

Disclaimer

© Copyright IBM Corporation [current year]. 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, DB2, and 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

DB2 for z/OS Adoption

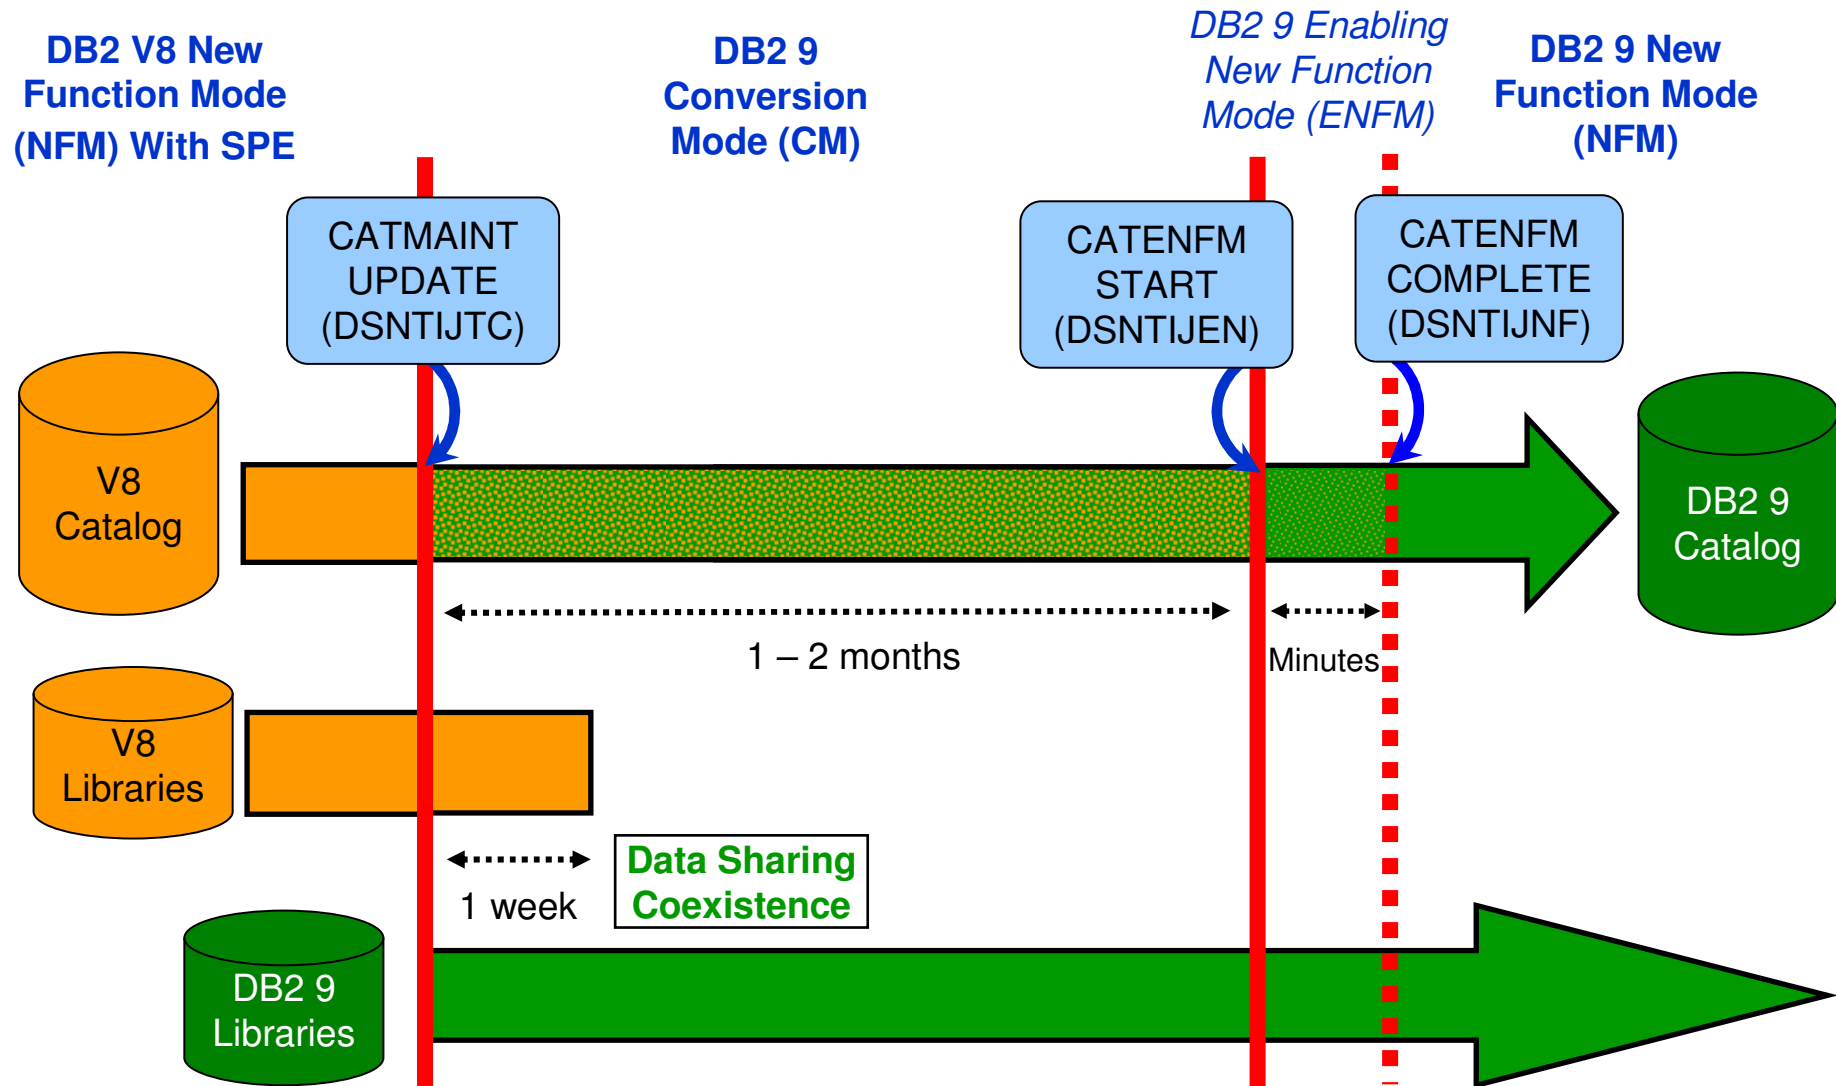
- **>85% WW Customers are Current**
- **DB2 V8: Majority have Migrated**
 - 100% of Top 100 >99% of Top 200
 - V7 End of Service: June 30, 2008
 - V8 Withdrawal from Marketing
 - Announced: Dec. 2, 2008
 - Effective: Sept. 8, 2009
- **DB2 9: Climbing Sharply**
 - 39% of Top 100 31% of Top 200
 - Over 500 customers
- **How's the Quality (compared to V8)?**
 - Lower overall PMR volume
 - Less Severity 1 APARs
 - Lower PE rate



DB2 9 Migration Experiences So Far

- Every customer experience is different
- Procedural upgrade easier
- Significant performance improvements
 - 5-10% CPU reduction
 - Up to 40% elapsed time improvement
- Aggressive optimization design
 - Significant performance improvement for majority of scenarios
 - Some cases of query performance regression
 - Fixes in Query analysis and re-write, Optimization, Cost estimation, and query runtime
 - Some access path selection instability due to host variables (static SQL) or parameter markers (dynamic SQL) in predicates
 - Some incorrout issues have been encountered
 - Complex queries with extensive use of In-Subquery or Exist Predicates
 - Tablespaces with predicate on partition column

Migration Overview



Conversion Mode

- After CATMAINT has been completed, DB2 will be running in CM mode
- Allow fallback/coexists with prior release.
- No new function is allowed.
- Objects created in this mode can be used from the lower level members.

Enabling New Function Mode


V8	V9
<ul style="list-style-type: none">▪ JCL Job – DSNTIJNE▪ CHAR -> VARCHAR expansion▪ Increasing page size▪ EBCDIC -> UNICODE	<ul style="list-style-type: none">▪ JCL Job – DSNTIJEN▪ SYSOBJ from 8K to 16K page size▪ SYSPACKAGE.SEQNO from SMALLINT to BIGINT▪ RTS tables become part of catalog (two new tables in catalog)

- There are several job steps in DSNTIJNE/DSNTIJEN.
- First step CATENFM START puts DB2 in ENFM mode.
- Once in ENFM mode, no fallback/coexistence is allowed.
- Incompatible catalog changes will be made when running the job.
- No new function is allowed yet.

New Function Mode

- After DSNTIJNE/EN completed all the enabling new functions, you can then run DSNTIJNF (CATENFM COMPLETE) to put DB2 in NFM
- If it is data sharing, all members will be notified in this step and the whole data sharing group will be put in NFM.
- New functions will be available in this mode.
- You can disable the new functions by going back to ENFM mode in V8, or CM*/ENFM* modes in V9.

V7-V8-V9 Catalog Migration – 700MB catalog



Plan Stability

- New function of DB2 9
 - Protects customers against access path regression
 - Allows for a “safe” way to REBIND (fall back)
 - Available even in DB2 9 (CM) as it can benefit migration and fallback
 - Strongly recommended
 - Make sure that the pre-conditioning APAR for Plan Stability (PK52522) is applied on all V8 (NFM) systems
- What is the problem?
 - REBINDs can cause access path changes
 - Most of the time, this improves query performance ...
 - ... But when it doesn't
 - No easy way to undo the REBIND
 - Can lead to a lot of grief to our customers and to IBM

Plan Stability ...

- Existing “solutions” inadequate
 - Preventing REBINDS altogether
 - REBINDing into alternate collections
 - Using hints
- What is the solution?
 - At REBIND, DB2 will save old copies of packages
 - In the event of a performance regression, users will have a way to fallback to an older copy

Plan Stability ...

- At REBIND, save old copies of packages
 - Catalog tables
 - Directory
- Two flavors
 - BASIC and EXTENDED
 - Controlled by new ZPARM PLANMGMT
 - Default is OFF
 - Also supported as REBIND options
- REBIND PACKAGE ...
 - PLANMGMT(BASIC)
 - 2 copies: Current and Previous
 - PLANMGMT(EXTENDED)
 - 3 copies: Current, Previous, Original
- Most bind options can be changed at REBIND
 - But a few must be the same ...
- REBIND PACKAGE ...
 - SWITCH(PREVIOUS)
 - Switch between current & previous
 - SWITCH(ORIGINAL)
 - Switch between current & original
- FREE PACKAGE ...
 - PLANMGMTSCOPE(ALL) – Free package completely
 - PLANMGMTSCOPE(INACTIVE) – Free all old copies
- Catalog support
 - SYSPACKAGE reflects active copy
 - SYSPACKDEP reflects dependencies of all copies
 - Other catalogs (SYSPKSYSTEM, ...) reflect metadata for all copies
- Invalidation and Auto Bind
 - Each copy invalidated separately
 - Auto bind replaces only the current copy – previous and original are not affected

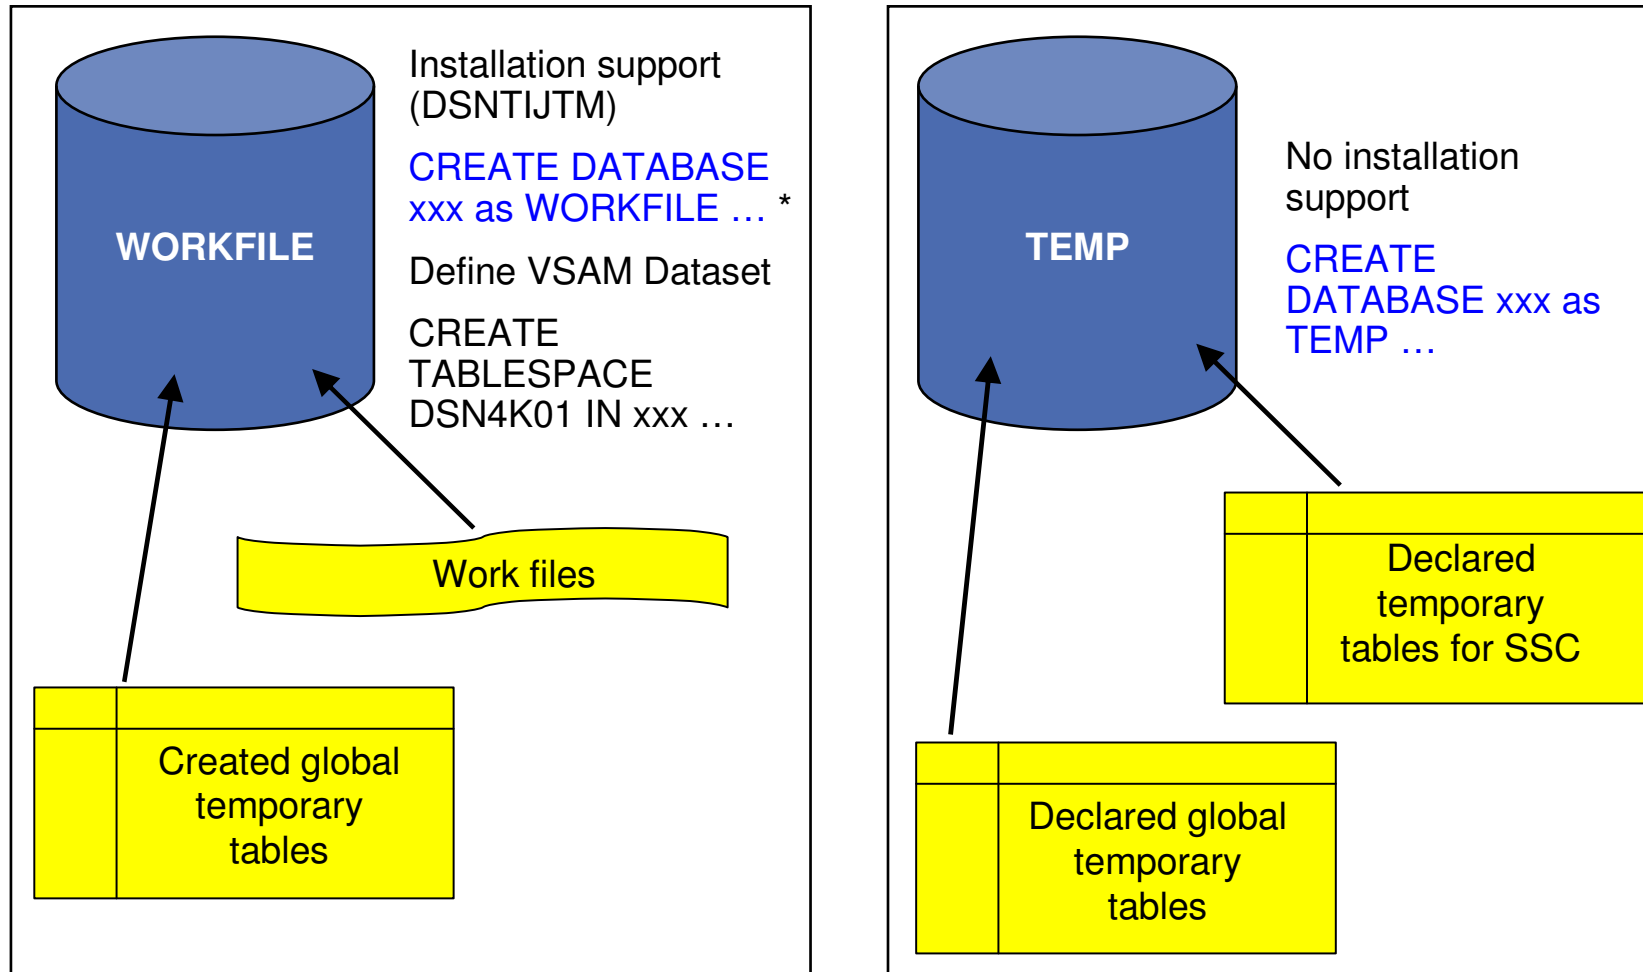
Sample Strategy for Migration using Plan Stability

- Migration strategy with Plan Stability ...
 - Delay rebind on DB2 9 (CM) until running stable
 - Do not rebind on DB2 9 (CM) until RUNSTATS has been run
 - Run with default of RUNSTATS TABLE(ALL) INDEX(ALL) KEYCARD
 - Make sure new ZPARM STATCLUS=ENHANCED (Default)
 - Introduces major change to CLUSTERRATIO calculation in DB2 9 and introduction of new statistic Data Repetition Factor (DRF)
 - Use DB2 supplied RUNSTATS utility if ISV utility does not yet support the enhanced statistics collection
 - Recommended Optimizer ZPARM changes
 - Set OPTIOWGT=ENABLE (APAR PK75643 will reflect this recommendation)
 - Set OPTIXOPREF=ON (APAR PK77426 will reflect this recommendation)
 - Several INCORROUT are related to the new global query block optimisation
 - Set OPTXQB=OFF (temporary) to disable this function

Sample Strategy for Migration using Plan Stability


- Why REBIND plans and packages?
 - As a general rule it is always a good recommendation to rebind on the new release to get the best performance
 - Almost all Optimizer enhancements available in DB2 9 (CM)
 - Re-enable SPROCs
 - Many plans/packages have SPROCs for fast column processing
 - All plans/packages with SPROCs that were bound prior to DB2 9 will be disabled
 - As a result DB2 will build SPROCs dynamically at execution time
 - Typical CPU performance impact in 0 to 10% range
 - Non-zero value for BYPASS COL indicator of problem
 - Performance trace of IFCID 224 identifies plans and packages which need rebinding to re-enable SPROCs
 - Exploit virtual storage constraint relief for static SQL

Temporary Space - The DB2 V8 Picture



* Only in a data sharing environment – in non-data sharing syntax is `CREATE DATABASE DSNDB07`

Temporary Space - The DB2 9 Picture



Installation and migration support
(REXX program called by DSNTIJTM)

```
CREATE DATABASE xxx as WORKFILE;
```

```
DSNTWFG DB41 DB2ADM xxx +  
3 10 16 BP0 SYSDEFLT +  
1 20 16 BP32K SYSDEFLT
```


- Declared Global Temporary Tables and Static Scrollable Cursors now use the WORKFILE database instead of the TEMP database
- Uses DB2-managed (instead of user-managed) storage in SYSDEFLT storage group
- Segmented table space organisation (user-defined SEGSIZE or default of 16)
- 4KB and 32KB page sizes only – no 8KB or 16KB

Planning For - Converged TEMP Space

- Migration from DB2 V8
 - To reclaim TEMP database storage, *YOU* must drop the TEMP database and reallocate the storage
 - Recommendation: Do not drop the TEMP database until you are sure that you will not return be falling back to V8, to avoid having to recreate it after fallback
- New installation panel for work file database definitions (DSNTIP9)
 - In migration mode, if you specify non-zero values
 - Migration job DSNTIJTM will create additional DB2-managed WORKFILE table spaces in the SYSDEFLT storage group → new REXX program DSNTWFG
 - DB2 does not take into account the existing work file table spaces
- Recommendation: set the 'DSVCI' ZPARM to YES to allow DB2 to match VSAM CI size to table space page size
- Ensure you have 32KB WORKFILE table spaces for Declared Global Temporary Tables and Static Scrollable Cursors

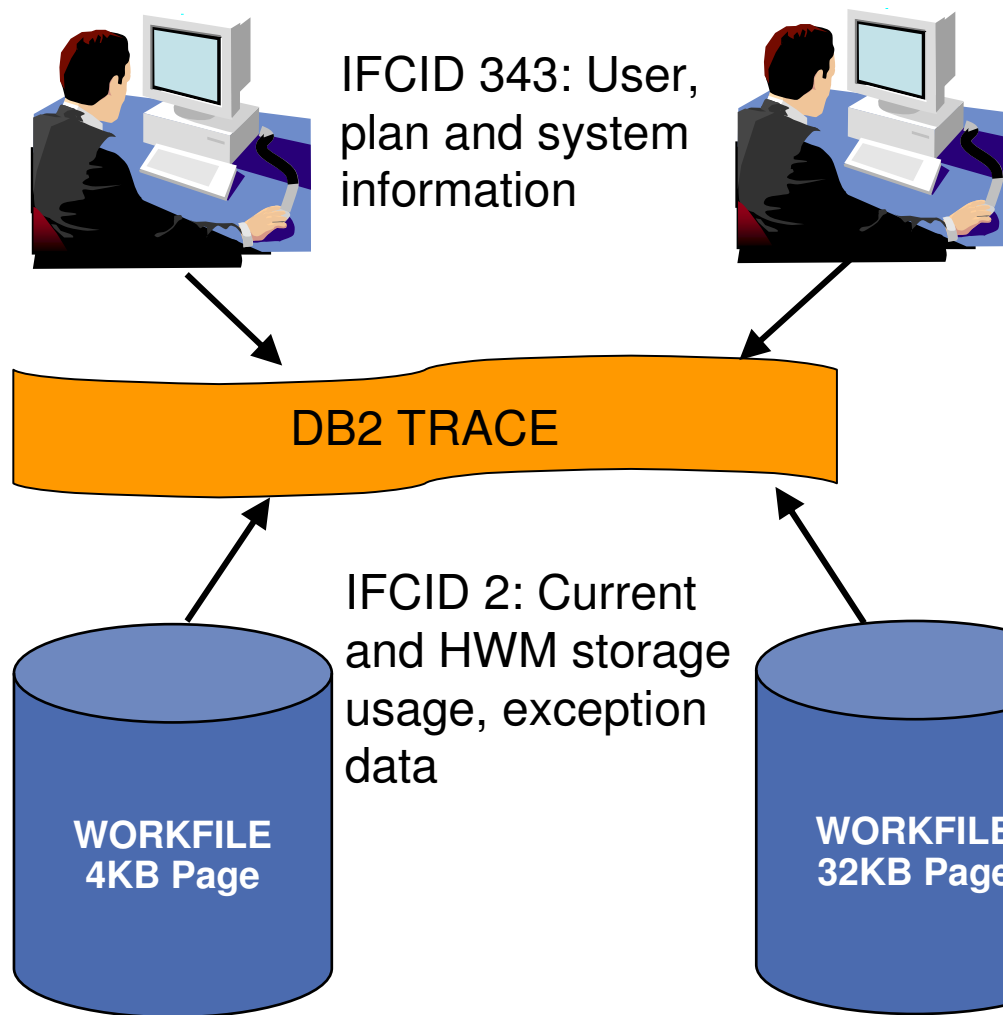
Controlling Temporary Space Utilization

- Control of temporary space utilization at the agent level
- New ZPARM: MAXTEMPS
 - Macro DSN6SYSP, panel DSNTIP9
- If MAXTEMPS is exceeded for any given agent:



```
SQLCODE = -904, ERROR:  
UNSUCCESSFUL EXECUTION CAUSED BY AN  
UNAVAILABLE RESOURCE.  
REASON 00C90305,  
TYPE OF RESOURCE 100, AND  
RESOURCE NAME = 'WORKFILE DATABASE'  
SQLSTATE = 57011
```

Monitoring Temporary Space Utilization



■ IFCID 2

- Updated with new counters
- Monitoring of temp space usage at the DB2 subsystem level, including the number of times MAXTEMPS has been exceeded

New IFCID 343 (Performance Trace Class 3)

- Written when the MAXTEMPS limit for an agent is exceeded

Separation DGTT and Sort Workfile

- APAR PK70060 to
 - For sort workfile, DB2 will first go after tablespaces with zero secondary quantity to find space and if none is available, will try to find space in the first 2GB region of tablespaces with non-zero secondary quantity and even if that is not possible, will use any available tablespace
 - For DGTT, DB2 will first try to find space in table space with non zero secondary quantity and if none is available will use any available table space
- To achieve separation for DGTT and Sort Workfile use, the recommendation will be to have some table spaces with zero secondary quantity and some with non-zero secondary quantity
- Maximum primary quantity of workfile tablespace is 2GB. But workfile tablespace can grow up to 64GB like any classic segmented table space if secondary quantity is not zero

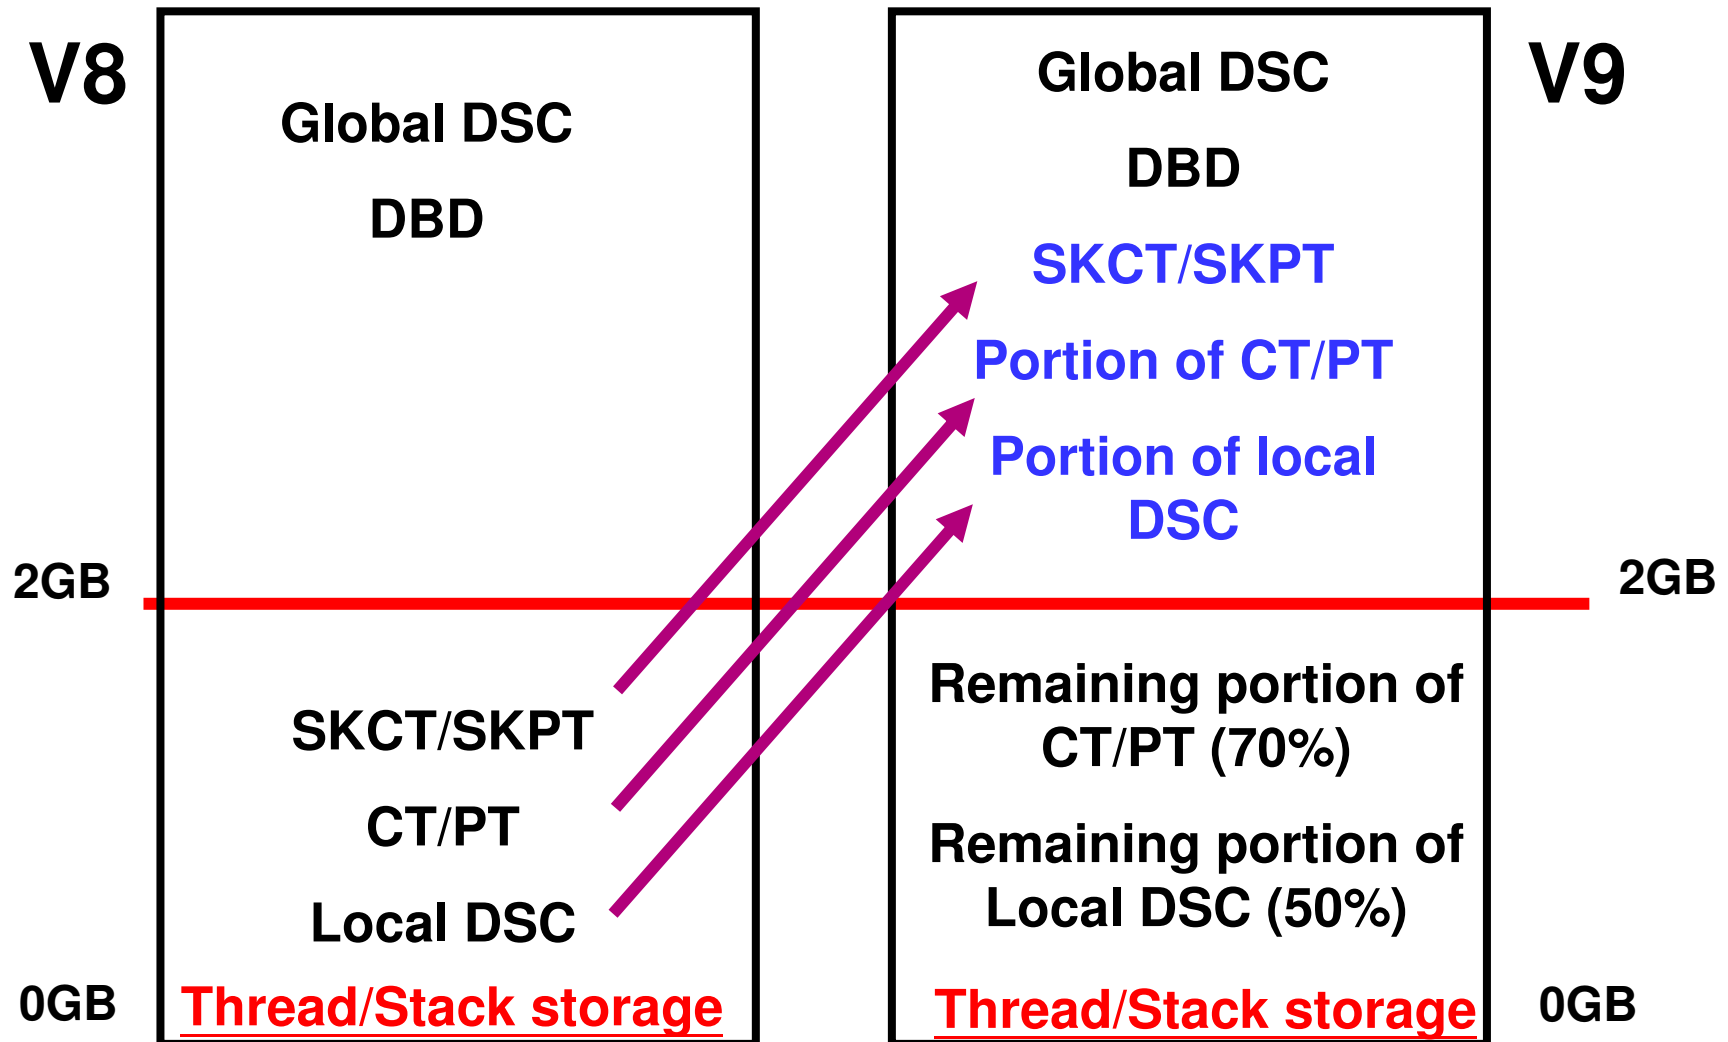
DB2 9 CPU Performance

- The target for DB2 9 CPU performance is to be roughly equivalent or marginally better relative to V8
 - Assumes z890, z990, z9 or z10
 - Assumes no access path regression
 - Assumes REBIND on the new release
- Mileage in terms of reduced CPU reduction will vary
- Customers running DB2 9 on old hardware (z800/z900) will definitely see CPU regression - may be 10%
- Data sharing customers running on DB2 9 (NFM) may see significant savings from reduced LC19 contention and less spin to get unique LRSN

DB2 9 CPU Performance ...

- Dynamic prefetch replaces all sequential prefetch in SQL calls (CM), except in tablespace scan
 - No rebind is required to switch to dynamic prefetch as this change is transparent to Optimizer
 - In V8 when pages are read in via dynamic prefetch they are marked as sequential and when they are subsequently touched by random getpage they are reclassified as random
 - In DB2 9 the pages are not re-classified following subsequent touch by random getpage

DBM1 Virtual Storage below 2GB



DBM1 Virtual Storage below 2GB ...


- DBM1 Virtual Storage Constraint Relief for static SQL users
 - EDM pool in DB2 9 – need REBIND
 - SKCT/SKPT moved above 2GB
 - A portion (close to 30%) of CT/PT moved above 2GB
 - Average estimated reduction of 60% but wide fluctuation from 20 to 90%
 - Only non-stealable components (CTs/PTs) are left in the EDM pool below 2GB
 - Plan on at least keeping $2 * 0.7 * \text{CT/PT}$ (V8 peak) for EDM Pool below 2GB
 - There are no stealable EDM components now below 2GB
 - Increase in package size after REBIND
 - Increased likelihood of hitting EDM Pool full condition
 - Intent of 2x factor is to provide some headroom
- DBM1 VSCR for dynamic SQL users
 - Local dynamic statement cache (KEEPDYNAMIC=YES)
 - Rough estimation of V9 = 50% of V8

DBM1 Virtual Storage below 2GB ...

- User thread storage, System thread storage, Stack storage
 - Current expectation of less than 10% difference overall
- Potential reduction can range from 0 to 300MB depending on thread/stack storage usage
 - Very few installations will see as much as a 200-300MB reduction!
 - Especially if you are an IMS/TM customer who has very large ECSA and a small Extended Private Region (i.e., 1GB)
 - No VSCR for dynamic SQL users unless they are using local dynamic statement cache as a result of using BIND option
KEEPDYNAMIC(YES)
 - VSCR from local dynamic statement cache savings will be small if the V8 size is heavily constrained by low MAXKEEPD

Partition By Growth Table Space

- Utilities can operate at the partition level
 - Except for the LOAD utility
- REORG PART may fail when
 - PCTFREE or FREEPAGE is non-zero or zero
 - A new ZPARM parameter to set PCTFREE=FREEPAGE=0
 - REORG TABLE to alleviate this problem
- Not recommend for tables with LOB
 - REORG is unable to move data rows across partition boundary
 - Use DSSIZE 64 GB to reduce the chance of hitting this problem
- Unable to use DSN1COPY to copy data for a table that has more than one partition



Partitioned Table Space (parts added as needed)

Basic Row Format and Reordered Row Format

```
CREATE TABLE TB1 (C1 INTEGER NOT NULL, C2 VARCHAR(10) NOT NULL, C3 CHAR(10) NOT NULL, C4 VARCHAR(20))
```

Basic Row Format

C1	C2	C3	C4
8000000A	0006 WILSON	ANDREW	0008 SAN JOSE

2-byte length

Reordered Row Format

C1	C3	O2	O4	C2	C4
8000000A	ANDREW	12	18	WILSON	SAN JOSE

offset to C2

offset to C4

- The offset is a hexadecimal value, indicating the offset for the column from the start of the data in the row.
- Note: in BRF, for a varying length field, DB2 logs from the first changed byte to the end of the row. In RRF, the first changed byte may be the offset field for the first varying length field following the one being updated.

Reordered Row Format (RRF)

- *Improved performance for tables that have variable length columns*
- *May reduce compression ratio for tables that have many small VARCHAR columns*
- *REORG/LOAD REPLACE to convert existing tables to RRF in NFM*
 - *Except if ROWFORMAT BRP is specified or if ZPARM RRF = DISABLE is set – PK87348*
 - *Can be done at PART level*

Utility CPU time reduction – primarily from index processing

- 5 to 20% in Recover index, Rebuild Index, Reorg Tablespace/Partition
- 5 to 30% in Load
- 20 to 60% in Check Index
- 35% in Load Partition
- 30 to 50% in Runstats Index
- 40 to 50% in Reorg Index
- Up to 70% in Load Replace Partition with NPIs and dummy input

Reorg Utility Performance

- Parallel unload/reload by partition
 - 10 to 40% faster in one measurement
- Eliminate Build2 phase in Online Reorg Partition with NPI (Non Partitioning Index) for better availability
 - But higher CPU time and elapsed time when few out of many partitions, especially with more NPIs, are Reorg'd as entire NPIs copied to shadow dataset
 - Additional temporary DASD space needed
 - NPIs are automatically Reorg'd also
- Fast Log Apply in Online Reorg

Index Compression

Difference between data and index compression

	Data	Index
Level of compression	Row	Page (1)
CPU overhead	In Acctg	In Acctg and/or DBM1 SRB
Comp in DASD	Yes	Yes
Comp in BP and Log	Yes	No
Comp Dictionary	Yes	No (2)
'Typical' Comp Ratio CR	10 to 90%	25 to 75% (3)

Large Index Page Size

- Pros and cons of large index page size (8K, 16K, 32K)
 - Reducing the number of index levels, reducing the number of getpages for index traversal, there by reducing CPU resource consumption
 - Reducing the number of index page splits which are painful especially for GBP-dependent index (data sharing)
 - Page size greater than 4K is an essential requirement for index compression
 - Note index compression is NOT a performance improvement
 - On the other hand large page size may possibly lead to either
 - Wasted space within page to maintain ability to compress down to 4K CI
 - Large page size can aggravate index buffer pool hit ratio for random access

What customers are saying about DB2 9



We are in great shape. **DB2 9 is very solid and stable.** Our reorgs are running great and **performance is the best that it has ever been.** We are a very happy customer!
-Mark Shackelford, VP Infrastructure Baldor Electric



"Like a good wine, DB2 has certainly gotten better with age!"
-Debra L Cook, ExxonMobil Global Services Company



"We needed a database that represented the future and DB2 9 is the future. DB2 9 compression capabilities are key in helping reduce the size of our databases — in one case by up to **83 percent.** This ultimately helps us **minimize storage costs and increase performance.**"
-Jean Holley, CIO, Tellabs, Inc.



"Not only is our IBM infrastructure **rock solid.** Our total cost of ownership is very low compared to other solutions. Our administrative framework for health education is **far superior in terms of cost compared to any other implementation** I've seen on campus. It took us a **fraction of the time** to develop than it would have on any other platform, and it's providing us with a reliable, fast solution for moving our health education program forward into the future."
- Steve Ware, Systems Coordinator, University of Florida Computing & Networking Services

What customers are saying about DB2 9



PETROL

"DB2 9 for z/OS will result in an **increased level of performance** with our data warehousing environment. We also look forward to using the new XML capabilities in DB2 9 to provide even **greater efficiency and consistency our XML applications.**"

- *Pavel Batista, IT Director, Petrol*



citi Citi never sleeps™

"The new security and compliance tracking capabilities of DB2 9 for z/OS will allow us to **simplify our existing process** for the management reporting of changes to our production database systems. In addition, we have evaluated the new **native XML support** that DB2 9 provides and view this as a **major technological advancement** to help us to integrate XML data into a format that will allow for more robust and efficient usage by applications."

- *Bob Perih, Senior Vice President, Citigroup*



SAP

"IBM and SAP have cooperated very closely on DB2 9 for z/OS and we look forward to supporting our customers with these new capabilities. The SAP Banking Services 5.0 release, which will be released soon, will be the first SAP product to take advantage of many new beneficial features in DB2 9 for z/OS."

- *Torsten Wittkugel, Vice President of Database and Operating System Platform Development, SAP*

THANK YOU