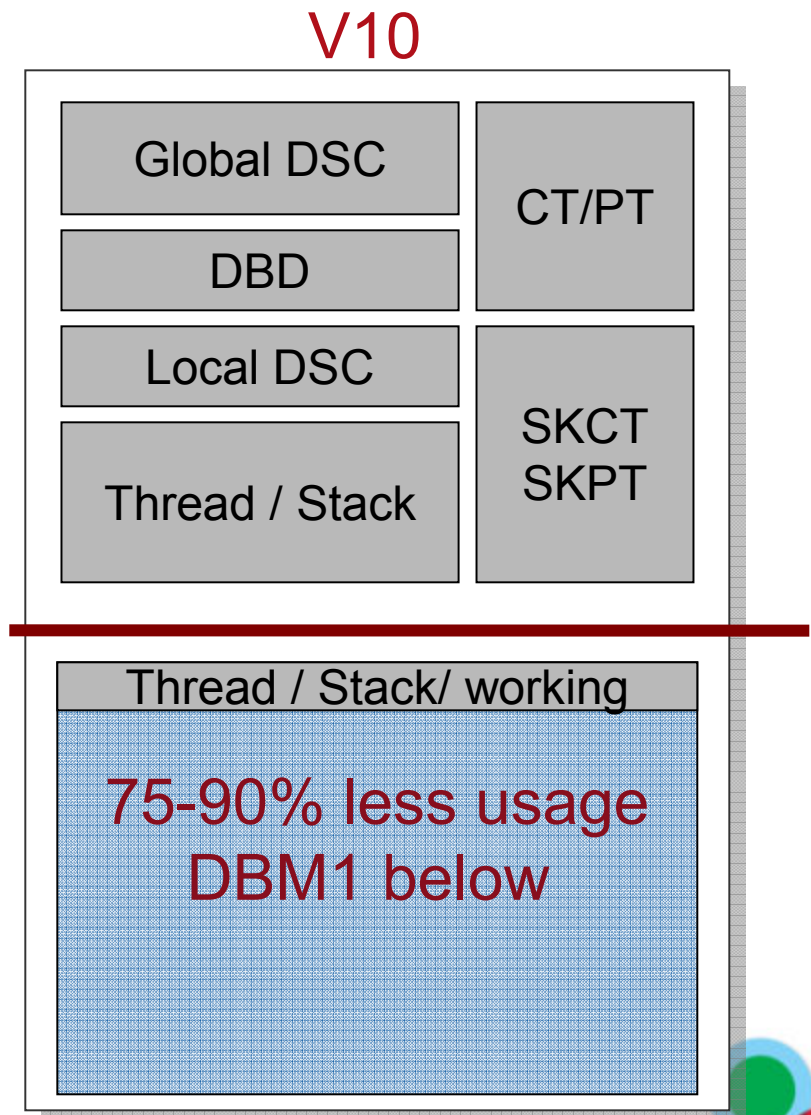
0	Off (default V8)
1	If (LDSC >=500MB & DBM1 Used >=75%) then free >= 100KB statement If DBM1 Used >=85% then free any statement (default V9)
2	If (LDSC >=500MB & DBM1 Used >=80%) then free >= 100KB statement If DBM1 Used >=88% then free any statement
3	If (LDSC >=350MB & DBM1 Used >=75%) then free >= 100KB statement If DBM1 Used >=88% then free any statement





DBM1 Virtual Storage Constraint Relief in DB2 10


- DBM1 below 2GB
 - 75-90% less usage in V10 compared to V9 after REBIND
 - Some of working storage (stack, xproc storage) stays below 2GB
- Supports larger number of threads
 - Possible data sharing member consolidation
- Opportunities to improve CPU by reversing VSCR tuning actions
 - More thread reuse
 - More release deallocate
 - High performance DBATs
 - Larger MAXKEEPD values for KEEP DYNAMIC=YES
 - Must provision additional real storage to back the requirement





REAL Storage Constraint


- Shortage of real storage
 - Can lead to excessive paging and severe performance issues
 - Important subsystems such as DB2 should not be paging IN from AUX (DASD)
 - Recommendation to keep page-in rates low (near zero)
 - Monitor using RMF Mon III
 - Monitor in DB2 page in for reads/writes and output log buffer paged in
 - Ultimately, can take the LPAR out
 - Once all AUX is consumed, the LPAR goes into a wait state
 - Can lead to long DUMP processing times and cause major disruption
 - DUMP should complete in seconds to make sure no performance problems ensue
 - Once paging begins, it is possible to have DUMP processing take 10s of minutes with high risk of system-wide or even sysplex-wide slowdowns
 - Could be running 'one DUMP away from a disaster'





REAL Storage Constraint


- Shortage of real storage can lead to ...
 - Wasted opportunities for CPU reduction
 - Reluctance to use bigger or more buffer pools
 - Reluctance to use buffer pool long-term page fix
 - Avoid the repetitive cost of page fix and free for each and every I/O
 - Up to 8% reduction in overall IRWW transaction CPU time
 - Many performance opportunities in DB2 10 require real storage
- General recommendation
 - Not good practice to configure REAL storage based on normal operating conditions and rely on DASD paging space to absorb the peaks
 - Need sufficient REAL storage to cover both the normal DB2 working set size and MAXSPACE requirement
 - With V9, MAXSPACE can typically up to 8-9GB
 - Undersizing MAXSPACE will result in partial dumps
 - Seriously compromise problem determination





Virtual vs. REAL Storage

- Virtual storage below 2GB bar is usually densely packed (as before in V7)
 - VIRTUAL=REAL is a fair approximation
- Virtual storage above the bar number may be misleading
 - Above-the-bar memory is sparsely populated
 - Backing rate is low for 64-bit storage
 - DB2 allocates very large memory objects
 - No need to back by REAL storage frames until first reference
 - VIRTUAL will not equal REAL






Monitoring REAL Storage

- Real storage needs to be monitored as much as Virtual storage
 - Need to pay careful attention to QW0225RL (Real frames in use by DBM1) and QW0225AX (Auxiliary frames)
 - Ideally QW0225RL should be significantly less than the amount of virtual consumed

REAL AND AUXILIARY STORAGE		QUANTITY
REAL STORAGE IN USE	(MB)	5958.66
AUXILIARY STORAGE IN USE	(MB)	0.00


- High QW0225AX should be investigated
 - Indicates that there are periods when the DB2 working set is pushed out to auxiliary storage (DASD)
 - » Need to understand what is the driver
 - Excessive amounts of storage on AUX may cause long DUMP times and severe performance issues
 - » Paging may become severe





How to Limit REAL Storage


- Hidden ZPARM SPRMRSMX ('real storage kill switch')
 - Delivered in APAR PK18354
 - Not widely broadcasted
- Prevents a runaway DB2 subsystem from taking the LPAR down
 - Should be used when there is more than one DB2 subsystem running on the same LPAR
 - Aim is to prevent multiple outages being caused by a single DB2 outage
 - Should be set to 1.5x to 2x normal DB2 subsystem usage
 - Kills the DB2 subsystem when SPRMRSMX value is reached





Summary

- Do not put your business at risk by running lean on memory
 - DBM1 31-bit virtual storage (V8/V9)
 - REAL memory provisioned on the LPAR (applies to all DB2 versions)
- Protect your DB2 subsystems by setting CTHREAD and MAXDBAT to realistic values that can be supported (applies to all DB2 versions)
- Provision sufficient REAL storage to cover both the normal DB2 working set size and MAXSPACE requirement (applies to all DB2 versions)
- Do not compromise potential CPU and performance benefits by starving your memory (applies to all DB2 versions)
 - Plan ahead to have enough memory to take advantage of DB2 10 capabilities





Questions ?





Disclaimer

© Copyright IBM Corporation 2011. All rights reserved.
U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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

IBM, the IBM logo, ibm.com, and DB2 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

