

# **DB2 9 For z/OS FEATURES AND BENEFITS GUIDE FOR PEOPLESOFT ENTERPRISE APPLICATIONS**

---

<b>DB2 9 FOR Z/OS FEATURES AND BENEFITS GUIDE FOR PEOPLESOFT ENTERPRISE APPLICATIONS.....</b>	<b>1</b>
<i>Introduction .....</i>	<i>2</i>
<i>Description of Columns in Guide .....</i>	<i>2</i>
<b>1. RESOURCE REDUCTION .....</b>	<b>4</b>
1.1 THROUGHPUT .....	4
1.2 OPERATIONAL .....	10
1.3 MEMORY .....	11
1.4 DISK .....	12
<b>2. PRODUCTIVITY IMPROVEMENT.....</b>	<b>13</b>
2.1 OPERATIONAL .....	13
2.2 DEVELOPMENT .....	16
2.3 ADMINISTRATIVE.....	18
<b>3. SECURITY.....</b>	<b>20</b>
<b>4. DSNZPARM Changes.....</b>	<b>21</b>
<b>5. FUNCTIONS NO LONGER SUPPORTED.....</b>	<b>24</b>

## Introduction

This Guide consolidates information from various different Redbooks and Manuals and presents it in a format that will assist in undertaking a cost benefit analysis and feature/function relevancy checklist. It also highlight features of DB2 9 that are of specific benefit to PeopleSoft.

## Description of Columns in Guide

### Relevancy to PeopleSoft

- All of the features benefit some or all DB2 applications.
  - Features that are not relevant to PeopleSoft are noted as N/A.
  - Features with specific relevance to PeopleSoft are described as such. Any benchmark numbers quoted in this column will be documented in detail in a White Paper due out Summer 2008. It will be published in the "IBM and Oracle Resource Library" at <http://www-03.ibm.com/solutions/businesssolutions/oracle/doc/jsp/resource/all/index.jsp>
  - Features that may or may not have relevancy depending on your environment are left blank.

### Product plans

Any known or planned enhancements to PeopleSoft to support the feature.

PT8.48.12 (EBCDIC), PT8.48.14 (Unicode), and PT8.49.04 are currently certified with DB2 9.  
PT8.50 is planned to support DB2 V8 and 9 only, not DB2 V7

No statement of intent yet regarding any future PeopleTools Release pre-reqing DB2 9.

Use of new SQL programming constructs are at the developers discretion. When DB2 9 is pre-req'ed, or if the developer can code a Release specific codeline, then they are free to use new function. Certain features can be supported by using metaSQL in AE and PeopleCode that tests the DB2 Version. A MetaSQL to test the DB2 version is targeted for PT8.50.

### Area of Benefit

#### *Resource reduction (75 Features)*

- Throughput (62 features)
  - Benefits

- Concurrency -- reduces batch window, allows more batch/utility and online concurrency.
  - Elapsed time and/or cpu time– runs faster/uses less mips
- Categories:
  - Hardware/Optimizer/Locking/Tracing/MQ/Utilities/Data Sharing/Backup/Governor
- Operational (3 Features)
  - Benefits:
    - Speeds up recovery/restart
  - Categories:
    - Backup/SMS
- Memory (8 Features)
  - Benefits:
    - Reduces Virtual Storage and/or below the bar memory usage
  - Categories:
    - Memory
- Disk (2 Features)
  - Benefits:
    - Reduces disk space
  - Categories:
    - Disk

### *Productivity improvement (67 Features)*

- Operational (34 Features)
  - Benefits:
    - Makes running and recovery easier
    - Utility enhancements
    - Monitoring and tool enhancements
  - Categories:
    - Backup/Utilities/Intallation/Monitoring/Tools
- Development (20 Features)
  - Benefits:
    - Gets development done faster
  - Categories:
    - SQL/Tools/Misc
- Administrative (13 Features)
  - Benefits:
    - Gets DBA/System Programmer tasks done faster
  - Categories:
    - DDL/Tracing/Commands/Data Sharing/TCP

### *Security Improvements (4 Features)*

- Benefits:
  - Improves Security
- Categories:
  - Security

## References

TECH -	DB2 9 for z/OS Technical Overview – IBM Redbook: SG24-7330-00
PERF -	DB2 9 for z/OS Performance Topics – IBM Redbook: SG24-7473-00
SAP -	Enhancing SAP by using DB2 9 for z/OS – IBM Redbook: SG24-7239-00
INDEX COMPRESSION -	Index Compression with DB2 9 for z/OS – IBM Red Paper
LOBs -	LOBs with DB2 for z/OS: Stronger and Faster – IBM Redbook: SG24-7270-00
WHITE PAPER -	PeopleSoft Enterprise: DB2 9 features explored in a System z environment. (IBM White Paper: Available Summer 2008 at <a href="http://www-03.ibm.com/solutions/businesssolutions/oracle/doc/jsp/resource/all/index.jsp">http://www-03.ibm.com/solutions/businesssolutions/oracle/doc/jsp/resource/all/index.jsp</a> )

## Disclaimer:

This is not an officially approved Oracle document. It is written to the best of my ability based on my knowledge and some discussion with Oracle and IBM Support personnel. If I've got anything wrong, it's my fault.

## 1. RESOURCE REDUCTION

### 1.1 THROUGHPUT

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
1.1.1 Synergy with new I/O	<ul style="list-style-type: none"> <li>▪ A data set is spread across as many as 8 disks enabling faster prestaging.</li> <li>▪ FICON channels are progressively much faster than ESCON channels</li> <li>▪ z990 introduced FICON Express 2</li> <li>▪ z9 introduced FICON Express 4</li> <li>▪ DS8000 introduced faster device adapters and host adapters</li> <li>▪ MIDAW (Modified Indirect Data Address Word) has increased the channel efficiency</li> <li>▪ MIDAW requires z9 (2094) and z/OS1.6</li> </ul>			Hardware	Tech 2.2
1.1.2 Synergy with z990	<ul style="list-style-type: none"> <li>▪ Faster cpus</li> <li>▪ z990 (2084) <ul style="list-style-type: none"> <li>• More than 2 times faster row-level encryption</li> <li>• V9 long displacement instruction hardware support, simulated by microcode on z900</li> <li>• Most impact on input and output column processing</li> <li>• V9 CPU vs V8 on z900: +5 to 10%, more if many columns</li> <li>• V9 CPU vs V8 on z990 or later: 0 to -10% for column intensive</li> </ul> </li> </ul>			Hardware	Tech 2.2
1.1.3 Synergy with z9	<ul style="list-style-type: none"> <li>▪ z9 (2094) <ul style="list-style-type: none"> <li>• MIDAW to improve I/O performance</li> <li>• More engines (up to 54)</li> </ul> </li> </ul>			Hardware	Tech 2.2
1.1.4 System z9 Integrated Information Processor and specialty engines	<p>Remote SQL processing</p> <ul style="list-style-type: none"> <li>▪ DRDA network-connected applications over TCP/IP</li> <li>▪ Parallel queries</li> <li>▪ DB2 utility index processing</li> </ul>	Can offload 50-90% of PeopleSoft SQL to zIIP. Can offload 100% of PeopleSoft mainframe non-DB2 workload to z/Linux on IFL.		Hardware	Tech 2.3 SAP 10
1.1.5 Synergy with z10	<ul style="list-style-type: none"> <li>▪ Z10 (2097) <ul style="list-style-type: none"> <li>• 4.4Ghz quad core chip</li> <li>• 1.5TB of memory</li> <li>• 64 engines</li> </ul> </li> </ul>			Hardware	
1.1.6 CPU	When you first move to V9, CPU time generally decreases from 0% to 10% for z9 and z990 systems			Hardware	
1.1.7 Buffer manager enhancements	<ul style="list-style-type: none"> <li>• Larger preformatting quantity when formatting table spaces. Increased from 2 cys in V8 to 16 cys in V9</li> <li>• Long-term storage page fix on the I/O work area that is used for compressed indexes and castout engine work areas.</li> <li>• Increase in number of open and closed service tasks from 20 to 40.</li> </ul>			Optimizer, Data Manager	Perf 4.13 Perf 4.17
1.1.8 Prefetch enhancements	<ul style="list-style-type: none"> <li>▪ Replace all sequential prefetch, except in tablespace scan, with dynamic prefetch in SQL calls. Dynamic prefetch is more intelligent and more robust.</li> <li>• Larger prefetch quantity and deferred write quantity are used in V9 for large buffer pools, which are defined as follows:</li> </ul>	25% reduction in prefetch requests doing a NMIS against PS_PAY_EARNINGS		Optimizer, Data Manager	Perf 2.2 Perf 5.8 White Paper

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
	<ul style="list-style-type: none"> <li>For sequential prefetch, if VPSEQT*VPSIZE&gt; 160 MB for SQL then 64 pages; 320 MB for utility then 128 pages</li> <li>For deferred write, if VPSIZE&gt; 160 MB for SQL, 320 MB for utility</li> </ul>				
<b>1.1.9 WORKFILE database enhancements</b>	The WORKFILE database and the TEMP database are converged into the WORKFILE database. The WORKFILE database has been optimized to select the best page size when using a workfile table space. If the workfile record length is less than 100 bytes, a table space with a page size of 4 KB is used. Otherwise, a table space with a page size of 32 KB is used. Work file access is often sequential. Therefore, using a larger page size can be more efficient.			Optimizer, Data Manager	Perf 5.9 Tech 4.14 SAP 13.1 SAP Appx D
<b>1.1.10 In-memory work file for small sorts</b>	In-memory workfile support is provided when the final sort output data that would have been stored in a work file is less than the 4 KB or 32 KB pagesize of the selected work file. This means that small sorts are not written out to the work file, but the results are presented directly from the WORKFILE 4 KB or 32 KB buffer.			Optimizer, Data Manager	Tech 13.16.2 Perf 4.14
<b>1.1.11 Large record sort using 32K workfile page</b>	A 32 KB buffer pool for larger sort record sizes may be used to gain improved performance. Using a larger page size reduces the number of I/Os. It is now possible to control temporary space utilization at the agent level. A new DSNZPARM, MAXTEMPS, is added to DSN6SYSP to specify the maximum amount of space that can be used by a single agent at any single time.			Optimizer, Data Manager	Perf 5.9 Tech 13.16.1 This doc 4.3.3
<b>1.1.12 Increase the number of user defined indexes on the catalog</b>	Allows 500 user defined indexes against the catalog			Optimizer, Data Manager	Tech 3.10.2
<b>1.1.13 Greater than 4 KB page size for indexes</b>	The size of an index page limits the number of index keys that the index page can accommodate and can cause contention in indexes that split frequently. V9 lifts the 4K page size restriction by offering expanded index page sizes of 8 KB, 16 KB, and 32 KB.	Number of pages needed reduced by 50% for 8K vs 4K indexes. Pages added through page splits reduced by up to 50% for large inserts.	Customer specific option. Use PeopleSoft Tablespace DDL Automation tool to generate default index buffer pool, or ALTER INDEX	Optimizer, Data Manager	Perf 5.4 Tech 3.10.6 White Paper
<b>1.1.14 Asymmetrical split of index pages</b>	This enhancement has the potential to allow index sizes to be reduced and allow for more densely packed indexes	50% fewer leaf pages added to some PS_PAY_EARNINGS indexes when inserting rows		Optimizer, Data Manager	Perf 5.4 Tech 3.10.7 White Paper
<b>1.1.15 Random indexes</b>	A randomized key order allows DB2 to spread out index keys within the whole index tree, instead of maintaining an ascending or descending order, thereby minimizing index page contention and turning a hot index page into a cool index page. Cannot be used for sequential access			Optimizer, Data Manager	Perf 5.4 Tech 13.15
<b>1.1.16 Index on expression</b>	Index on expression allows you to create an index on a general expression	Useful for case-insensitive searches that use UPPER function such as NAME. PMR52007 currently open – APAR to open about 6/10/08.	Custom index required. Potential for inclusion as delivered index(es)	Optimizer, Data Manager	Perf 2.16 Tech 7.2 SAP 4.2 White Paper

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.1.17 Index look-aside</b>	The objective of the index look-aside technique is to minimize the number of getpage operations that are generated when an individual SQL statement or DB2 process is executed repeatedly and makes reference to the same or nearby pages. Index look-aside was introduced in DB2 V4 for SELECT. In DB2 V8, index look-aside was added for clustering an index during INSERT. In DB2 9, it is possible for more indexes to use the index look-aside function for DELETE			Optimizer, Data Manager	Perf 4.16 SAP 4.3
<b>1.1.18 Last used index</b>	SYSIBM.SYSINDEXSTATS provides information about the last usage of an Index.	Can be used to identify unused indexes, although unused delivered indexes are not officially recommended for deletion		Optimizer, Data Manager Disk space	SAP 5.2 White Paper
<b>1.1.19 REOPT AUTO based on parameter marker change</b>	If you specify the new REOPT(AUTO) bind option, DB2 automatically determines if a new access path is required to further optimize the performance of a statement for each execution	Has the potential to improve access paths due to widespread use of parameter markers (bind variables). APAR PK27981 DB2 V8 PTF UK33370 DB2 9 UK33371 required. REOPT (AUTO) fails in Payroll. PMR opened as of 5/2008.		Optimizer, Data Manager	Perf 2.12 Tech 13.3 White Paper
<b>1.1.20 Histogram statistics</b>	RUNSTATS normally collects frequency statistics for single-column or single multi-column data. Because catalog space and bind time performance concerns make the collection of these types of statistics on every distinct value found in the target column or columns very impractical, such frequency statistics are commonly collected only on the most frequent or least frequent, and therefore most biased, values. Histogram statistics is a way of summarizing data distribution. This technique divides up the range of possible values in a data set into intervals, such that each interval contains approximately the same percentage of the values.	Improves access paths in some circumstances for LIKEs and BETWEENs		Optimizer, Data Manager	Per 2.17 Perf 6.3 Tech 13.5 SAP 11.4 White Paper
<b>1.1.21 Reordered row format</b>	You can specify the column order however you wish and DB2 automatically reorders the columns within the row and places all the variable length columns at the end of the physical row within the data page.	App Designer already complies with this		Optimizer, Data Manager	Perf 4.12 Tech 3.7 SAP 11.6
<b>1.1.22 APPEND</b>	The APPEND processing option instructs DB2 to ignore clustering during SQL insert and online LOAD processing. Rather than attempting to insert rows in cluster-preserving order, rows are appended at the end of the table or appropriate partition	10% drop in cpu time when doing many inserts on PS_PAY_EARNINGS	Customer specific mod. Use ALTER TABLE.	Optimizer, Data Manager	Tech 4.4 SAP 12.3 White Paper
<b>1.1.23 Optimization of complex query</b>	Enhancements to the page range screening (or limited partition scan) function, which result in fewer partitions being accessed: <ul style="list-style-type: none"> <li>▪ Join predicates</li> <li>▪ Non-matching predicates</li> </ul>			Optimizer, Data Manager	Perf 2.4 Tech 13.8
<b>1.1.24 Global query optimization</b>	When a subquery does not meet the conditions for transformation into a join, the DB2 optimizer might instead transform the subquery, for more efficient processing, by correlating or de-correlating the queries.	7 sample subqueries used for testing against DB2 V5,6,7,8 now all produce correct access path in DB2 9		Optimizer, Data Manager	Perf 2.3 Tech 13.9 SAP 11.1 SAP 16.3.2 White Paper

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.1.25 Generalized sparse index and in-memory data cache</b>	In-memory data caching is extended to joins other than star join. In theory, all tables, which lack an appropriate index or enough statistics could benefit from sparse index / in-memory data caching: <ul style="list-style-type: none"> <li>▪ Temporary tables</li> <li>▪ Table expressions</li> <li>▪ Materialized views</li> </ul>			Optimizer, Data Manager	Perf 2.9 Tech 13.10 SAP 4.4 SAP 16.4
<b>1.1.26 Dynamic index ANDING for star join queries</b>	This enhancement introduces a different join approach within a star join group, called Pair-Wise Join. It requires a one-column index on Fact table to support each dimension table join columns. The proposed join method relies heavily on the resource of Rid Pool.	N/A		Optimizer, Data Manager	Perf 2.1 Tech 13.11 SAP 16.1
<b>1.1.27 Large Object (LOB/XML) flow optimization</b>	LOB (and XML) data retrieval has been enhanced so that it is more effective for small and medium size objects, and still efficient in its use of locators to retrieve large amounts of data.	N/A		Optimizer, Data Manager	Perf 3.3 LOB 4.3 Tech 13.12
<b>1.1.28 Group By sort improvement</b>	Sort enhancement for both DISTINCT and GROUP BY with no column function			Optimizer, Data Manager	Tech 13.16.4
<b>1.1.29 Distinct sort improvement when only non-unique index available</b>	DISTINCT sort avoidance with a non-unique index	This should improve some materialized "DISTINCT VIEW" access paths		Optimizer, Data Manager	Perf 2.1 Tech 13.16.5 White Paper
<b>1.1.30 Multi-row fetch for USS</b>	Enables block fetching of rows	AE processes under USS will benefit. Requires DB2 V8 APAR PK15288	Enhancement is targeted for PT8.50	Optimizer, Data Manager	
<b>1.1.31 Improved parallel query optimization</b>	With DB2 V8 the DB2 optimizer chooses the lowest cost "sequential" plan, and then determines how to "parallelize" the access path. With V9 the optimizer considers multiple sequential plans for parallelism and then executes the one with the lowest costs.	Makes parallelism more attractive, especially because of zIIP offload capability		Optimizer, Data Manager	SAP 16.3.1
<b>1.1.32 Package performance</b>	The overhead of executing packages with a single or a few short running SQL statements is an increase of CPU time when compared to DB2 V8	N/A		Optimizer, Data Manager (-)	Perf 5.12
<b>1.1.33 Open dataset ahead of use via command</b>	Command to pre-open or prime a data set ahead of first thread access. -ACCESS DB(dbname) SPACE(spacename) PART(n) MODE(OPEN)			Optimizer, Data Manager	Perf 8.13 Tech 5.9
<b>1.1.34 Latch class contention relief</b>	Addresses a number of performance issues with the high usage of certain latches and the underlying problem that was causing them.			Lock	Perf 4.10 Tech 3.5
<b>1.1.35 Skip locked rows</b>	The SKIP LOCKED DATA option allows a transaction to skip rows that are incompatibly locked by other transactions. Not the same as UR			Lock	Tech 7.6
<b>1.1.36 Optimistic concurrency control and update detection</b>	Optimistic concurrency control, also known as optimistic locking, represents a faster, more scalable locking alternative to database locking for concurrent data access. It minimizes the time for which a given resource is unavailable for use by other transactions. When an application uses optimistic concurrency control, locks are obtained immediately before a read operation and released immediately, as opposed to when cursor moves to next row or page. Update locks are obtained immediately before an update operation and held until the end of the transaction. Optimistic concurrency control uses the RID and a row change token to test whether data has been changed by another transaction since the last read operation.	N/A		Lock	Perf 5.13 Tech 7.3

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.1.37 Accounting trace overhead</b>	Equivalent or reduced CPU time for accounting trace classes versus DB2 V8.			Trace	Perf 4.11
<b>1.1.38 Improve the operations of Audit Trace function</b>	Adds options to reduce the volume and cost of AUDIT trace.			Trace	Tech 3.10.5
<b>1.1.39 MQ Messaging Interfaces user-defined function</b>	The MQ Application Messaging Interface (AMI) is no longer available and is replaced by the new MQ Message Queue Interface (MQI).			MQ	Perf 7.2
<b>1.1.40 Online CHECK LOB</b>	CHECK LOB has been enhanced to run in a SHRLEVEL CHANGE mode, in addition to the default SHRLEVEL REFERENCE			Utilities	Perf 6.7 LOB 6.12 Tech 11.12 SAP 8.7
<b>1.1.41 CHECK INDEX SHRLEVEL REFERENCE parallelism</b>	When more than one index is specified in the CHECK INDEX SHRLEVEL REFERENCE option, CHECK INDEX checks the indexes in parallel unless constrained by lack of storage or sort work files.			Utilities	Perf 6.1.1 Tech 13.7 SAP 8.7.3
<b>1.1.42 REPAIR LOCATE SHRLEVEL CHANGE</b>	The REPAIR utility has been enhanced so that LOCATE can be run against indexes, index spaces and table spaces with SHRLEVEL CHANGE. This does not apply to LOB table spaces.			Utilities	Tech 11.13
<b>1.1.43 Online REBUILD INDEX</b>	REBUILD INDEX can be run SHRLEVEL CHANGE			Utilities	Perf 6.5 Tech 11.15 SAP 8.7.4
<b>1.1.44 Image Copy, RECOVER INDEX</b>	10 to 20% cpu reduction . Reduced CPU overhead for COPY TABLESPACE with CHECKPAGE to be almost negligible.			Utilities	Perf 6.1 Tech 13.6
<b>1.1.45 CHECK INDEX performance</b>	20 to 60% cpu reduction for CHECK INDEX 36% reduction in both CPU time and elapsed time when performing CHECK INDEX with default options on the partitioning index			Utilities	Perf 6.1.1 Tech 13.6
<b>1.1.46 LOAD performance</b>	5 to 30% cpu reduction for LOAD INDEX 30% cpu reduction for LOAD PARTITION up to 70% cpu reduction for LOAD REPLACE PARTITION with dummy input Improvement of 33% when loading a single partition in comparison to DB2 V8			Utilities	Perf 6.1.2 Tech 13.6
<b>1.1.47 REBUILD INDEX performance</b>	5 to 30% cpu reduction REBUILD INDEX CPU reduction of 19% when rebuilding a single partitioning index			Utilities	Perf 6.1.3 Tech 13.6
<b>1.1.48 REORG performance</b>	40 to 50% cpu reduction for REORG INDEX 9% CPU time improvement when reorganizing a table space with a single NPI defined	10% cpu reduction for 8m row PS_PAY_EARNINGS. 6% cpu reduction for 26m row PS_PAY_DEDUCTION		Utilities	Perf 6.1.4 Tech 13.6 White Paper
<b>1.1.49 RUNSTATS performance</b>	30 to 40% cpu reduction for RUNSTATS INDEX 33% reduction in both CPU and elapsed times when performing RUNSTATS on the partitioning index	50% cpu reduction for 8m row PS_PAY_EARNINGS		Utilities	Perf 6.1.5 Tech 13.6 White Paper
<b>1.1.50 Index key generation improvements</b>	All utilities take advantage of the improvements that are made to index key generation with VARCHAR columns varying in length from 1 to 14. There is a reduction of 6% to 66% in both CPU time and elapsed time in comparison to V8 depending on whether NOT PADDED is used	Benefits Unicode implementations		Utilities	Perf 6.1.6 Tech 13.6
<b>1.1.51 Other V9 that are Enhancements to V8 via APARS</b>	<ul style="list-style-type: none"> <li>▪ CHECK INDEX SHRLEVEL CHANGE</li> <li>▪ Cross loader support for LOBs</li> <li>▪ LOAD and UNLOAD of LOBs</li> <li>▪ Automatic display of blocking claimers for REORG</li> <li>▪ Online REORG enhancements</li> </ul>			Utilities	LOB 6.3.3 LOB 6.1 and 6.3 Tech 11.21

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.1.52 Data sharing logging improvement</b>	Any LRSN update spins are done without holding the log latch, so other work is not blocked. This saves both CPU cycles and reduces log latch contention.			Data Sharing	Perf 8.6 Tech 5.1
<b>1.1.53 Reduction in LOB locks</b>	Changes to LOB locking mean LOB data must be externalised for GBP-dependent LOB table spaces, before locks are released.	N/A		Data Sharing	Perf 8.7 LOBS 4.5.2 Tech 5.3
<b>1.1.54 Locking constraint relief</b>	With z/OS 1.7 the XES lock limit on the number of locks IRLM can use (about 3.5 million) in the CF structure has been increased. (Data Sharing)			Data Sharing	Tech 5.4
<b>1.1.55 Improved group buffer pool write performance</b>	For group buffer pools that are duplexed, DB2 9 eliminates cross invalidations as a result of the secondary being updated.			Data Sharing	Perf 8.9 Tech 5.5
<b>1.1.56 Improved workload balancing within the same LPAR</b>	Improved workload balancing within the same LPAR			Data Sharing	Perf 8.11 Tech 5.7
<b>1.1.57 Group buffer pool dependency removal by command</b>	Command to remove Group buffer pool dependency to avoid data sharing overhead			Data Sharing	Perf 8.12 Tech 5.8
<b>1.1.58 Allow table level retained locks to support postponed abort URs+A166</b>	Maintain the checkpoint information at table level so that each table in a segmented table space can be tracked independently. With each table now having it's own lock, an application is not blocked from using other tables within a multi-table table space, due to one table being subject to a postponed abort. (Data Sharing)			Data Sharing	Perf 8.4 Tech 5.14
<b>1.1.59 Simplification of the special open processing</b>	During the forward and backward log phases of restart, data sets are opened and locks obtained. There is lock contention around taking these conversion locks, which can prolong restart times. This enhancement is to remove the need for conversion locks during the "special open". (Data Sharing)			Data Sharing	Tech 5.15 Perf 8.5
<b>1.1.60 Coupling Facility Control Code (CFCC) Level 15</b>	This enhancement includes increasing the allowable tasks in the Coupling Facility (CF) from 48 to 112. (Data Sharing)			Data Sharing	Tech 5.16
<b>1.1.61 Opening data sets earlier in restart processing</b>	This enhancement identifies pagesets not yet opened during the reading of the logs and schedules an asynchronous open			Backup/Recover /Restart	Perf 8.3 Tech 5.13
<b>1.1.62 RLF enhancements</b>	Middleware servers accessing DB2 through DDF use the same plan name, DISTSERV. DB2 RLF function is enhanced to provide the middleware servers a better way to govern their resources. The enhancements are provided by using a new Resource Limit Middleware Table (RLMT) and a unique index in the existing database DSNRLST	PeopleSoft uses DISTSERV plan for all SQL not coming from z/OS.		Governor	Tech 13.4 SAP 13.6

## 1.2 OPERATIONAL

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
1.2.1 Log I/O enhancements	<ul style="list-style-type: none"> <li>The number of active log input buffers is increased from 15 to 120.</li> <li>Archive log files are now read by Basic Sequential Access Method (BSAM) instead of Basic Direct Access Method (BDAM), so that the archive log files can use Data Facility Storage Management Subsystem (DFSMS) striping and compression.</li> <li>The number of archive log files input buffers is increased from 1 to N. N is proportional to the number of stripes. For each stripe, DB2 uses 10 tracks worth of buffers, regardless of the block size.</li> <li>Convert archive log processing from AMODE(24) to AMODE(31).</li> <li>Exploit z/OS DSNTYPE=LARGE disk data set support</li> </ul>			Backup/Restore/Restart	Tech 3.4 Perf 5.6
1.2.2 Deferring the updates of SYSLGRNX till after end of restart	The updating of the SYSLGRNX entries are now triggered by the first system checkpoint following restart, therefore allowing restart to complete quicker. Note this benefits DB2 restart in a non data sharing environment as well.			Backup/Restore/Restart	Perf 8.2 Tech 5.12
1.2.3 SMS constructs	Allows you to utilize SMS storage classes on the CREATE STOGROUP statement when defining DB2 managed datasets. This increases the flexibility of handling these data sets, while it minimizes the manual effort involved. You can now manage large "hot" DB2 data sets using DFSMS.			SMS	Tech 7.13

## 1.3 MEMORY

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.3.1 Real storage</b>	Real storage usage has grown by up to 10% . Adequate real storage is important in DB2 performance.			Memory (-)	Perf 4.4
<b>1.3.2 Virtual storage management</b>	Almost all of DDF runs in 64-bit mode and now uses above-the-bar storage in the z/OS Shared Memory Facility. The virtual storage usage in the DIST address space was reduced by 39% for both the CLI and embedded SQL interface. Reduction in CPU due to reduced data moves between the DBM1 and DIST address spaces.	All SQL not coming from z/OS runs through DDF		Memory	Perf 4.5, 4.6 SAP 13.2 Tech 3.1
<b>1.3.3 Virtual storage constraint relief</b>	More DB2 structures above the bar in the DBM1 address space to provide extra VSCR for the DB2 structures that still need to reside below the bar			Memory	Perf 4.3 SP 13.2
<b>1.3.4 EDM pool changes for static SQL statements</b>	DB2 now places portions of the bound /prepared DML statements (SQL statement text, SQLDA for DESCRIBE output, part of the native SQL PL package, and some other areas) into storage above the 2 GB bar.			Memory	Perf 4.3.1 SAP 13.2
<b>1.3.5 EDM pool changes for dynamic SQL statements</b>	Individual dynamic SQL statements have their storage split between above-the-bar and below-the-bar portions.			Memory	Perf 4.3.2 SAP 13.2
<b>1.3.6 Below-the-bar EDM pool</b>	Below-the-bar EDM pool now has no least recently used (LRU) objects, which reduces the EDM latch class 24 serialization contention. However because of the removal of the LRU objects, the below-the-bar EDM pool must be sized to contain peak usage plus a cushion for fragmentation.			Memory	Perf 4.3.3
<b>1.3.7 CACHEDYN_FREELOCAL</b>	There is the potential that a high number of active dynamic SQL statements can cause a critical below-the-bar storage shortage. CACHEDYN_FREELOCAL indicates whether DB2 can free cached dynamic statements to relieve DBM1 below-the-bar storage	Default is 1 (Yes). Only applies to KEEP DYNAMIC(YES ) plans ie Cobol and SQR and only to local cache.		Memory	Perf 4.3.4
<b>1.3.8 WLM assisted buffer pool management</b>	z/OS 1.8 delivers new WLM services which can assist DB2 in making dynamic buffer pool size adjustments based on real-time workload monitoring. CM exploits these new services to allow for dynamic buffer pool size adjustments so that the system's memory resources can be more effectively utilized to achieve workload performance goals.			Memory	Perf 4.8 Tech 13.14, Tech 3.2 SAP 11.5 White Paper

**1.4 DISK**

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>1.4.1 NOT LOGGED table spaces</b>	Allow for the reduction of the amount of information written to the log in situations where absolute recoverability of the data is not required.	10% cpu reduction on update, insert and delete of 26m row PS_PAY_DEDUCTION . About 7GB of log space saved. Could be used for upgrade, temp tables, large loads, non-restartable processes with image copies wrapped around them	No plans. Customer specific mod. Using ALTER TABLESPACE	Disk space	Perf 5.7 Tech 3.3 SAP 12.4 White Paper
<b>1.4.2 Index compression</b>	Can compress your indexes but you need to define your indexes with page sizes larger than 4 KB ie point them to larger pagesize bufferpools..	25-50% space savings on HR and PT indexes that were tested	No plans. Customer specific mod. .	Disk space	Index Compression Perf 5.5 Tech 3.8 SAP 4.1.4 White Paper

## 2. PRODUCTIVITY IMPROVEMENT

### 2.1 OPERATIONAL

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>2.1.1 Recovery to point in time with consistency</b>	The QUIESCE utility can be used to establish a point of consistency that can subsequently be used. The -ARCHIVE LOG MODE(QUIESCE), could perform a similar function. But difficult to get a QUIESCE point. The RECOVER utility is enhanced to automatically detect the uncommitted work at the point in time selected for the recovery. DB2 will then roll back the changes on objects being recovered. After the recover, all the objects involved in the recovery will be in a transactionally consistent state			Backup/Recover/Restart	Tech 4.3 Perf 6.4.2 SAP 7.3
<b>2.1.2 Conditional Restart Log Truncation by Timestamp</b>	The only method to restore your whole DB2 subsystem to a previous point in time without ending up with inconsistent data is to use a conditional restart with log truncation. Up to DB2 V8, you were forced to specify a specific RBA or LRSN at which the log truncation was to be performed. You can now specify timestamps for both, normal conditional restart and for the RESTORE system utility.			Backup/Recover/Restart	Tech 4.5
<b>2.1.3 Support of large block interface</b>	For tape data sets the large block interface is now supported (> 32760 bytes). This can significantly improve COPY and the restore phase of RECOVER.			Backup/Recover/Restart	Tech 11.18
<b>2.1.4 DSN1COPY RESET enhancement</b>	DSN1COPY RESET is changed to reset the dictionary version			Backup/Recover/Restart	Tech 11.19
<b>2.1.5 Toleration of DBET errors during DB2 restart</b>	When you restart DB2, one important part of restart phase two, Current Status Rebuild, is to rebuild the database exception table (DBET). Starting with DB2 9, the way DB2 handles DBET errors during restart of the DB2 subsystem has changed in order to reduce problems due to DBET abends. DB2 now tolerates the DBET error and puts the page set into a restrictive exception state instead.			Backup/Recover/Restart	Tech 4.10
<b>2.1.6 Add Keyword RECOVERBEFORE to RECOVER utility</b>	The Recover utility has been enhanced by providing a new RESTOREBEFORE (RBA or LRSN) parameter option, which allows you to specify an RBA or LRSN value which tells DB2 to not use any Image Copy SYSCOPY entry that has a higher START_RBA than the specified value.			Backup/Recover/Restart	Tech 4.11.2
<b>2.1.7 Avoid copypending restrictive state when broken page detected</b>	Avoid copypending restrictive state when broken page detected			Backup/Recover/Restart	Tech 4.11.6
<b>2.1.8 New ability to cancel DB commands</b>	Cancel database commands that have been issued from the console.			Backup/Recover/Restart	Tech 4.11.7 SAP 14.3.1
<b>2.1.9 Initiate Auto GRECP recovery at the end of restart</b>	Initiate Auto GRECP recovery at the end of restart			Backup/Recover/Restart	Perf 8.1 Tech 5.11 SAP 8.1

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>2.1.10 BACKUP and RESTORE SYSTEM</b>	These utilities are enhanced to use new functions available with z/OS V1R8 DFSMSHsm. <ul style="list-style-type: none"> <li>Allows individual table spaces or index spaces to be recovered.</li> <li>Allows the backup to be implemented directly to tape.</li> <li>The physical copying to disk in the background can be improved by the use of incremental FlashCopy.</li> </ul>			Backup/Recover/Restart	Perf 5.3 Perf 6.4.1/6.4.2 Tech 11.3 SAP 7.1
<b>2.1.11 Additional RECOVER enhancements</b>	The RECOVER utility has been enhanced in two other areas: <ul style="list-style-type: none"> <li>Displaying progress of RECOVER</li> <li>More messages during LOGAPPLY phase of recover utility</li> </ul>			Backup/Recover/Restart	Tech 11.4.1 Tech 4.11.5
<b>2.1.12 COPY improvements</b>	Enhancements are: <ul style="list-style-type: none"> <li>Optionally copy only objects that are in copy pending or informational copy pending state by specifying the SCOPE PENDING option</li> </ul>			Backup/Recover/Restart	Perf 6.9 Tech 11.7
<b>2.1.13 MODIFY RECOVERY</b>	Instead of deletion by date or age, retention criteria can be specified. The new keyword RETAIN is added to the MODIFY RECOVERY syntax.			Backup/Recover/Restart	Perf 6.2 Tech 11.9 SAP 7.2
<b>2.1.14 DSN1LOGP message enhancement</b>	DSN1LOGP has been enhanced to detect the situation where the entire log range is no longer recorded in the BSDS, because for example the archive logs have rolled of. and give a RC 4 along with the message			Backup/Recover/Restart	Tech 11.10
<b>2.1.15 REORG enhancements</b>	There are a series of enhancements to REORG utility, including a significant improvement to availability when reorganising a subset of partitions within a partitioned table space. <ul style="list-style-type: none"> <li>REORG elapsed time reduction</li> <li>Removal of the BUILD2 phase of Online REORG</li> <li>Change to REORG SHRLEVEL REFERENCE by part</li> <li>Online REORG usability and keyword changes</li> <li>LOB REORG SHRLEVEL REFERENCE</li> <li>Intra-REORG parallelism</li> <li>RETRY_DELAY improvement</li> </ul>			Utilities	Perf 6.6 LOB 6.9 Tech 11.1 SAP 8.2 SAP 8.3 Sap 8.4 SAP 8.5
<b>2.1.16 LOAD enhancements</b>	The LOAD utility has been enhanced: <ul style="list-style-type: none"> <li>LOAD restarting restrictions when there are LOBs involved are lifted</li> <li>LOAD and UNLOAD to support rounding with DECFLOAT new data type</li> <li>UNLOAD to skip locked rows</li> </ul>			Utilities	Tech 11.5
<b>2.1.17 Online CHECK DATA</b>	CHECK DATA is enhanced to allow the new CLONE, LOBERROR, and XMLERROR options.			Utilities	Perf 6.7 Tech 11.11 SAP 8.7
<b>2.1.18 DB2I and DSNU CLIST enhancement</b>	The "DB2 Utilities Panel" in DB2I has been enhanced to allow the SDSNLOAD library to be specified.			Utilities	Tech 11.14
<b>2.1.19 Timestamps in messages</b>	For operational clarity, all utility messages now include the day number from the Julian date, plus a timestamp			Utilities	Tech 4.11.3 Tech 11.16
<b>2.1.20 Timeout or deadlock condition documented in utility output</b>	Timeout or deadlock condition documented in utility output			Utilities	Tech 4.11.2
<b>2.1.21 DSN1PRNT output for broken pages formatted</b>	DSN1PRNT output for broken pages formatted			Utilities	Tech 4.11.4

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>2.1.22 Large format data sets</b>	Large format data sets (>65535 tracks) are automatically supported when they are input data sets to utilities. They are supported as output for utilities if the DD card specifies table spaces that are created with DSNTYPE=LARGE.			Utilities	Tech 11.17
<b>2.1.23 RTS in catalog</b>	New catalog tables: SYSIBM.SYSTABLESPACESTATS and SYSIBM.SYSINDEXSPACESTATS are provided, which contain Real Time Statistics for table and index spaces.			Utilities	SAP 5.1
<b>2.1.24 DSNJU004</b>	Shows the output for the checkpoint queue of a DB2 9 subsystem.			Utilities	Tech 11.20
<b>2.1.25 TEMPLATE switching</b>	The new template switching function allows image copies of varying sizes to have different characteristics. This provides significant flexibility in terms of the data set names and attributes, for example device types			Utilities	Perf 6.8 Tech 11.8
<b>2.1.26 REFRESH for EARLY code</b>	An IPL is no longer needed for early code. Instead you can make use of the new REFRESH DB2 command			Install	Tech 4.6
<b>2.1.27 Automated memory monitoring</b>	Time-based memory monitoring function that can help to alert you when certain storage utilization thresholds are reached			Monitor	Perf 4.3.5
<b>2.1.28 Instrumentation and processes for virtual storage monitoring</b>	Instrumentation has been added to report on the usage of these new and changed virtual storage pools. The information that is collected and presented in IFCID 225 is summary information.			Monitor	Perf4.3.6
<b>2.1.29 z/OS shared memory usage measurement</b>	New fields have been added to IFCID217 and IFCID225 that can be used to report on the shared memory usage.			Monitor	Perf 4.5.1
<b>2.1.30 Automatic identification of latch contention and DBM1 below-the-bar virtual storage</b>	Uses memory monitor			Monitor	Perf 4.9 Tech 3.5 Perf 4.3.5
<b>2.1.31 DB2 Optimization Expert for z/OS</b>	DB2 Optimization Expert for z/OS, V1.1 supports monitoring and tuning of SQL statements that run as part of a workload on your DB2 for z/OS subsystem.			Tool	Perf 4.19 Tech A.1.1 SAP 5.3
<b>2.1.32 DataQuant for z/OS</b>	DataQuant for z/OS, V1.1 delivers a comprehensive query, reporting, and data visualization platform for both Web and workstation-based environments. While remaining compatible with the QMF product line, DataQuant introduces a variety of powerful business intelligence capabilities.			Tool	Tech A.1.2
<b>2.1.33 DB2 Utilities Suite for z/OS</b>	DB2 Utilities Suite for for z/OS, V9.1 supports all major new functions and structural changes in the DB2 for z/OS product			Tool	Tech A.1.3
<b>2.1.34 DB2 Accessories Suite for z/OS</b>	DB2 Accessories Suite for z/OS, V1.1 is made up of a series of independent components designed to enhance your use of the DB2 for z/OS data server <ul style="list-style-type: none"> <li>▪ Optimization Service Center for DB2 for z/OS improves the performance of DB2 SQL queries at both the individual query and full SQL workload level</li> <li>▪ Spatial Support for DB2 for z/OS contains a set of spatial data types, user-defined functions, and stored procedures for spatial related queries</li> <li>▪ The International Components for Unicode for DB/2 for z/OS (ICU) is a set of C/C++ and Java libraries for Unicode support and software internationalization</li> </ul>			Tool	Tech A.1.4

## 2.2 DEVELOPMENT

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
2.2.1 New data types	BIGINT, BINARY, VARBINARY, DECFLOAT			SQL	Perf 2.14 Tech 6.1 SP 3.1,2,3
2.2.2 INSTEAD OF triggers	INSTEAD OF triggers are triggers that are processed instead of the update, delete or insert operation that activates the trigger. Unlike other forms of triggers that are defined only on tables, INSTEAD OF triggers can only be defined on views.			SQL	Perf 2.13 Tech 6.2
2.2.3 MERGE	The MERGE statement updates a target (a table or view, or the underlying tables or views of a fullselect) using the specified input data. Rows in the target that match the input data are updated as specified, and rows that do not exist in the target are inserted.		Other platforms support this statements, so a MetaSQL could be coded. Potential future enhancement. No target release at this time.	SQL	Perf 2.5 Tech 6.3 SAP 12.1
2.2.4 SELECT FROM MERGE/UPDATE/DELETE	Provides support for SELECT FROM MERGE, SELECT FROM UPDATE and SELECT FROM DELETE, in addition to V8 SELECT FROM INSERT			SQL	Perf 2.6 Tech 6.4
2.2.5 ORDER BY and FETCH FIRST in subselect	Allows specification of these clauses as part of subselect in addition to V8 support for fullselect.			SQL	Perf 2.7 Tech 6.5 SAP 11.2 SAP 16.6
2.2.6 TRUNCATE	With V8, to empty a table you have to either do a mass delete (that is, use DELETE FROM table-name without WHERE clause) or use LOAD Utility with REPLACE REUSE and LOG NO NOCOPYPEND. If there is a delete trigger on the table, using the DELETE statement requires you to DROP and subsequently recreate the deleted trigger to empty the table without firing that trigger. LOAD REPLACE works on a table space level instead of on a table level. You cannot empty a specific table if the belonging table space contains multiple tables. TRUNCATE statement addresses these problems	Only of benefit if there are Delete Triggers. None known of in HCM or FIN.	Potential for MetaSQL to support EPM. No target release at this time.	SQL	Perf 2.8 Tech 6.6 SAP 12.2 SAP 13.5 SAP 16.5
2.2.7 INTERSECT and EXCEPT	The UNION, EXCEPT, and INTERSECT clauses specify the set operators union, difference, and intersection. UNION is already supported. For DB2 family compatibility, DB2 9 introduces EXCEPT and INTERSECT.			SQL	Per 2.11 Tech 6.7
2.2.8 Built in functions	ASCII_CHR, ASCII_STR COLLATION_KEY DECRYPT_BINARY DIFFERENCE EBCDIC_CHR EBCDIC_STR EXTRACT LOCATE_IN_STRING LPAD MONTHS_BETWEEN NORMALIZE_STRING OVERLAY RID RPAD SOUNDEX TIMESTAMPADD TIMESTAMPDIFF TIMESTAMP_FORMAT TIMESTAMP_ISO UNICODE UNICODE_STR VARCHAR_FORMAT CORRELATION COVARIANCE COVARIANCE_SAMP			SQL	Tech 6.8
2.2.9 CURRENT SCHEMA	Allows the value of CREATE SCHEMA to different from the value in CURRENT SQLID.			SQL	Tech 7.5
2.2.10 LOB file reference	The purpose of file reference variables is to import or export data between a LOB column and an external file outside of the DB2 system			SQL	LOB 2.4 Tech 7.7

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>2.2.11 FETCH CONTINUE</b>	This enhancement introduces extensions to the FETCH SQL statement that provide a convenient method for applications to read from tables that contain LOB or XML columns, when the actual length of the LOB or XML value is not known or is so large that the application cannot materialize the entire LOB in memory.	N/A		SQL	LOB 4.6.2 Tech 7.8
<b>2.2.12 Spatial support</b>	for location and Geographic data	N/A		SQL	Tech 7.9 SAP 15.1
<b>2.2.13 Native SQL procedures</b>	Provides support for both external SQL stored procedures, which still require a C compiler and still exist as external load modules outside of DB2, and native stored procedures which run entirely within DB2 and do not require a WLM Stored Procedure address space	N/A		SQL	Perf 4.15 Tech 3.9
<b>2.2.14 Nested compound statements in SQL stored procedures</b>	You can now use: <ul style="list-style-type: none"> <li>A compound statement within a condition handler.</li> <li>Nested compound statements to define different scopes for SQL variables, cursors, condition names, and condition handlers.</li> </ul>	N/A		SQL	Tech 7.15
<b>2.2.15 New stored procedure related special registers</b>	Two new special registers are available with native SQL procedures: CURRENT DEBUG MODE and CURRENT ROUTINE VERSION	N/A		SQL	Tech 7.16
<b>2.2.16 PureXML</b>	The new PureXML support in DB2 manages XML in a hierarchical format, representing the XML data model in a natural way. This provides key advantages over existing offerings that require XML data to be stored as large objects which can be expensive to search and update, or decomposed across multiple relational tables which requires complex schema mappings and query statements.	N/A		SQL	Tech 8.2
<b>2.2.17 SOA</b>	The SOAP UDF extends the flexibility of DB2 to be efficient in performance and in connectivity to partners by using the SOAP interface. The Developer Workbench is a comprehensive development environment that can help you develop and maintain your SOA applications	N/A		SQL	Perf 7.3
<b>2.2.18 Changes to DB2-supplied dynamic SQL processing applications</b>	Prior to V9, SQL procedures have to be processed by either the precompiler or the SQL procedure processor DSNTPSMP stored procedure. Now you can easily use SPUFI, DSNTEP2, or DSNTEP4 to process your SQL procedures.	N/A		Tool	Tech 7.18
<b>2.2.19 Unified Debugger</b>	With the Unified Debugger, you can observe the execution of SQL procedure code, set breakpoints for lines, and view or modify variable values. The Unified Debugger supports external and native SQL procedures, including nested stored procedures. To use the Unified Debugger with SQL procedures, you must include breakpoints in your routines or executable files.	N/A		Tool	Tech 7.20
<b>2.2.20 Cultural sort</b>	Addresses national alphabets which have special two character characters which are not correctly sorted in current DB2 sorting algorithm.	Unicode related	No known plans currently.	Misc	Tech 7.4

## 2.3 ADMINISTRATIVE

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
2.3.1 Universal table space	Table spaces that are both segmented and partitioned are called Universal Table Spaces™ (UTS). Two flavors: partition-by-growth (PBG) and partition-by-range (PBR). Simple tables spaces are deprecated. You will not be able to create new simple table spaces.		Under consideration. Potentially could be built into PeopleSoft Tablespace DDL Automation tool	DDL	Perf 5.1 Tech 3.6 SAP 12.3 SAP 13.4 SAP 14.2
2.3.2 Clone table support	Provides you with the ability to generate a table with the exact same attributes as a table that already exists. This new table is called the clone table of the base table and it is created in the same table space as the base table, with indexes, views, triggers and security persisting. Once the clone table is created, you can load or insert rows into it. Then EXCHANGE DATA.	Better approach for upgrade. Do not lose indexes, views etc.	Under consideration for future Tools release	DDL	Perf 5.2 Tech 4.2 SAP 8.5
2.3.3 RENAME INDEX	You can use the extended RENAME statement to rename existing indexes			DDL	Tech 4.7 SAP 4.1
2.3.4 RENAME COLUMN	Under some circumstances you are now able to rename columns of an existing table without having the need to drop and recreate the object.	Already supported on some other db platforms	Under consideration for future Tools Release	DDL	Tech 4.8
2.3.5 UPDATE SCHEMA	This enhancement supports changing schema, creator, owner and VCAT name of objects using CATMAINT utility. This is a system-wide function meaning that all objects are altered from one schema name to another. Creator, owner and VCAT names are also updated in a similar manner.			DDL	Tech 4.9
2.3.6 Automatic creation of objects	You can now create the following objects implicitly depending on how you code your CREATE TABLE SQL statement: <ul style="list-style-type: none"> <li>• Database</li> <li>• Table space</li> <li>• Enforcing primary key index</li> <li>• Enforcing unique key index</li> <li>• ROWID index if the ROWID column is defined as GENERATED BY DEFAULT</li> <li>• LOB table space, auxiliary table, auxiliary index</li> </ul>	Allows for DDL platform compatibility	Under consideration for future PeopleTools release	DDL	Perf 2.15 Tech 7.1 SAP 13.3
2.3.7 Allow ALTER TABLE ... LONG VARCHAR TO VARCHAR	You can alter the data type of LONGVARCHAR and LONG VARGRAPHIC columns to VARCHAR and VARGRAPHIC respectively using the ALTER TABLE ALTER COLUMN syntax and specify the length stored in the catalog table for the LONG VARCHAR or LONG VARGRAPHIC column.			DDL	Tech 7.10
2.3.8 Trace filtering	You can be more specific on the -START TRACE command by specifying additional include parameters as well as exclude parameters			Trace	Tech 3.10.3
2.3.9 Allow EXCLUDE qualification in DSNWQAL for READS	Omegamon/DB2 PE include/exclude threads			Trace	Tech 3.10.4
2.3.10 DB2 and DSN commands for native SQL procedures	DIS procedure, start/stop procedure, rebind package			Command	Tech 7.17
2.3.11 Enhanced messages when unable to get p-locks	Enhanced messages when unable to get p-locks			Data Sharing	Perf 8.14 Tech 5.10

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
2.3.12 Internet Protocol Version 6 support	DB2 9 now supports IPV6 which is the next generation of the internet protocol designed to replace the current version, IPv4.			TCP	Tech 9.1 SAP 9.1
2.3.13 Run DDF without SNA when using TCP/IP only	DDF can be setup to no longer activate its SNA/APPC support and thus no longer requires VTAM be installed (VTAM may still need to be installed for other reasons).			TCP	Tech 9.3

### 3. SECURITY

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>3.1.1 SSL Network Encryption</b>	DB2 has a new DRDA secure port, which allows use of Secure Socket Layer(SSL) authentication methods.			Security	SAP 6.2 Tech 10.9.3
<b>3.1.2 Network trusted context</b>	Trusted context provides a technique to work with other environments more easily than before, improving flexibility and security. A trusted context addresses the problem of establishing a trusted relationship between DB2 and an external entity, such as a middleware server, eg: <ul style="list-style-type: none"> <li>• IBM WebSphere® Application Server</li> <li>• IBM Lotus® Domino®</li> <li>• SAP NetWeaver®</li> <li>• Oracle® PeopleSoft®</li> <li>• Siebel® Optimizer</li> </ul>			Security	Perf 7.1 Tech 10.9.2
<b>3.1.3 More security options with INSTEAD OF triggers</b>	You can use INSTEAD OF triggers to encode and decode data from the database within a view.			Security	Tech 10.3
<b>3.1.4 Audit Management</b>	This tool can be used by DBAs and auditors to provide all the information needed for an audit while also maintaining database security.			Security	Tech 10.5

#### 4. DSNZPARM Changes

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>4.1</b> <i>These parameters are eliminated:</i>					Tech 12.9.1
4.1.1 DSN6SPRM.RELCURHL:	Release at COMMIT any page or row locks on which a cursor WITH HOLD is positioned. The field RELEASE LOCKS has been removed from panel DSNTIP8				
4.1.2 DSN6SYSP.DBPROTCL:	Default protocol (DRDA or PRIVATE) to be used when DBPROTOCOL is not explicitly specified for the bind of a plan or a package. The field DATABASE PROTOCOL has been removed from panel DSNTIP5.				
4.1.3 DSN6SYSP.STORPROC:	Name of the address space procedure for the DB2-established SPAS. The field DB2 PROC NAME has been removed from panel DSNTIPX.				
4.1.4 DSN6SPRC.SPRMIFS:	Maximum allocated storage area for IFI READS.				
4.1.5 DSN6SPRM.PARAPAR1:	Enables optimizer enhancements added by APAR PQ87352.				
4.1.6 DSN6SPRM.TABLES_JOINED_THRESHOLD:	Minimum number of table joins in a query to cause DB2 to monitor the resources consumed when determining the optimum access path for that query.				
4.1.7 DSN6SPRM.MAX_OPT_ELAP:	Maximum amount of elapsed time in seconds to be consumed by the DB2 Optimizer.				
4.1.8 DSN6SPRM.SJMXPOOL:	Max MB of the virtual memory pool for star join queries. The field STAR JOIN MAX POOL has been removed from panel DSNTIP8.				
4.1.9 DSN6SPRM.SUPPRESS_TS_CO NV_WARNING:	Suppress table space conversion warnings (hidden keyword)				
4.1.10 MORE_UNION_DISTRIBUTION					
<b>4.2</b> <i>These parameters are modified:</i>					
4.2.1 DSN6SPRM.MGEXTSZ:	Use sliding secondary quantity for DB2-managed data sets. Default changes from NO to YES. On panel DSNTIP7, OPTIMIZE EXTENT SIZING				
4.2.2 DSN6SYSP.IDXBPOOL:	Default bufferpool for CREATE INDEX. Now accepts 8 KB, 16KB, and 32 KB page size BPs in addition to 4 KB.				
4.2.3 DSN6SYSP.WLMENV:	Name of the default WLM environment for DB2. Length has increased from 18 to 32 characters.				
4.2.4 CACHEDYN_FREELOCAL	The default has changed from 0 in DB2 V8 to 1 in DB2 9				
4.2.5 MAXARCH	The default has changed from 1000 to 10000.DSNTIPA, RECORDING MAX				
4.2.6 MINSTOR	The default has changed from NO to YES. DSNTIPE, MANAGE THREAD STORAGE				
4.2.7 NPGTHRS	-1 is no longer supported. Values must be 0 or any integer greater than 1				

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
<b>4.3</b> <b>These parameters are added:</b>					
<b>4.3.1</b> <b>DSN6SYSP.IMPDSDEF:</b>	Define the data set when creating an implicit base table space or implicit index space. On panel DSNTIP7, DEFINE DATA SETS. Range: YES or NO.Default: YES.				
<b>4.3.2</b> <b>DSN6SYSP.IMPTSCMP:</b>	Enable data compression an implicit base table space. On panel: DSNTIP7, USE DATA COMPRESSION. Range: NO or YES. Default: YES.				
<b>4.3.3</b> <b>DSN6SPRM.MAXTEMPS:</b>	Max MB of temp storage in the work file DB that can be used by a single agent at any given time for all temporary tables. On panel: DSNTIP9, MAXTEMP STG/AGENT. Range: 0 - 2147483647. Default: 0 (meaning no limit ~ previous releases).				
<b>4.3.4</b> <b>DSN6SYSP.MAXOFILR:</b>	Maximum number of open files for LOB file references. On panel: DSNTIPE, MAX OPEN FILE REFS. Range: 0 - value of MAX USERS field. Default: 100.				
<b>4.3.5</b> <b>DSN6SYSP.TBSBP8K:</b>	Default 8 KB BP for CREATE TABLESPACE. On panel: DSNTIP1,DEFAULT 8-KB BUFFER POOL FOR USER DATA. Range: BP8K0 - BP8K9. Default: aBP8K0				
<b>4.3.6</b> <b>DSN6SYSP.TBSBP16K:</b>	Default 16 KB BP for CREATE TABLESPACE. On panel:DSNTIP1, DEFAULT16-KB BUFFER POOL FOR USER DATA. Range: BP16K0 - BP16K9.Default: BP16K0				
<b>4.3.7</b> <b>DSN6SYSP.TBSBP32K:</b>	Default 32 KB BP for CREATE TABLESPACE. On panel:DSNTIP7, DEFINE DATA SETS. Range: YES or NO. Default: YES.				
<b>4.3.8</b> <b>DSN6SPRM.SYSTEM_LEVEL_B ACKUPS:</b>	Whether the RECOVER Utility should use system-level backups as a recovery base for object level recoveries. On panel: DSNTIP6,SYSTEM-LEVEL BACKUPS. Range: NO or YES. Default: NO.				
<b>4.3.9</b> <b>DSN6SPRM.RESTORE_RECOV ER_FROMDUMP:</b>	Specifies for the RESTORE SYSTEM and the RECOVER Utilities whether the system-level backup that has been chosen as the recovery base should be used from the disk copy of the system-level backup (NO), or fromthe dump on tape (YES). On panel: DSNTIP6, RESTORE/RECOVER. Range: NO or YES.Default: NO.				
<b>4.3.10</b> <b>DSN6SPRM.UTILS_DUMP_CLA SS_NAME:</b>	DFSMSHsm dump class for RESTORE SYSTEM to restore from a system-level backup dumped to tape. On panel: DSNTIP6, DUMP CLASS NAME. Range: Valid SMS class name. Default: blank.				
<b>4.3.11</b> <b>DSN6SPRM.RESTORE_TAPEU NITS:</b>	Max tape units that RESTORE SYSTEM can allocate to restore a system-level backup from tape. On panel: DSNTIP6, MAXIMUM TAPE UNITS. Range: 0 - 256 or NOLIMIT. Default: NOLIMIT.				
<b>4.3.12</b> <b>DSN6SPRM.STATCLUS:</b>	RUNSTATS clustering statistics type. On panel: DSNTIP6, STATISTICS CLUSTERING. Range: ENHANCED or STANDARD. Default: ENHANCED. The ENHANCED clustering statistics result in better SQL access paths when there are duplicate values of the clustering index or when the data is in reverse clustering order. These enhanced clustering statistics may cause many access paths to change. The STANDARD option results in the same clustering statistics as				

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
	Version 8.				
4.3.13 DSN6SPRM.MXDTCACH:	Max MB of virtual memory allocated for data caching. On panel: DSNTIP8, MAX DATA CACHING. Range: 0 - 512. Default: 20. If you specify zero, DB2 does not use data caching during query execution and you can apply sparse index only. If you specify a value between 1 and 512, DB2 allocates memory from above the 2 GB bar pool up to the specified size for data caching.				
4.3.14 DSN6SPRM.EDM_SKELETON_POOL:	Minimum size in KB of the EDM skeleton pool. On panel: DSNTIPC, EDM SKELETON POOL SIZE. Range: 5000 - 1048576. Default: 102400 The skeleton pool, which is located above the bar, is a component of EDM for storing skeleton cursor tables (SKCTs) and skeleton package tables (SKPTs), where each SKCT represents an application plan and each SKPT represents a package.				
4.3.15 DSN6SPRM.OPTXQB:	Enable/Disable global query optimization. No panel. Range: ON or OFF. Default: ON				
4.3.16 DSNHDECM.DEF_DECFLOAT_ROUND_MODE:	Default rounding mode for decimal floating point data type. On panel: DSNTIPF, DECFLOAT ROUNDING MODE. Range: ROUND_HALF_EVEN, ROUND_HALF_UP, ROUND UP. Default: ROUND_HALF_EVEN ROUND_CEILING, ROUND_DOWN, ROUND_FLOOR, ROUND_HALF_DOWN, DSN6SPRM.ADMTPROC ( JCL proc name used to start the task address space for the DB2 administrative scheduler)				
4.3.17 DSN6SPRM.HONOR_KEEPDICTIONARY	Specifies whether DB2 should honor or ignore the KEEPDICTIONARY parameter on LOAD REPLACE or REORG request that converts a table space from basic row format to reordered row format. Valid settings are NO and YES. The default is NO.				
4.3.18 DSN6SPRM.MAX_CONCURRENT_PKG_OPS	Specifies the maximum number of concurrent package requests that can be processed simultaneously.				
4.3.19 DSN6SYSP.TBSBPLOB	Name of the bufferpool to be used by LOB table spaces that are implicitly created. Any valid DB2 bufferpool name is accepted. The default is BP0				
4.3.20 DSN6SYSP.TSBPXML	Name of the bufferpool to be used by XML table spaces that are implicitly created Any valid 16K BP name is accepted - the default is BP16K0				

---

## 5. FUNCTIONS NO LONGER SUPPORTED

Feature	Short Description	Relevancy to PeopleSoft	Product Plans	Category	Reference
5.1 <i>Functions no longer supported</i>	<a href="http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.inst/db2z_functionsnotsupp.htm">http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.inst/db2z_functionsnotsupp.htm</a>				