



DB2 9 for z/OS

Migration Planning and Experiences

Michael Dewert, DB2 for z/OS Development
dewert@de.ibm.com



INNOVATE.OPTIMIZE.
PERFORM.

DB2 9 for z/OS Roadshow
Paris – 6 April 2009

Disclaimer

© Copyright IBM Corporation 2009. 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



Objectives

- Share lessons learned, surprises, pitfalls
- Provide hints and tips
- Address some myths
- Provide additional planning information
- Information on new enhancements



Agenda

- Quick hits
 - Preparing for the migration
- Migration
 - Overview
 - Plan stability
 - Converged TEMP space
- What to expect?
 - DB2 9 CPU performance
 - DBM1 Virtual Storage relief below the 2GB bar
- More quick hits
 - New Functions



Quick Hits

- No need to fear migration with proper planning and testing
 - Follow step by step approach adopted by other successful customers
- Migrate only from DB2 V8 (NFM) with fallback SPE (PK11129)
 - See DB2 9 migration info APAR (II12423) for list of prereq APARs and PTFs
 - Make sure that the pre-conditioning APAR for Plan Stability (PK52522) is applied on all V8 (NFM) systems
- z/OS 1.8 required for
 - Full RACF support of trusted context and roles
 - Volume-level utilities enhancements
- Minimum required for DB2 Connect:
 - V9.1 FP1, V8.2 FP6, V8.1 FP13, V9.5
 - V9.5 FP3 introduces Sysplex Failover/Workload balancing and XA support for DB2 Connect Client
- No plan to allow V7 Precompiler function under DB2 9



Quick Hits ...

- Sound maintenance strategy is essential for all customers
 - Recommended to exploit CST/RSU process
 - Apply 2 to 3 preventative service drops annually
 - Exploit Enhanced HOLDDATA to be vigilant on HIPERs and PEs
 - No one-size-fits-all strategy
 - Review installation guide and the material supplied to ensure that RSU only service is installed
 - Can enforce installing RSU only service by adding the SOURCEID (RSU*) option in the supplied APPLY and ACCEPT jobs
 - Note '*' will pull ALL RSUs off of a particular tape



Quick Hits ...

- PDSE required for SDSNLOAD and SDSNLOD2
 - PDSE may only be shared by z/OS systems which are part of the same z/OS Sysplex
 - In other cases, a separate copy must be created for each z/OS system
 - See II14067 (z/OS 1.7) or II14255 (z/OS 1.8) for related maintenance
 - See APAR PK59069 for DB2 Precompiler
 - See APARs OA24410, UA35620 for PDSE
- Make PDSE address space restartable
 - Recommendation only applicable if datasets (e.g., SDSNLOAD) not in linklist
 - Allows user restart (refresh) of PDSE address space to circumvent hangs, abends and user errors
 - Must specify PDSE_RESTARTABLE_AS(YES) in IGDSMSxx parmlib member
 - If a PDSE issue is discovered with datasets then use the following command to refresh the SMSPDSE1 address space:
 - V SMS,PDSE1,RESTART
 - During IPL, SMSPDSE starts by default and SMSPDSE1 will start if PDSE_RESTARTABLE_AS(YES) is specified



Quick Hits ...

- Expand BSDS using V8 stand-alone utility DSNJCNVB
 - Or let DSNTIJUZ run the DSNJCNVB module during migration to DB2 9 (CM)
- Need to configure HVSHARE (64-bit shared private storage) in IEASYSxx parmlib at a minimum of 128GB per DB2 subsystem running on LPAR
 - Even if DDF is not used
 - Use DISPLAY VIRTSTOR,HVSHARE or D VS,HVSHARE to check



Quick Hits ...

- Need to reset advisory REORG-pending (AREO*) status on all table spaces and indexes before migrating to DB2 9 (CM)
- Rebind all packages containing SQL LOCK TABLE statements in DB2 9 CM
 - Proactive rebinds under V8 with the PTF UK14099 for APAR PK21103 will avoid the problem
- Rebind all plans and packages that have not been rebound since DB2 V3
- No need for data sharing group-wide shutdown after entry to DB2 9 (NFM) to switch to Locking Protocol 3 (LP3)
 - Apply APAR PK62027 PTF UK38906
 - Will automatically get the LOB locking improvements on entry to NFM
 - With this fix on there is no LP3 any more



Quick Hits ...

- DB2 9 did not eliminate DDF Private Protocol
 - Plan is to eliminate in DB2 9+1 release
 - If do not convert from Private to DRDA protocol, will not be able to migrate to DB2 9+1 release
 - DSNTDP2DP (Private to DRDA Protocol Catalog Analysis Tool) introduced with APAR PK27413 to assist the conversion
 - REXX program which looks at the DB2 Catalog
 - Generates CREATE ALIAS statements for remote locations that will probably need 2-part aliases
 - Generates the commands to convert packages and plans which have a remote location dependency (or can be detected to have))



Quick Hits ...

- DB2 9 did not eliminate DDF Private Protocol ...
 - There may be packages or plans remaining with DBPROTOCOL(PRIVATE) where there are no SQL statements within the respective package or plan which refer to remote object i.e., remote location dependency cannot be detected
 - For example
 - SQL statements in the Catalog which do not refer to any 3-part name directly
 - Aliases the SQL statements reference which do not
 - Possible reasons
 - Use of dynamic SQL
 - Package bound from a remote requestor/client
 - These packages will have to be rebound with DBPROTOCOL(DRDA) to flip the package/plan from private protocol access to DRDA protocol access



Quick Hits ...

- DDF functionality
 - See DB2 9 Info APAR II14203 for DDF/DRDA related maintenance
 - If an IP address is specified, it must only be done in one place:
 - Either in the BSDS (via DSNJU003) – New in DB2 9
 - DDF will listen on INADDR_ANY
 - Or, on the PORT statement of TCP/IP profile (via the BIND keyword)
 - DDF will listen on a specific IP address
 - BSDS (INADDR_ANY) & PORT (BINDSPECIFIC) are mutually exclusive
 - Once DDF binds to a specific IP address, it cannot listen on INADDR_ANY
 - SSL support:
 - DDF only supports secure ports when it is listening on the INADDR_ANY
 - If a secure port (SECPORT) is defined in the BSDS, no IP address can be specified on the PORT statement in the TCP/IP profile



Quick Hits ...

- Removed functions and incompatibilities
 - All stored procedures must be modified to become WLM enabled
 - The legacy JDBC driver will not work
 - In DB2 9, existing WLM environments configured to use the legacy driver will encounter a failure on address space initialization. This will happen in CM as well as NFM.
 - Data sharing groups with V8 in coexistence with DB2 9 (CM) will experience failure if a Java routine is invoked on any DB2 9 members where the WLM-SPAS JCL does not reference the Universal JDBC driver
 - Recommendation: Modify your WLM-SPAS JCL to use the Universal JDBC Driver while still on V8, prior to migration to DB2 9



Quick Hits ...

- Removed functions and incompatibilities ...
 - Simple table space creation support removed in DB2 9
 - New defaults:
 - Implicitly created TS: Segmented (CM) or UTS - Partition by Growth (NFM)
 - Explicitly created TS: Segmented
 - Aim at converting existing simple table spaces to segmented or UTS or partitioned
 - The DBPROTCL ZPARM is removed
 - The default distributed protocol for BIND will always be DRDA
 - The RELCURHL=NO ZPARM option is removed
 - Incompatibility for applications dependent on retaining page or row locks across commits for a WITH HOLD cursor

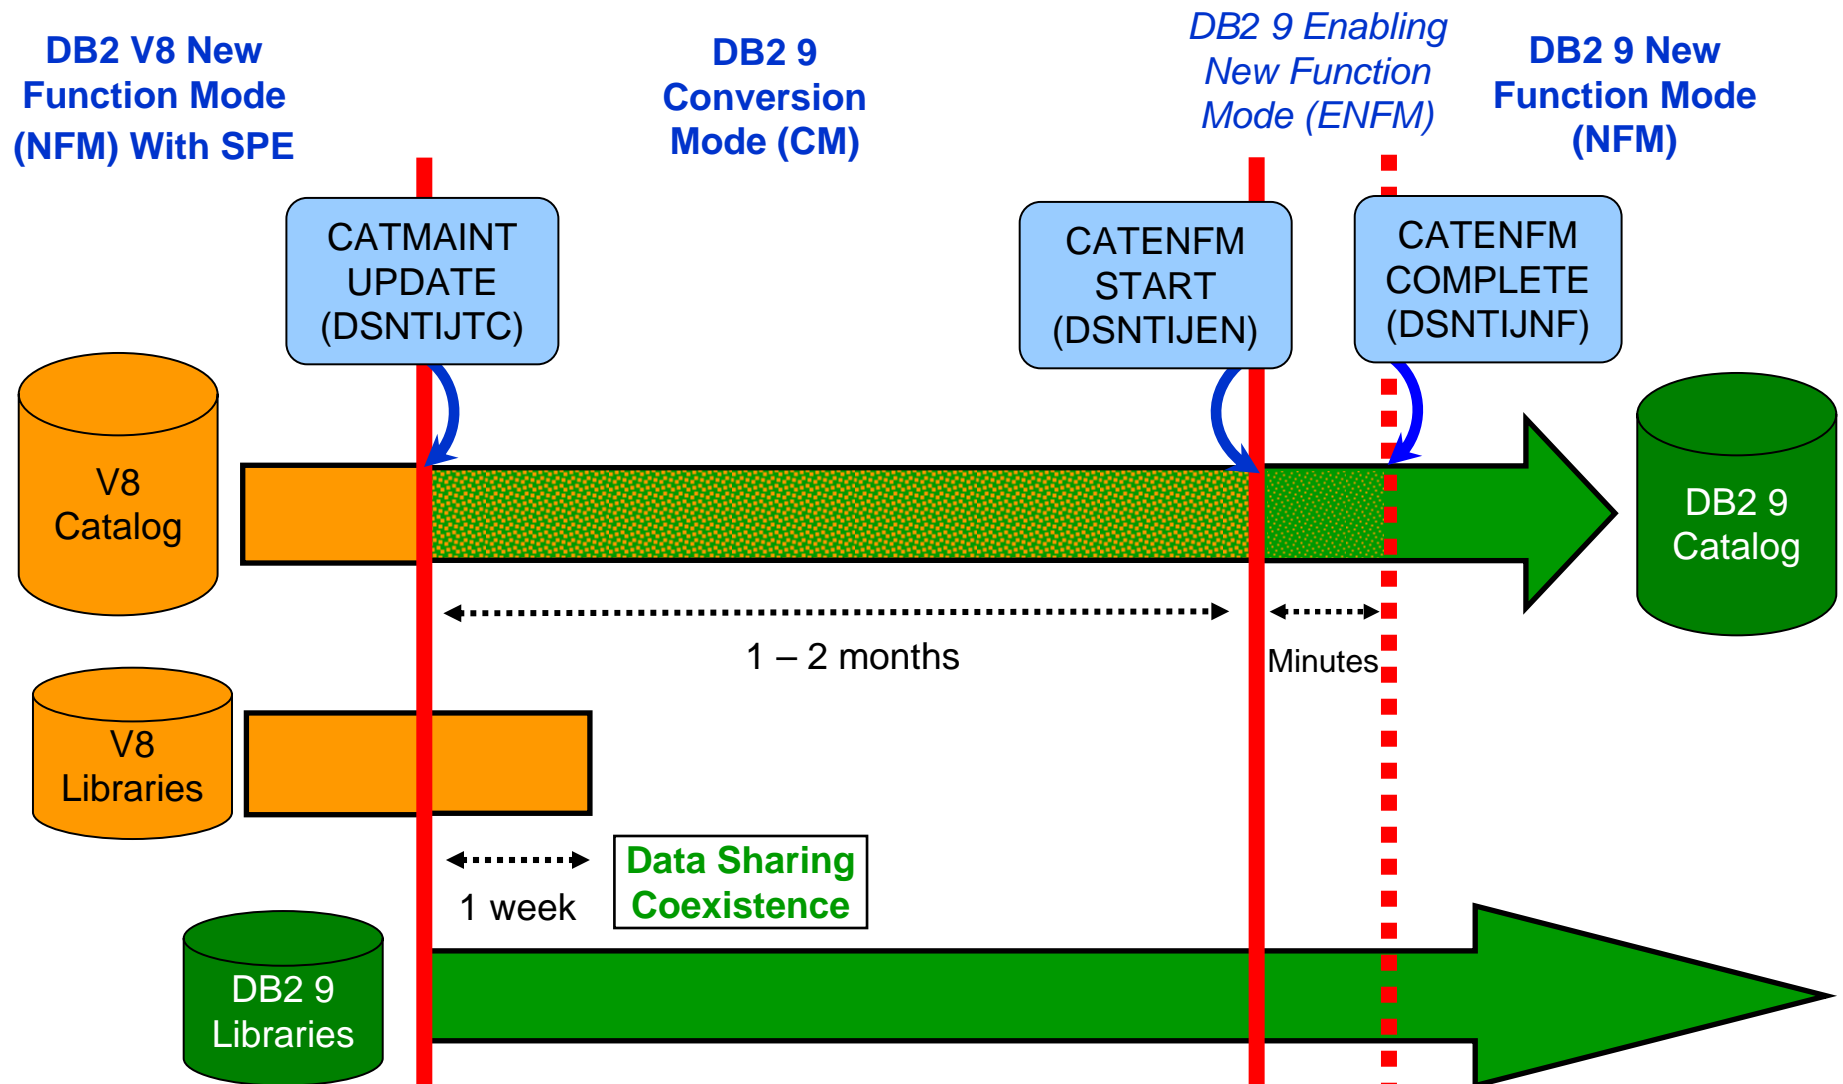


Quick Hits ...

- **Deprecated functions**
 - Plans containing DBRMs, ACQUIRE(ALLOCATE)
 - Consider switching to packages
 - Simple table spaces
 - Simple table space creation support removed in DB2 9
- **Changes to defaults**
 - Default setting for MGEXTSZ (secondary extent allocation) is YES
 - Changed by the installation CLIST from NO to YES
 - Default BIND options: ISOLATION(CS), CURRENDDATA(NO)
 - Not changed for REBIND, i.e. use the existing value



Migration Overview



Plan Stability

- New function of DB2 9 (PK52523)
 - 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



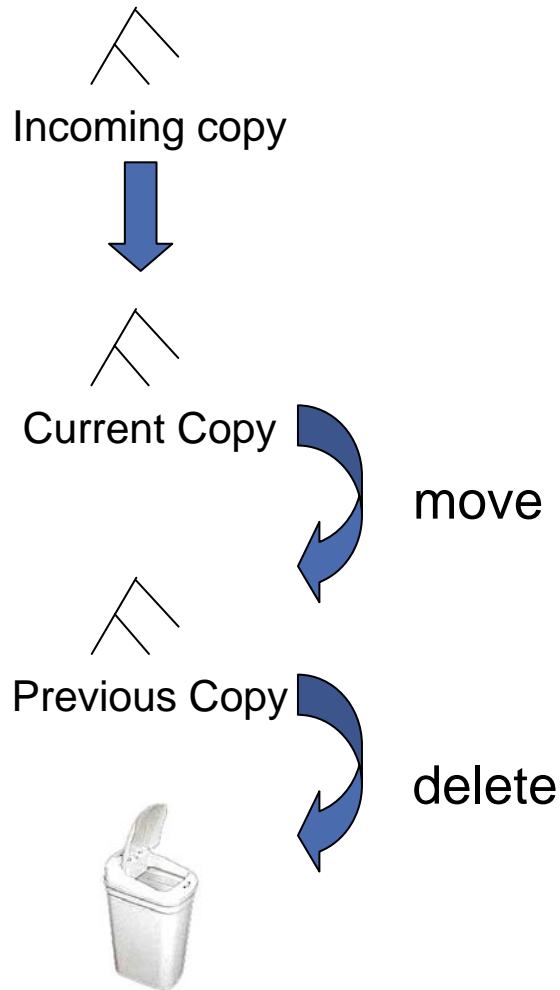
Plan Stability ...

- 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

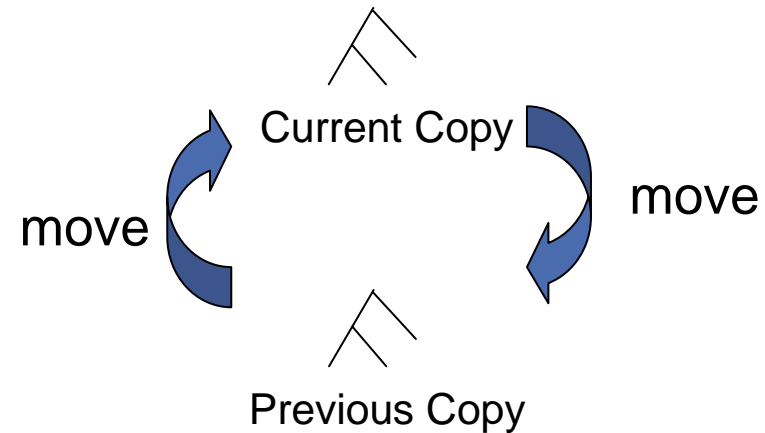


Plan Stability ...

REBIND ... PLANMGMT(BASIC)



REBIND ... SWITCH(PREVIOUS)



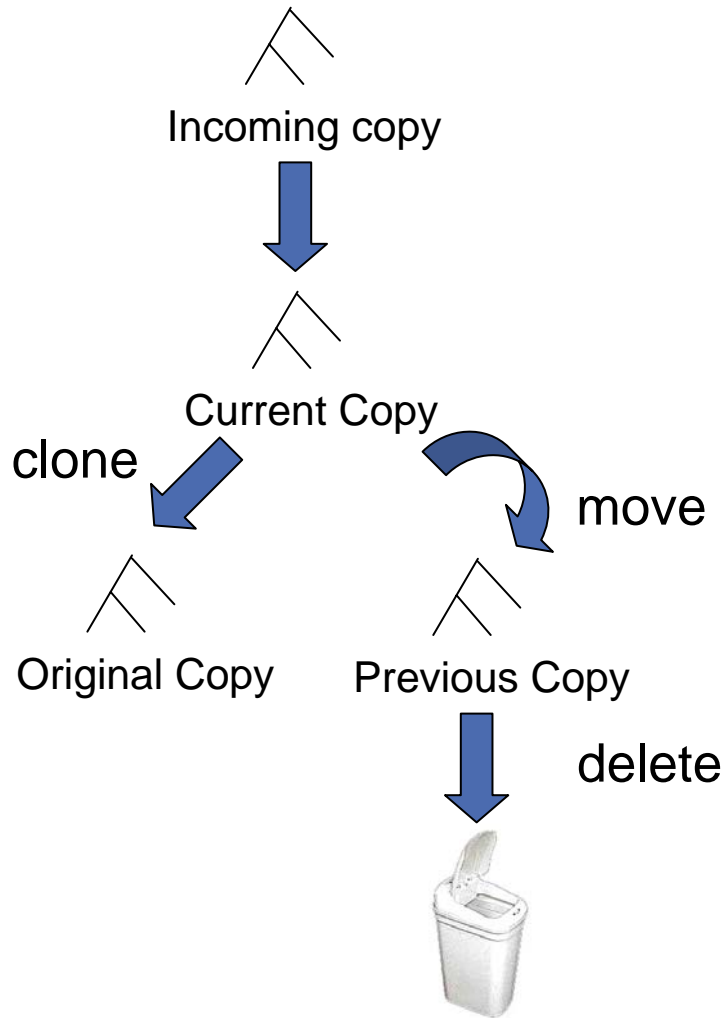
Similar processing for:

- EDM Sections
- SYSPACKDEP dependencies
- SYSPACKAGE record saved in EDM

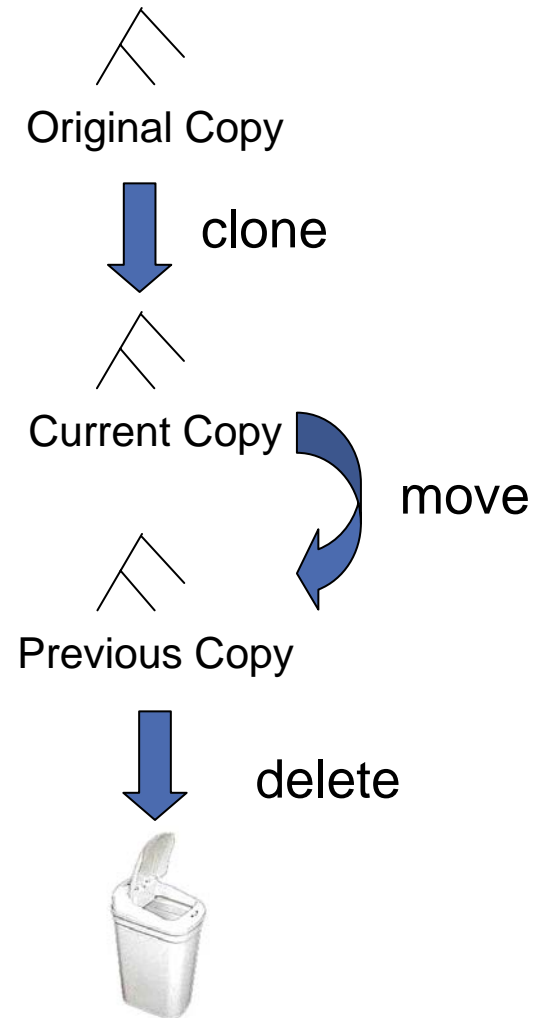


Plan Stability ...

REBIND ... PLANMGMT(EXTENDED)



REBIND ... SWITCH(ORIGINAL)



Sample Strategy for Migration using Plan Stability

- Migration strategy with Plan Stability
 - Before migrating to DB2 9 (CM), ensure V8 plan table information is available
 - On migration to DB2 9 (CM), set ZPARM PLANMGNT to EXTENDED
 - Objective: Make sure that the V8 version of package is kept as the original in case a fallback to DB2 V8 is required
 - Restriction: EXTENDED means that DB2 always keeps 3 versions of the package (even if they are the same). Watch out for SPT01 growth (limit is still 64GB with DB2 9)



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(RID) 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)
 - To avoid special DB2 handling to compensate for the different package structure in DB2 9 to avoid “puffing” of CTs/PTs/SLTEs
 - 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



Sample Strategy for Migration using Plan Stability ...

- REBIND in DB2 9 (CM)
 - DB2 now stores 3 versions of the package
 - Initial REBIND
 - Current version = DB2 9 version
 - Previous version = Original version = V8 version
 - If needed because of V9 access path regression, use REBIND PACKAGE ... SWITCH(PREVIOUS) to fallback to the V8 version
 - Subsequent REBINDs
 - Current version = “New” DB2 9 version
 - Previous version = Latest DB2 9 version
 - Original version = V8 version
 - Use REBIND PACKAGE ... SWITCH(PREVIOUS) to fallback to the previous V9 version

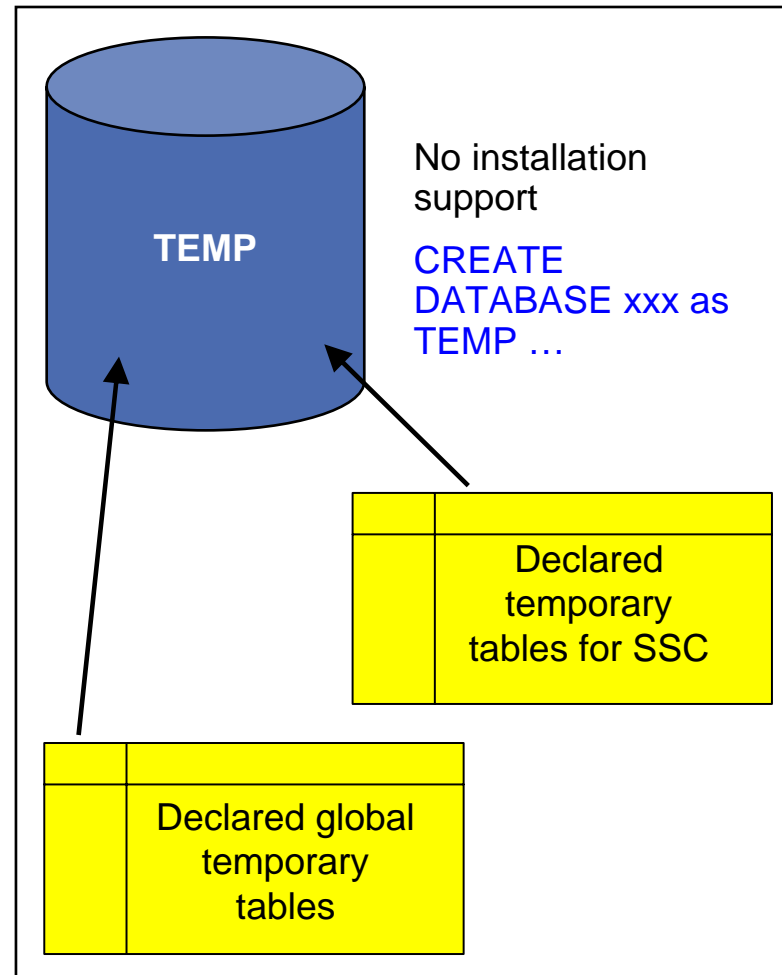
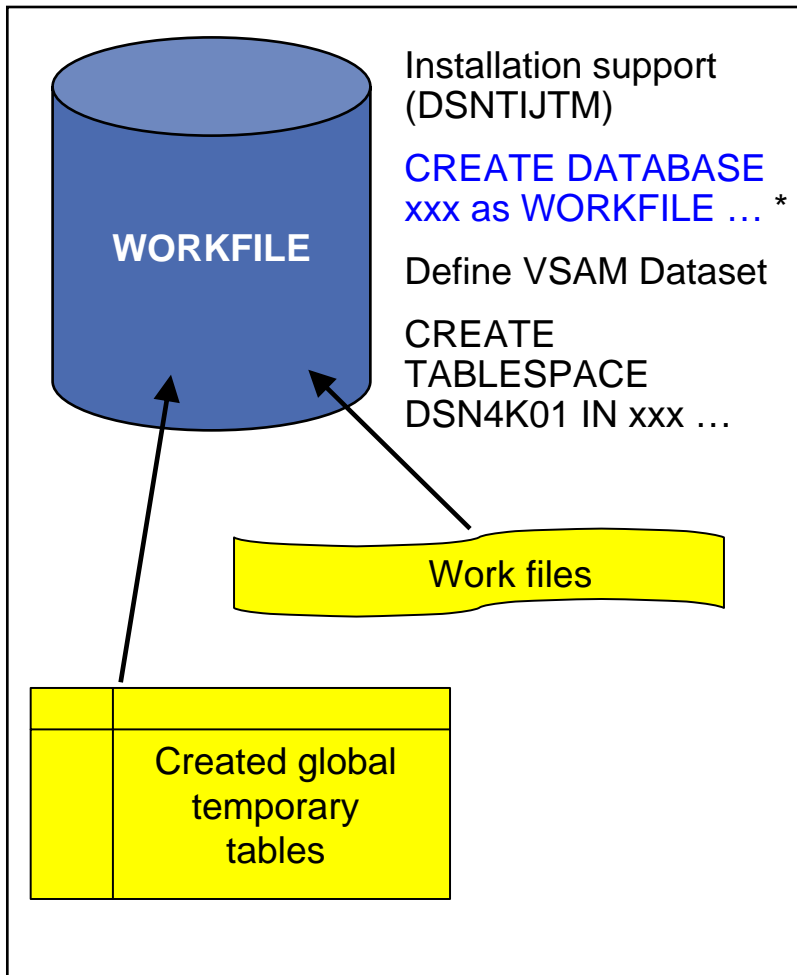


Sample Strategy for Migration using Plan Stability ...

- In case of fallback to DB2 V8, before falling back to V8
 - Use `REBIND PACKAGE ... SWITCH(ORIGINAL)` to fall back to original version of package (V8 version)
- In the future, to establish a new original version e.g to move forward to the version after DB2 9
 - Use `FREE PACKAGE ... PLANMGMTSCOPE(INACTIVE)` to free off the original (will also free off the previous version)



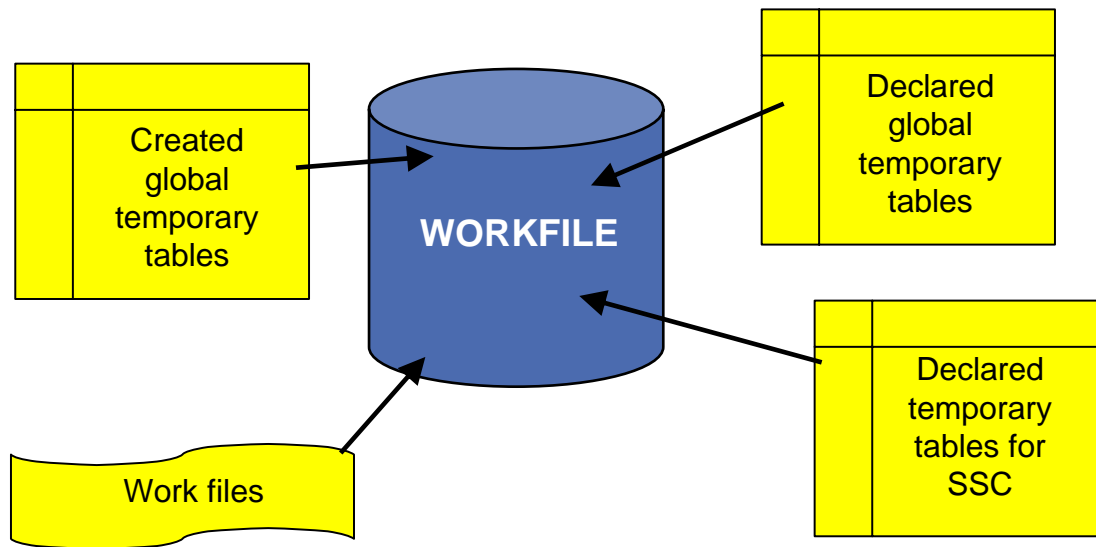
Temporary Space – The DB2 V8 Picture



* 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



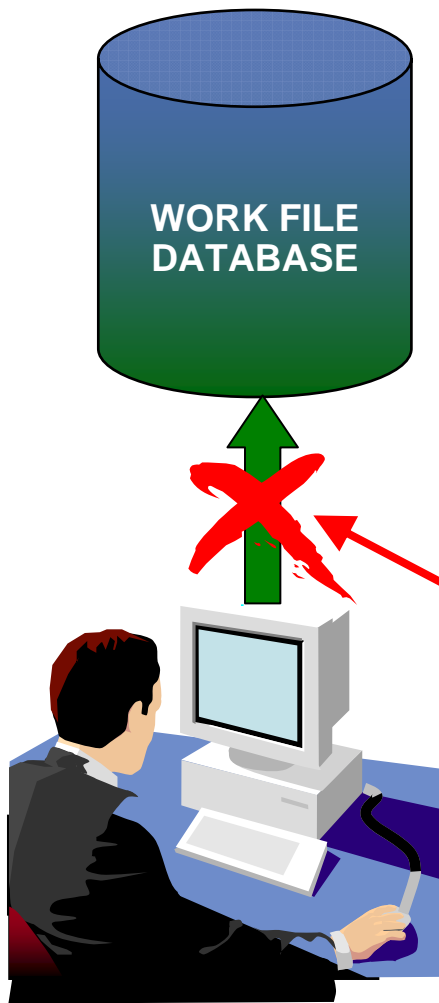
Planning For Converged TEMP Space

- 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

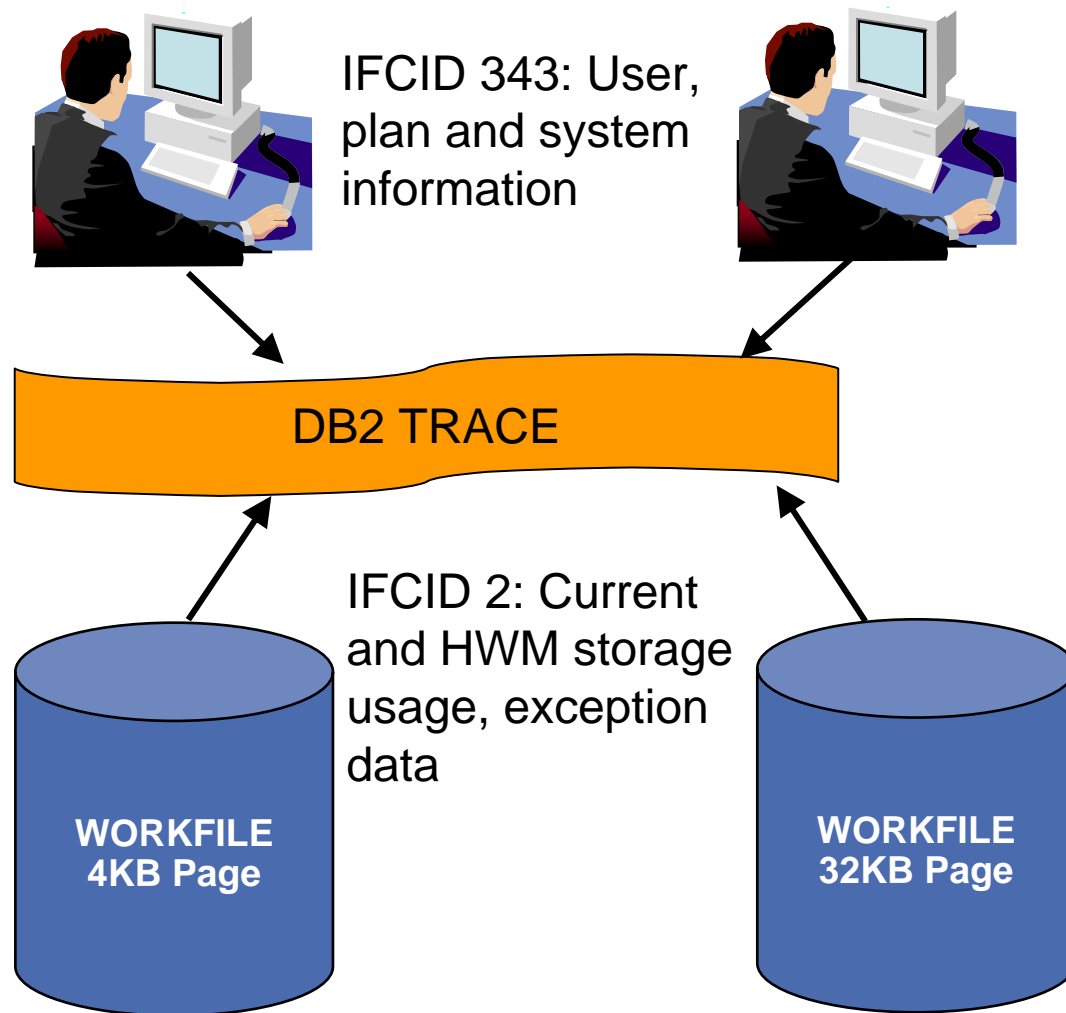
- Control of temporary space utilization at the agent level
- New ZPARM: MAXTEMPS
 - Macro DSN6SYSP, panel DSNTIP9
- If MAXTEMPS is exceeded for any given agent:



```
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



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
- Do not spend the benefits until you see them
- 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

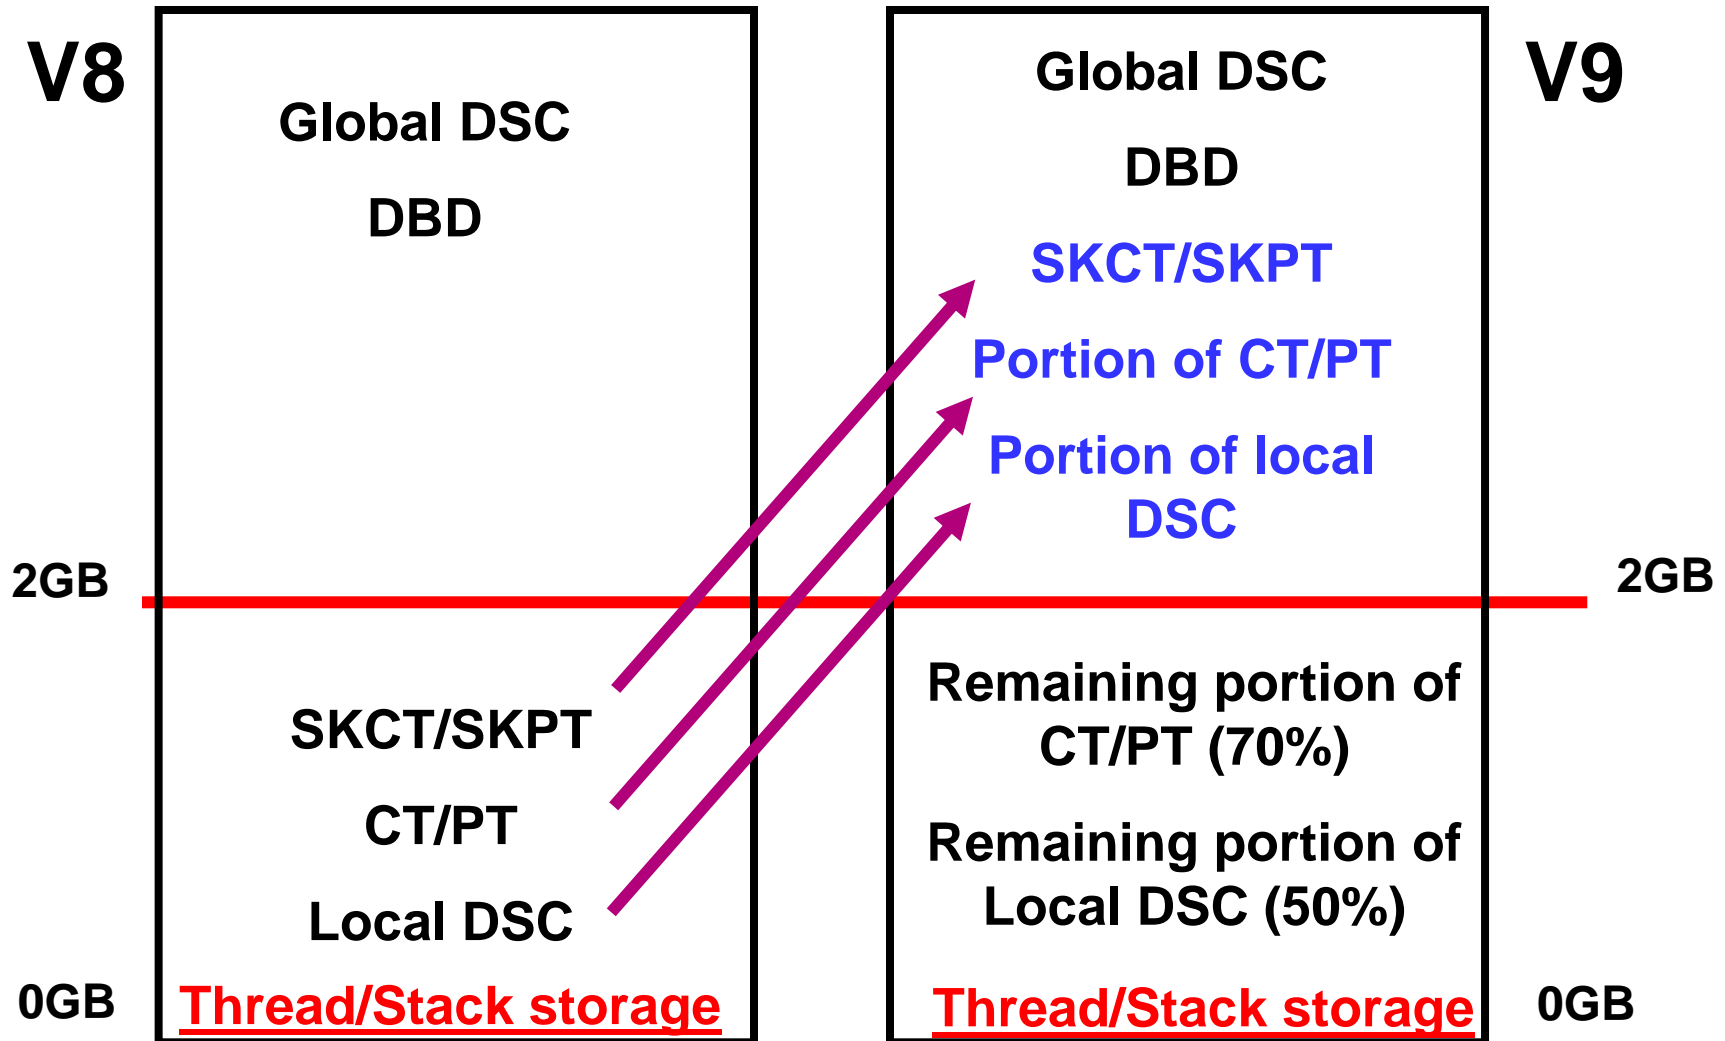


DB2 9 CPU Performance ...

- 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 KEEP DYNAMIC(YES)
 - VSCR from local dynamic statement cache savings will be small if the V8 size is heavily constrained by low MAXKEEPD



DBM1 Virtual Storage below 2GB ...

- V8 APAR PK20800 8/07
 - DISPLAY THREAD(*) SERVICE(STORAGE)
 - Produces 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 CT Long storage pool and the CT Very Long storage pool values. These reflects virtual storage consumption below the 2GB bar. They may be used to identify poorly behaved applications or DB2 code issues.



More Quick Hits

- Object level RECOVER from system level backups is available as soon as DB2 9 (CM), whereas object level RECOVER to PIT with consistency is not available until DB2 9 (NFM)
- Nice package of DB2 9 COPY enhancements
 - MRU to avoid throwing away useful random pages
 - Automatic implicit CHECKPAGE on table space pages and no setting of COPYPEND when damage found
 - All with less CPU



More Quick Hits ...

- Pros and cons of large index page size (8K, 16K, 32K)
 - Potential for reducing the number of index levels and there by reducing CPU resource consumption because the number of getpages for index traversal is reduced
 - 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
 - Aggravated index buffer pool hit ratio for random access



More Quick Hits ...

- Reordered Row Format (RRF)

- Potential for significant reduction in CPU resource consumption when accessing many rows with many varying length columns
 - Implements CPU tuning recommendation to place fixed length columns ahead of varying length columns
 - Provides for direct access to each varying length column
 - In some few cases may lead to increased logging volume

Prefix	Fixed Length Cols	Varchar Indicators	Varying Length Cols
--------	-------------------	--------------------	---------------------

- On by default in DB2 9 (NFM) and applies to all tablespace types
- A pageset will be automatically converted from Basic Row Format (BRF) to RRF when pageset is re-allocated (e.g., REORG, LOAD REPLACE)
- Application transparent – even for SELECT *
- REORG and LOAD REPLACE utilities override KEEPDICTIONARY during first time migration where data compression of variable-length rows



More Quick Hits ...

- Reordered Row Format (RRF) ...
 - Be careful when using DSN1COPY to export/import data during the transition period when all DB2 systems and data sharing groups have not yet been migrated to DB2 9 (NFM)
 - Should consider disabling RRF on a temporary basis before migrating first system or data sharing group to DB2 9 (NFM), and re-enabling after all the systems and data sharing groups have been migrated to DB2 9 (NFM)
 - How to disable RRF
 - Find &SPRMRRF parameter in the macro DSN6SPRC in SDSNMACS
 - Look for “&SPRMRRF SETC ‘1’ 1=Enable High Perf Row Option”
 - Edit the macro and set &SPRMRRF to ‘0’ and rerun DSNTIJUZ job or similar
 - Disabling RRF with SPRMRRF has the following effect:
 - Any tablespaces created in DB2 9 NFM will be created as BRF
 - When a BRF tablespace is reorganised it will remain in BRF
 - Note carefully pageset/partitions already in RRF will remain in RRF
 - REORG will not convert to BRF



More Quick Hits ...

- Reordered Row Format (RRF) ...
 - Important DB2 maintenance
 - APAR PK78958 disables RRF conversion for compressed pagesets
 - APAR PK78959 adds new DIAGNOSE options and will also disable RRF for CREATEs of compressed pagesets
 - APAR PK79127 stops inadvertent conversion to RRF during CM mode when support is not there



More Quick Hits ...

- Universal Table Space (UTS)
 - LOCKSIZE ROW (default) is not a recommendation
 - No MEMBER CLUSTER support
 - Hard prerequisite for CLONE TABLE
 - Migration path: UNLOAD/DROP/CREATE/LOAD
- V7/V8 combination of MEMBER CLUSTER+FREEPAGE 0+PCTFREE 0 is still an alternative to DB2 9 APPEND option of CREATE/ALTER TABLE
 - Advantage: Fills in holes left behind by bulk delete - no need for REORG to reclaim space
 - Only applies to partitioned table spaces
 - MEMBER CLUSTER is required



More Quick Hits ...

- Effectiveness of Asymmetric leaf page split function
 - Design point is to provide performance relief for classic sequential index key problem
 - Asymmetric split information is tracked in the actual pages that are inserted into, so it is effective across multiple threads across DB2 members
 - Prior to APAR PK62214, DB2 only remembered the last insert position and a counter
 - APAR PK62214 introduces changes to the tracking and detection logic, and it should work much better for data sharing
 - The new approach remembers an insert 'range' and tolerates entries being slightly out of order
 - It may still not be effective for large key sizes (hundreds of bytes), or if entries come in very bad order (i.e., they do not look sequential)
 - But for simple cases like 3, 2, 1, 6, 5, 4, 9, 8, 7, 12, 11, 10 ... DB2 will be able to determine that the inserted entries are ascending



More Quick Hits ...

- Enhanced index look-aside
 - In DB2 9 potential for extra usage of index look-aside during INSERT, DELETE, UPDATE processing
 - Applies to non-clustering indexes where CLUSTERRATIO is equal to or greater than 80%
- Cannot perform object level RECOVER from system level backup if dataset has moved away to a different volume since the backup was taken
 - For example
 - PIT recovery to before point of REORG with inline copy
 - Storage administrator decides to move datasets around to defrag volumes
 - Migrate to new DASD
 - Consider use of Recovery Expert Tool, or ISV tool which provides similar function



More Quick Hits ...

- BACKUP SYSTEM, RESTORE SYSTEM, object level RECOVERY from system level backup
 - If the production volumes are PPRC/XRC primary, the RECOVER utility will use the normal copy (i.e. not FlashCopy) when restoring data from a CopyPool backup
 - With z/OS 1.8 APAR OA23849, DFSMSHsm provides new options which allow FlashCopy to PPRC primary on FRBACKUP or FRRECOV
 - Once FlashCopy to PPRC primary is enabled via FRBACKUP PREPARE, then the BACKUP SYSTEM , RESTORE SYSTEM, and the RECOVER utilities can use FlashCopy to backup/restore data to/from DASD volumes that are PPRC primary
 - However, the PPRC duplex-pending condition will be set until the background copy has completed
 - GDPS Hiperswap will not tolerate PPRC volumes that are in duplex-pending state
 - So with APAR OA23849, there will still be an issue for customers that use GDPS/PPRC Hiperswap solution for high availability
 - For DB2 customers that use XRC, the RESTORE SYSTEM utility cannot use FlashCopy to restore the entire DB2 system from a CopyPool backup
 - But RESTORE SYSTEM can use a system level backup on tape
 - For XRC customers that wish to use FlashCopy to restore the entire DB2 system to a PIT, they need to disable XRC before running the RESTORE SYSTEM utility



More Quick Hits ...

- Optimistic Locking is not just for WebSphere!
 - Provides a build in Timestamp for each Row/Page
- Real Time Statistics (RTS) always enabled
 - In DB2 9 the RTS tables are now in the DB2 Catalog
 - In-memory statistics are always externalised to the RTS tables
 - Whereas in V8 the process had to be explicitly started
- MODIFY RECOVERY AGE n
 - Will set COPY PENDING for tablespaces which are defined with DEFINE NO and where SYSIBM.SYSTABLEPART.SPACE = -1
 - i.e., underlying VSAM pageset does not exist
 - Circumvention START DB(name) SPACENAM(name) ACCESS(FORCE)
 - Fixing APAR PK69427



More Quick Hits ...

- Declared Global Temporary Tables (DGTTs)
 - Up to 20x increase possible in both CPU resource consumption and elapsed time without fixing maintenance
 - APAR PK67301 to reduce excessive logging for sequences of insert and mass delete within the same commit scope
 - With more frequent commits, the performance degradation is less and PK67301 makes it even less
 - APAR PK62009 to reclaim the space emptied by mass delete after commit for ON COMMIT DELETE DGTT
 - However, if you have a loop of Inserts/Mass Delete without commit, then Insert performance will steadily deteriorate and that is where DELETE WHERE can help
 - APAR PK70347 to reduce DBM1 31-bit storage usage when Performance Trace Class 8 active



More Quick Hits ...

- Declared Global Temporary Tables (DGTTs) ...
 - 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



More Quick Hits ...

- SQL TRUNCATE
 - Efficient mass delete of data rows only with classic segmented or universal tablespace
 - No efficient mass delete of index entries
 - Performance is ugly with classic partitioned tablespace and/or many indexes
- REOPT(AUTO)
 - Only applies to dynamic SQL



More Quick Hits ...

- Do I still need QUIESCE points after DB2 9(NFM)?
 - PIT recovery will still be faster to a quiesce point assuming you can get one!
 - If you are using ISV tools which may not handle URs but do handle QUIESCE points
 - For example to allow a consistent PIT recovery to a different target
 - To force changed pages to DASD for whatever reason
 - Record a log RBA/LRSN in SYSCOPY for whatever reason



More Quick Hits ...

- Different Group Attach behavior when running multiple DB2 members of the same data sharing group on the same LPAR
 - DB2 in V8 would attach to first available member based on IEFSSNxx parmlib specification
 - Usermod available on V8 to change the behaviour and spread connections across available DB2 members on the same LPAR
 - DB2 in V9 will spread connections across available DB2 members on the same LPAR
 - Until now there was no way to re-instate the old V8 behaviour
 - New co-requisite APARs PK79228 and PK79327 will provide such an option
 - Enhancement to the ZPARM macro (DSN6GRP)
 - Can specify new option RANDOMATT=NO
 - Activate via –SET SYSPARM RELOAD
 - From that point forward, the DB2 member will be considered "ineligible" for Random Group Attach
 - DB2 member will only be selected if no other members are available
 - To completely restore V8 behaviour, must specify RANDOMATT=NO for every member of the data sharing group



More Quick Hits ...

- Native SQL Procedures

- Potential for significant reduction in CPU resource consumption by avoiding
 - Overhead of stored procedure invocation
 - Overhead of roundtrip between WLM and DBM1 address spaces for each SQL call
- Short running SQL procedure could achieve up to an 40% ITR improvement
- But little or no improvement for long-running SQL procedure
- zIIP-eligible if DRDA as it runs in DBM1, not WLM, address space under DDF enclave SRB
- Easy to code, develop and manage
- Conversion to native SQL procedures is easy



More Quick Hits ...

- Stored Procedures - Performance of different languages
 - Environment Configuration
 - z/OS 1.9
 - DB2 9 for z/OS
 - Universal Driver 3.52.76
 - JDK 1.4.2 (SQLJ/JDBC stored procedures)
 - 3 CP's
 - 2 zIIP's
 - 2 zAAP's
 - IRWW OLTP workload



More Quick Hits ...

- Stored Procedures - Performance of different languages ...

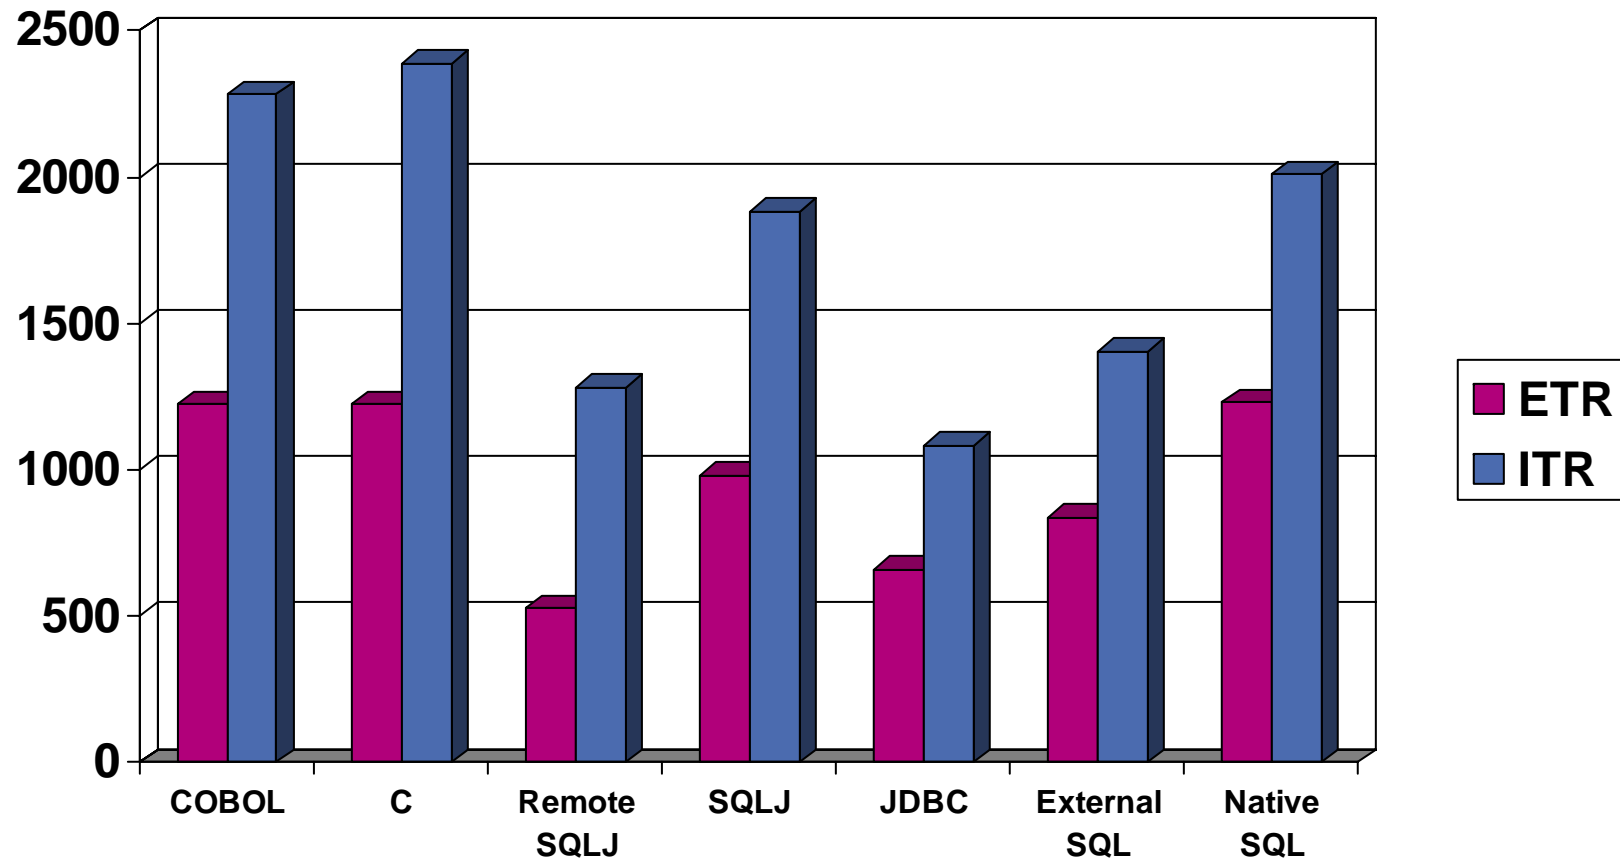
Language/API	Base CPU/tran Cost	Billable CPU/tran Cost after zIIP and/or zAAP redirect
COBOL stored proc	1X (BASE)	.88x
C stored proc	.95x	.83x
SQLJ stored proc	1.7x	1.15x (zIIP + zAAP)
JDBC stored proc	2.95x	1.76x (zIIP + zAAP)
External SQL stored proc	1.62x	1.49x
Native SQL stored proc	1.14x	.65x
Remote SQLJ	1.78x	1.06x

56



More Quick Hits ...

- Stored Procedures - Performance of different languages ...



More Quick Hits ...

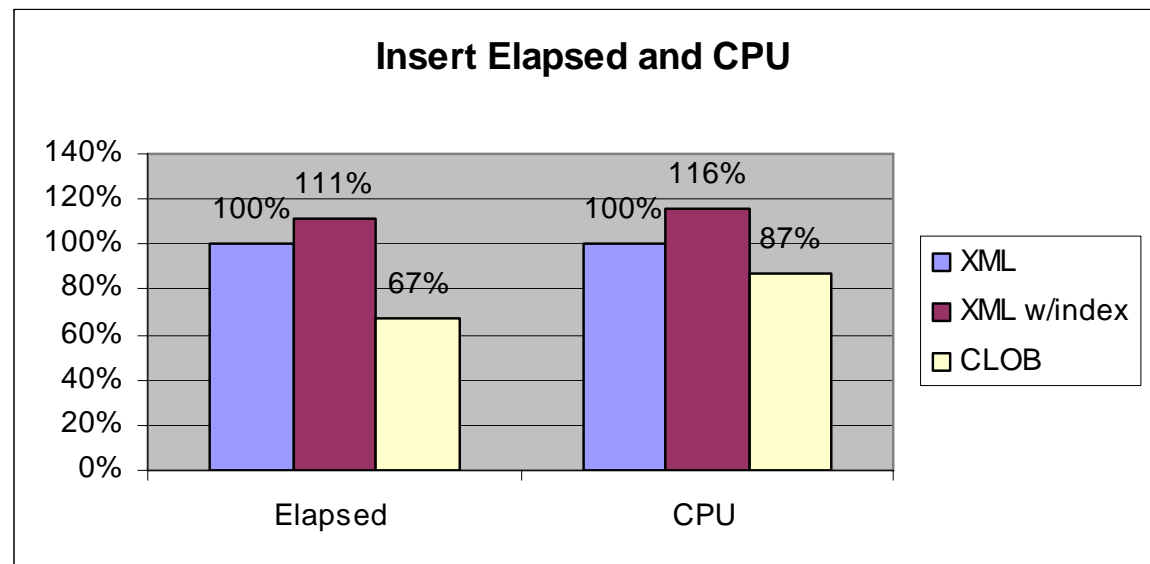
- Checklist for XML-related Configurations
 - Basic XML parsing requires z/OS XMLSS: z/OS 1.8 or z/OS 1.7 with APAR OA16303
 - XML schemas requires IBM 31-bit SDK for z/OS, Java 2 Technology Edition, V5 (5655-N98), SDK V1.5. And Java stored procedure setup
 - zparms for virtual storage:
 - XMLVALA and XMLVALS
 - Default: 200MB and 10GB
 - Also LOBVALA and LOBVALS impact bind-in and bind-out of XML
 - Buffer pool for XML tables (default BP16K0), requires authorisation for users who create or alter tables with XML columns



Some Guidelines to XML Usage

- The best table design is typically hybrid tables with relational and XML columns
- Heavily used columns, such as primary key and foreign keys, are better kept in regular columns
- XML columns are better kept to small docs if sub- documents are frequently accessed, not the whole large documents

- Insert cost of XML indexes is 15-20% per index in general
 - Define only necessary indexes (same as relational case)



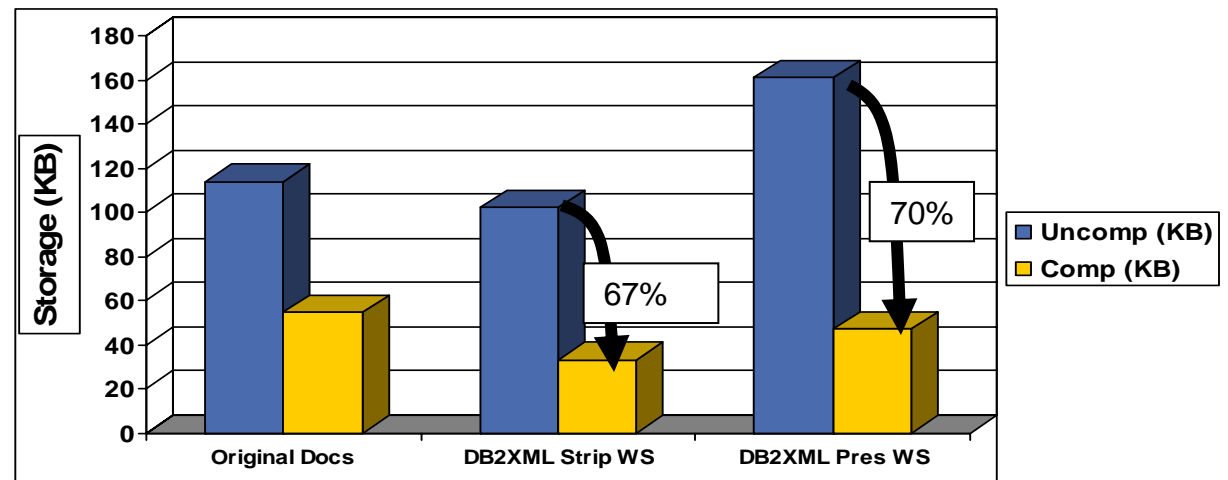
(average of 1K to 10M document insert performance)

59



Some Guidelines to XML Usage ...

- Cost of validation: 1.6 to 2.4x compared to non-validation
 - but highly depends on the schema complexity
- Compression – good ratio on XML tablespace, CPU impact is similar to relational



96 sample documents
Strip WS: Strip Whitespaces
Pres WS: Preserve Whitespaces

