

Compliments of SoftBase Systems, Inc.
www.softbase.com 1-800-669-7076

DB2 UDB for z/OS, Version 8: New Release Highlights

Introduction

According to IBM, DB2 Universal Database (UDB) for z/OS, Version 8 is the biggest release in DB2 history. DB2 UDB Version 8 became generally available three years after Version 7, with more than a thousand additional person years of work. This reference card is designed to provide you with a quick overview of key new features that can make a large impact in your company or organization.

DB2 for MVS, OS/390, & z/OS Release History

- **Relational Model Theory:** E.F. Codd (IBM) - 1970
- **SQL/DS - VM:** 1st IBM Relational DBMS - 1981
- **DB2 Version 1** - 1983
- **DB2 Version 2** - 1988
- **DB2 Version 3** - Dec. 1993
- **DB2 Version 4** - Nov. 1995
- **DB2 Version 5** - Jun. 1997
- **DB2 Version 6** - Jun. 1999
- **DB2 Version 7** - Mar. 2001
- **DB2 Version 8** - Mar. 2004

Prerequisites for DB2 UDB for z/OS, Version 8

- DB2 UDB for z/OS Version 7 for migration to DB2 UDB for z/OS Version 8
- zSeries Processor that supports z/Architecture
(**new hardware purchase may be required**)
- z/OS V1R3 Base Services or later

Key Open Standards for DB2 Version 8

- SQL, Java, J2EE, XML, ODBC, JDBC, SQLJ, DRDA, Unicode, and Web Services

Performance Enhancements

- **64 Bit z/Architecture** - allows for Virtual Storage expansion up to 16 Exabytes per address space. DB2 Buffer Pools, EDM Pools, Sort Pools, RID Pools, and Compression Dictionaries can now be increased in size and located above the 2 GB bar. Bigger buffer pools and additional memory provide dramatic overall performance improvements for DB2 systems.

SQL Enhancements

- **GET DIAGNOSTICS** - replacement for utilizing the SQLCA. This new function retrieves additional information about each SQL statement, its condition and connections. GET DIAGNOSTICS is also required to retrieve information about the new DB2 V8 SQL multi-row functions below and to accommodate the longer names for DB2 V8 objects.
- **Multi-row FETCH and INSERT** - allows users to execute multiple FETCHES or INSERTs from one SQL statement. This new enhancement will be very useful for DB2 scrolling applications and for retrieving blocks of DB2 data into a program at one time. For example:

```
INSERT INTO T1 FOR :host variable ROWS
VALUES( :ARRAY1, :ARRAY2) ATOMIC;
```

- **SEQUENCES** - new feature that allows DB2 to generate recoverable and guaranteed-unique sequential key values for new rows being added to DB2 tables. These values can be generated across multiple tables and even data sharing environments.
- **Multiple DISTINCT** - the DISTINCT keyword can now be used in multiple columns functions with different expressions in the same SQL statement:

```
SELECT SUM (DISTINCT C1), AVG(DISTINCT C2)
FROM TABLE_1;
```

- **SELECT from INSERT** - new statement allows you to both insert a new table row with IDENTIFY, SQUENCE, or user defined data values and to then retrieve the values of those columns with a single SQL operation. Following is an example of the new statement:

```
SELECT C3 FROM
INSERT (C1, C3, C5) INTO T1
VALUES('ABC', CURRENT DATE, 'DEF');
```

Data Warehousing Enhancements

- **Materialized Query Tables (MQTs)** - New MQTs boost Data Warehouse performance by pre-calculating, aggregating, or summarizing data one time and then saving this information for future use. Results saved in MQTs can then be repeatedly accessed by later Data Warehouse queries. Elapsed execution times of several hours can be reduced to minutes or sometimes even seconds. The new MQT feature of DB2 Ver. 8 is tailor-made to dramatically improve performance in large multi-dimensional DB2 Data Warehouse environments.

- **Common Table Expressions and Recursive SQL** - this new feature allows SQL statements to construct a result set and then repeatedly utilize this information to construct answers to many common Data Warehouse type queries.
- **New Parallel Sorting** - the sort process has been enhanced to allow multi-table sorts to be run in parallel. Sort parallelism can substantially reduce elapsed time for large Data Warehouse queries.
- **New Star Join Improvements** - the new sparse indexing enhancement builds an index of only qualified rows - this new feature eliminates unqualified data rows earlier in the access path process, dramatically improving query response time and CPU performance in large Data Warehouse queries
- **Dedicated Virtual Pool for Star Join Workfiles** - repeated access to composite data in new Virtual Memory Pools can substantially improve performance of Star Joins and parallel sorts required by large DB2 Data Warehouse OLAP queries.

ERP Enhancements (SAP, PeopleSoft)

- **System Level Point-In-Time Recovery** - this new enhancement provides the capability to recover the DB2 system to any point-in-time very quickly, regardless of the presence of uncommitted units of work. This enhancement significantly improves recovery times for DB2 applications such as SAP that have very large numbers of tables.
- **Volatile Table Recognition** - The new VOLATILE table DDL key word identifies tables to the DB2 V8 optimizer that are empty or have widely varying amounts of data. SAP and PeopleSoft applications frequently contain many extra tables that are not used and can be accessed more efficiently with this new enhancement. Also, ERP and CRM software packages with temporary work tables will perform better when this option is specified for these tables.

More DB2 UDB for z/OS, Version 8 Resources

- **IBM DB2 UDB for z/OS, Version 8 Homepage:**
<http://www.ibm.com/software/db2zos>
- **Softbase DB2 UDB for z/OS, Version 8 Customer Care Package:**
<http://www.softbase.com/carepackage/>
- **IBM DB2 UDB for z/OS, Version 8 Redbook:**
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg246079.html?Open>

IBM & DB2 are the registered trademarks of International Business Machines Inc. The names of actual companies and products mentioned herein may be the registered trademarks or trademarks of their respective owners.

Performance Enhancements (continued)

- **Improved SQL Query Optimization** – improved Stage 1 and Stage 2 processing will allow SQL statements to execute faster without any alterations or changes required by developers. Also, RUNSTATS now collects statistics about data distribution for non-indexed columns, enabling the DB2 optimizer to select improved access paths. **Note:** Optimization improvements require Dynamic SQL or a BIND or REBIND of SQL statements.
- **Compare for Unlike Data Types** – DB2 can now treat predicates comparing columns/values of different data types as indexable and Stage 1 processable. DB2 Version 8 can now also use indexes and Stage 1 processing to join tables using columns with different data types.
- **Support for Backward Index Scan** – DB2 can now utilize indexes in reverse order. This enhancement removes the need to construct both ascending and descending indexes on the same column.
- **Long and Variable Length Index Keys** – DB2 V8 now allows true varying length keys to be stored in index structures rather than padding VARCHAR and GRAPHIC type columns to their maximum length.
- **New Control Interval Sizes** – DB2 control intervals can now be 8K, 16K, or 32K in addition to the standard 4K size. Large sequential data access performance will be noticeably improved in DB2 V8.
- **Multi-row FETCH, INSERT, and UPDATE** – new group row processing can provide up to a 50% or even higher performance boost under some circumstances – for a significant fraction of the work.

Security Enhancements

- **Multilevel Security** – new enhancement uses security labels to provide row-level granularity for access control.

Availability Enhancements

- **Dynamic Partitioning** – DB2 UDB for z/OS, Version 8 has the ability to now immediately add partitions, rotate partitions, and change the partitioning key values of table-controlled partitioned tables. Newly added and rotated partitions are immediately available for usage.
- **Online Schema Changes** – Many DB2 object attributes can now be altered while DB2 is still available for use. The following Schema changes are now allowed by DB2 UDB for z/OS, Version 8:

- > Extend CHAR(n) column lengths
- > Change type within character data types (CHAR, VARCHAR)
- > Change type within numeric data types (SMALLINT, INTEGER, FLOAT, REAL, DOUBLE, DECIMAL)
- > Change type graphic data types (GRAPHIC, VARGRAPHIC)
- > Allow column data type changes for columns that are referenced within a view
- > Allow these column changes for columns that are part of an index
- > Add a column to an index
- > Drop the partitioning index
- > Change the clustering index
- > Create/alter an index to have non-padded varying length character columns within a key
- > Allow an alter of identity columns
- > Add a partition to the end of a table which extends the limit value
- > Rotate partitions
- > Support automatic rebalancing of partitions during REORG
- > Support REORG of parts in REORG pending states
- > Loosen the restrictiveness of indexes in recover or rebuild pending states
- **Online Index Reorganization** – indexes can now be REORGed while the DB2 is still available for access.
- **Partitioned Tablespace Maintenance Improvements** – REORG, LOAD, and RECOVER of individual partitions is now possible – all other partitions are still available for access and update. DB2 V8 now allows for up to 4096 partitions for tables, which allows for more granular separation of data for performance and administration reasons. Archiving and rebalancing production data can be accomplished much faster without losing database availability.
- **Online DSNZPARMs** – DB2 Version 8 allows additional DSNZ-PARMS to be altered online and to take effect without having to recycle the DB2 system.

Scalability/Compatibility Enhancements

- **64 bit Architecture** – 31 bit addressing 2**31 changing to 2**64 allows for dynamic growth of DB2 environments and exploits the power and capacity zSeries and z/OS.

- **SQL Statement Size Increased** – SQL statements can now be up to 2 MB in length.
- **DB2 Table Join Limit Increased** – Up to 225 tables can now be joined at a time in DB2 V8 for z/OS without any previous restrictions imposed by DB2 V7.
- **DB2 Table, View, and Alias Name Sizes Increased** – DB2 table name lengths have increased from 18 to 128 characters. Longer DB2 table, view and alias names will allow for more descriptive naming standards in shops.
- **DB2 Column Name Size Increased** – DB2 column name lengths have increased from 18 to 30 characters. Longer DB2 column names will allow for more descriptive naming standards in shops.
- **Index Keys Expanded** – Index key sizes have been expanded from 255 to 2,000 characters.
- **DB2 Table Partitions Increased** – the number of available partitions has increased from 256 to 4,096. Also, partitions can now be administered and maintained completely independently.
- **Maximum DB2 Table Size Increased** – DB2 32K page table spaces have been expanded in data capacity from 16TB to 128TB.
- **DB2 Active and Archive Log Increases** – The number of DB2 active logs has been increased from 31 to 93. Also, the number of archived DB2 logs has increased from 1,000 to 10,000. The increase of available logs allows DB2 systems with high transaction volumes to remain recoverable for much longer periods of time.
- **Character Literals Expanded** – Character literals have been expanded from 255 to 32,704 characters .
- **New Unicode Support** – Unicode support has been enhanced substantially, with the ability to parse SQL in Unicode, to join tables with different encodings (ASCII, EBCDIC or Unicode), and to convert the DB2 catalog to Unicode.
- **New DB2 XML Functions** – DB2 V8 now provides a set of functions that allows DB2 on z/OS to generate XML data for exchange with other applications and database platforms.
- **New IBM DB2 Universal Driver for SQLJ and JDBC** – provides an open and consistent set of protocols to access data from DB2 on z/OS the same way as from UNIX and Windows database platforms.