

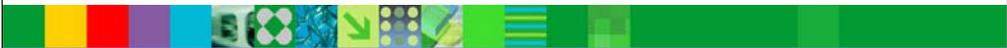


IBM Software Group

Reasons to Migrate to DB2 UDB for z/OS Version 8, Part 4 of 4: Incremental Scalability, Migration & References

DB2 UDB for z/OS Development Team,
IBM Silicon Valley Lab

DB2 Information Management Software



IBM **DB2 Information Management**
Technical Roadshows

Spring, 2005

@business on demand software

© IBM Corporation 2005

Find out how DB2® Universal Database for z/OS® Version 8 has been re-engineered, with fundamental changes in architecture and structure that will help you manage your very large databases more easily and cost effectively. "Reasons to Migrate to DB2 UDB Version 8" is a day long seminar suited to database technical managers, database administrators, applications developers and systems programmers. It focuses on four critical areas: availability, integration, application development productivity, and flexible growth and incremental scalability. Its information-packed sessions will familiarize you with the enhancements that are enabling organizations to streamline database management, respond more flexibly and quickly to business needs and ensure short and long-term growth capacity.

Agenda

- **Availability:** (Roadshow Part 1)
Now not even structural changes can stop DB2
- **Integration:** (Roadshow Part 2)
Increasing reliability, security and flexibility
- **Productivity:** (Roadshow Part 3)
Faster, easier application development
- ➔ • **Incremental Scalability:** (this presentation)
Capacity exactly when you need it
- **Migration Planning:** (this presentation)
What do you need to do to get ready for V8?

Many of the items I'm talking about today improve performance, availability, productivity and scalability. We'll try to keep the items that are closely related together, but categories don't always work when you need to choose and the real answer is both.

The agenda for today is roughly

8:15 am Registration and Continental Breakfast

8:45 am Welcome and Introduction

9:00 am Availability: Now not even structural changes can stop DB2

10:30am Break

11:00 am Integration: Increasing reliability, security and flexibility

12:30 pm Lunch

1:30 pm Productivity: Faster, easier application development

3:00pm Break

3:30 pm Incremental Scalability: Capacity exactly when you need it

Migration Planning: How to get there

5:00 pm Close of Program

Customer Value Highlights: Incremental Scalability Theme

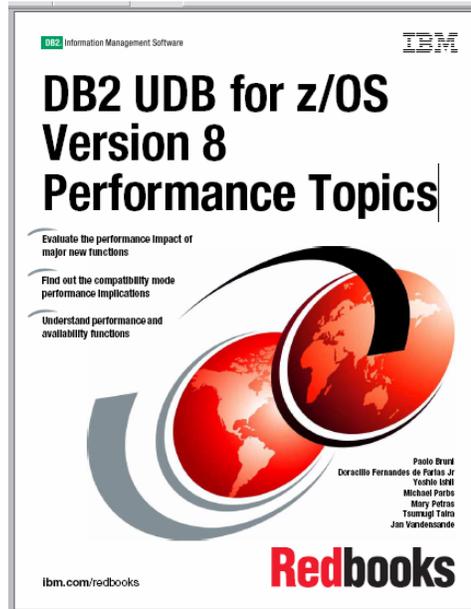
- **Add more applications and users than ever:** system constraints expanded to let you grow beyond current imagination
- **Use more storage, more effectively, with less complexity:** 64-bit virtual addressing delivers
- **Better performance for frequently-used operations:** Multi-row functions reduce overhead for repeated operations, index improvements return the right data faster, and materialized query tables provide pre-calculated summary data quickly

In this section we will talk about the scalability and performance improvements in DB2 UDB for z/OS Version 8.

**Performance
Details**

Measurements

462 pages



This new IBM Redbook is on the web now, with the detail about Version 8 performance. Check the web, Here are the URLs for Performance Topics:

<http://www.redbooks.ibm.com/abstracts/sg246465.html?Open>

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246465.pdf>

For all recent updates, start on the Support page, click on Redbooks, sort by date – newest first

<http://www.ibm.com/software/data/db2/zos/support.html>

Scaling DB2 for z/OS



Image of Earth from Moon, Source: NASA (Public Domain)

Breaking through limitations: scaling

- **Virtual Storage 2 GB** 2^{31} to 2^{64}
- **Partitions** 254 to 4096
- **Active logs** 31 to 93
- **Archive logs** 1000 to 10,000
- **Maximum table size** 16 TB to 128 TB (partitioned, 32K page)
- **Maximum data sets** 32K to 64K

One of the keys to reengineering is breaking through the limits of the current architecture. Increasing some limits improves scalability. Increasing other limits improves productivity, portability & family consistency.

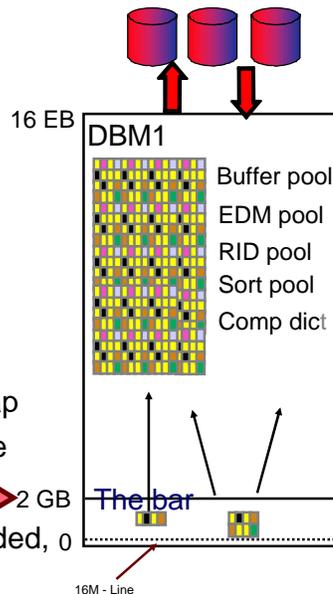
Increasing the amount of virtual storage we can address directly can help with the ability to scale and simplify management for virtual storage. It will require more real memory, but permit increased scalability and availability.

Increasing name sizes & SQL statement lengths makes porting from other DBMS much easier and improves DB2 family compatibility. Increasing the maximum number of partitions helps DB2 scale farther and makes management much easier when you can have one partition per day for 11 years.

With changes in V8 and in z/OS 1.5, the number of concurrently open data sets can increase to 64K, depending upon virtual storage below the 16 MB line. See APAR PQ96189?

64 bit evolution

- zSeries, z/OS, z/Architecture & large real storage
 - Real storage for performance
 - ▶ V6 data space advantages
 - zSeries & 64 bit addressing
 - Version 8 64 bit virtual
 - ▶ z/OS 64-bit Virtual Storage Roadmap
 - ▶ Improve scalability, availability, ease
 - ▶ Move above the bar
 - ▶ Hiperpool/dataspace no longer needed, 0 no ECSA for locks



The biggest impact of the zSeries architecture on DB2 is the ability to use more memory more effectively. Prior to the zSeries, customers were limited to 2 GB real storage due to the 31-bit addressing of the S/390 architecture. The real storage limit of 2 GB is a leading performance inhibitor for many high end customers. Another performance inhibitor is the 2 GB virtual storage limit for the main DB2 (DBM1) address space. If you have zSeries & OS/390 V2R10 64-bit mode or z/OS, use V6 buffer pools in data spaces, but not otherwise. See V7 Performance Topics red book & the web. See What's New? for V8 use of 64 bit virtual storage.

There are more steps as real & virtual memory sizes increase, moving more above the line and above the bar. See the Roadmap, GM13-0076-01 updated June 2002.

ibm.com/servers/eserver/zseries/library/whitepapers/gm130076.html

4096 Partitions

- Maximum number of parts raised from 254 to 4096
 - ▶ Table spaces and indexes
 - ▶ Table space use DSSIZE to go beyond 254 parts
- ALTER TS ADD PART adds partitions to the end
- Maximum table size remains 16 TB for 4 KB pages
- Data set naming convention
 - ▶ 'Axxx' - partitions 1-999 'Bxxx' - partitions 1000-1999
 - ▶ 'Cxxx' - partitions 2000-2999 'Dxxx' - partitions 3000-3999
 - ▶ 'Exxx' - partitions 4000-4096
- Maximum partitions allowed depends upon page size & DSSIZE
 - 4K page size, DSSIZE=1GB => 4096 parts, 4 TB max table
 - 4K page size, DSSIZE=64GB => 256 parts, 16 TB max table

The maximum number of partitions goes from 254 to 4096, so that you can have one partition per day for more than 11 years. Do not define all of the partitions, since you can add new partitions at the end.

While the maximum size remains 16 terabytes for 4K pages, larger page sizes can exceed 16 terabytes for a single table. DSSIZE is recommended, although LARGE also allows more than 254 parts.

This change requires changes in the data set naming convention.

Performance: DB2 UDB for z/OS

Performance Enhancements



- Ability to use indexes more often
 - More matching in predicates
 - Index backward scan
 - Index varying length
- Distribution statistics on non-index columns
- Materialized Query Tables
- Multi-row fetch & insert
- Long Term Page Fix Bufferpool option
- Data Sharing improvements

Many of the performance enhancements come from the enhanced index capabilities. Comparing a decimal column to a floating point number could not be done in the past. Being able to compare these values can mean choosing a better index or even an index access instead of a table space scan. An index can be used for a backward scan, so some indexes may not be required, reducing the overhead for inserts & deletes. Support for varying length indexes can save space and can use index-only access with a varchar. Having distribution statistics for columns which are not part of the index can give the optimizer better information so it can perform better optimization. Gathering these statistics was a separate DSTATS program, and is now part of RUNSTATS.

Materialized query tables can provide a one or two order of magnitude performance improvement by rewriting queries to use the precalculated information.

Stage 1 Indexable Unlike-types

- DB2 enhanced to allow index access when host variable & target column are not the same data type
- Deals with programming languages that don't support the full range of SQL data types
 - ▶ C/C++ has no DECIMAL data type
 - ▶ Java has no fixed length CHAR data type
- Significant performance improvement for many applications
- Simplifies application programmer & DBA tasks

It is sometimes necessary for customers to join tables on columns with different data types, or provide a search value with a data type that doesn't match the column's definition. For example, allow C double predicate to be indexable with a SQL DECIMAL(p,s). As z/OS data is accessed from PCs, UNIX, Linux and Java, we need to have a C data type of DOUBLE be indexable with a SQL data type of DECIMAL(p,s) or NUMERIC(p,s).

Currently, for many types of predicates, if the data types of the predicate operands do not match, then the predicate is residual. This can have a large negative effect on the performance of the query.

This change will improve the performance of queries that involve predicates with mismatched data types by allowing those predicates to be Stage 1, and possibly Indexable, (subject to certain restrictions).

RUNSTATS Distribution Statistics Enhanced

- **Non-uniform distribution statistics on non-indexed columns can improve performance 10X, 100X**
- **RUNSTATS improvement allows optimizer to consider non-uniform distribution statistics on columns that aren't part of an index**
- **Current technique is separate DSTATS program**
- **Significant performance improvement**
- **RUNSTATS with UPDATE NONE REPORT NO**
 - **Fast invalidation of dynamic statement cache**
- **HISTORY statistics without updating main statistics**

Skewed data distributions are responsible for a high proportion of performance problems with DB2 queries, especially in ad hoc queries. Symptoms can be less than optimal join sequences, too much synchronous I/O, and long response times. When there is asymmetrical distribution of data, not having distribution statistics on non-leading indexed columns and/or non-indexed columns can cause DB2 to make sub-optimal table join order and table join method decisions.

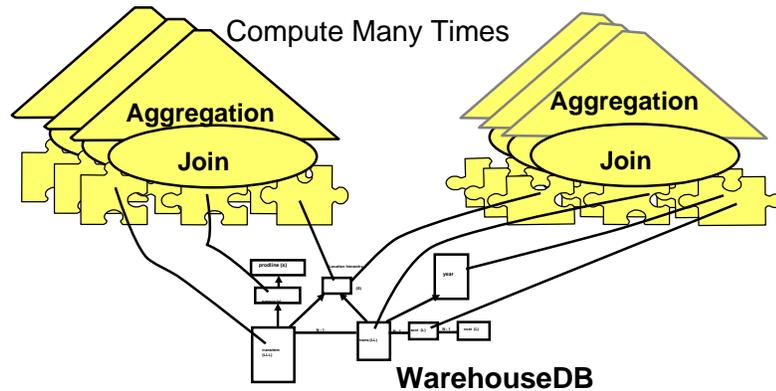
Collecting distribution statistics for non-leading indexed columns and/or non-indexed columns allows DB2 to use these statistics for better access path selection. Better index selections can be made, when there are screening predicates or there are matching in-list / in-subq predicates which break up matching equals predicates.

RUNSTATS with keywords REPORT NO and UPDATE NONE allows users to invalidate dynamic SQL caching for the table space and/or index space without the overhead of collecting statistical information and without generating reports or updating catalog tables.

Without Materialized Query Tables: Each Query Re-Computes!

Q11,Q12, ...

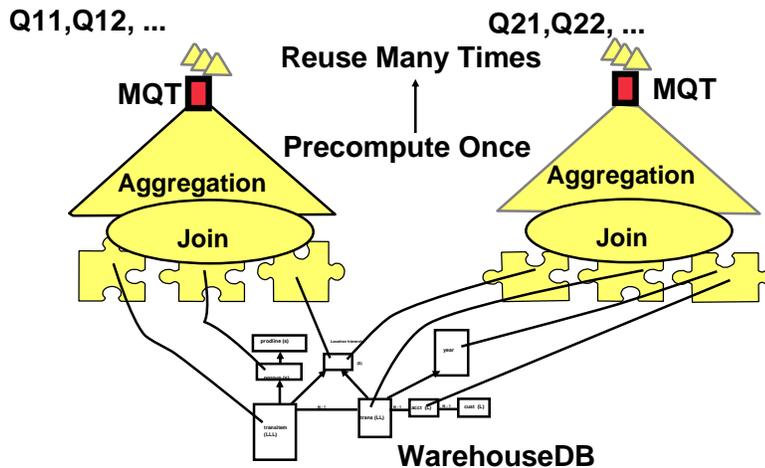
Q21,Q22, ...



Part 1 11

A materialized query table (MQT) can avoid redundant work of scanning, aggregation and joins. Multiple levels of summary tables have been used in warehouses and complex applications for years. One of the major issues is communicating the summaries to the users. In some cases, the users want to query the base data. With MQTs, the query users do not have to be aware of the MQT.

With Materialized Query Tables: Avoid Redundant Computation



Even though the query is submitted for the base table, the optimizer can rewrite the query to use the MQT. Using the precalculated information can improve subsequent queries by as much as two or three orders of magnitude. Materialization or precalculation and parallelism resolve the long response times.

A database administrator can use an MQT much as she or he would use an index for optimization. Controls for usage, initial loading and refresh are part of the definition.

Materialized Query Tables

- Sometimes called Summary Tables, Automatic Summary Tables, Automatic Materialized Query Tables, Materialized Views, ...
- Optimizer can rewrite queries to access MQT instead of base table or view
- Pre-computed information, very significant performance improvement
- Managed by user or system (SQL REFRESH)
- Automatic rewrite or manual
- Informational Referential Integrity (not enforced)

Many names are used for the MQTs in various implementations, including automatic summary tables, aggregate summary tables, automatic materialized query tables and materialized views. They all have a common objective, to precompute information for later use by other statements.

The objective is to make the use as transparent as possible. Customers may choose to have the MQTs managed and refreshed by the system or manually. Most customers will have the optimizer rewrite queries, but MQTs can be names in SQL statements.

Multi-row INSERT



- Inserts multiple rows on one API call
- Can be ATOMIC or NOT ATOMIC
- Can be static or dynamic SQL
- Significant performance boost

```
INSERT INTO T1
VALUES( :ARRAY1, :ARRAY2)
FOR :hv ROWS
NOT ATOMIC CONTINUE ON SQLEXCEPTION;
```

Being able to insert multiple rows helps in application portability and can improve performance, particularly across a network.

You can have atomicity (all rows must be successfully inserted or all are rolled back) or insert and then resolve any errors.

This change supports host language arrays, adding another in a range of customer enhancement requests.

Multi-row FETCH



- Returns multiple rows on one API crossing
- "wide" cursor with locks on multiple rows
- Supports scrollable and non-scrollable, static and dynamic SQL
- Significant performance boost

```
DECLARE C1 CURSOR
  WITH ROWSET POSITIONING
  FOR SELECT COL1, COL2 FROM T1;
OPEN C1;
FETCH NEXT ROWSET FROM C1
  FOR :hv ROWS INTO :ARRAY1, :ARRAY2;
```

Multiple row FETCH also helps with application portability and performance. It provides a new concept called a "wide" cursor, which contains multiple rows rather than just one. A rowset is a set of rows that is retrieved through a multiple-row fetch.

Being able to FETCH multiple rows at once can make a larger percentage improvement in performance, compared to INSERT. The FETCH statement requires less processing than INSERT, generally.

Multirow performance

- **FETCH:**
 - Up to 50% cpu savings
 - DSNTIAUL change made
 - DSNTEP4 usage
 - Distributed processing
- **INSERT**
 - Up to 30% cpu savings
 - Best for CICS, distributed, class 2 accounting, many rows, few columns



Here are some program changes already implemented to save cpu time. If you use DSNTIAUL, then the changes are included. If you use DSNTEP2, then changing to DSNTEP4 can save a lot of cpu time if you fetch large numbers of rows. Multirow fetch is used by distributed processing. If you have coded ODBC with the array interface, multirow fetch is used. The improvements can vary a lot, but the biggest savings will be where the processing is simple, in the CICS (non OTE) environment, with distributed processing where latency and TCP/IP processing can be avoided, with class 2 accounting on, with large numbers of rows, but few columns. For example, one customer measurement was a 76% improvement.

Long Term Page Fix

- ALTER BUFFERPOOL
 - New option PGFIX(YES)
 - By buffer pool
 - Use where IO rate is high
 - Must have real storage
 - Up to 8% cpu savings
 - Allowed in all modes, e.g. compatibility mode



ALTER BUFFERPOOL has a new option that most customers should use for subsystems which read or write frequently. Recommendation: Alter your DB2 Version 8 buffer pools which have frequent page reads or writes to use PGFIX YES if you have sufficient real storage available for these buffer pools. Fixing the buffer page **once and keeping them fixed** in real storage avoids the processing time that DB2 needs to fix and free pages **each time there is an I/O**. In some cases, this processing time can be as much as 10% for I/O intensive workloads. To use this option, issue the following command:

```
ALTER BPOOL(bpname) VPSIZE(vpsize) PGFIX(YES)
```

where *bpname* is the name of the buffer pool and *vpsize* is the size of the virtual pool.

REOPT(ONCE) BIND Option

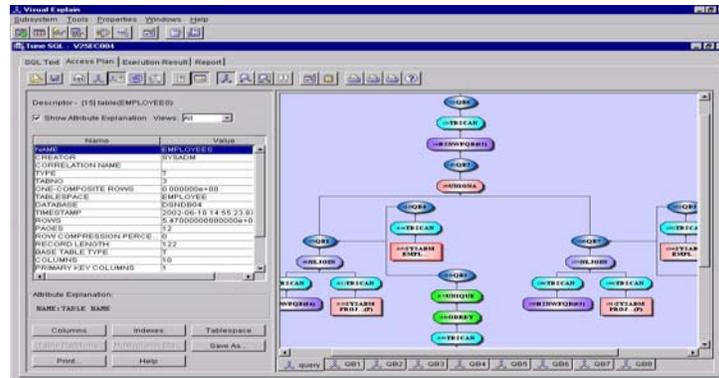
- Controls when DB2 builds access path for dynamic SQL
- Default, access path is calculated at PREPARE.
- REOPT(ONCE)
 - Defers access path selection until OPEN
 - Values of host variables on OPEN used to calculate access path
 - Resulting access path cached in global prepare cache



REOPT(ONCE) is a new bind option that tries to combine the benefits of REOPT(VARS) and dynamic statement caching. For an SQL statement with input host variables, no matter it is static or dynamic, the access path chosen by the optimizer during bind time (before the values of host variables are available) is not optimal in general. Bind option REOPT(VARS) solves this problem by (re)preparing the statement at run time when the input variable values are available, so that the optimizer can re-optimize the access path using the host variable values. However, for frequently called SQL statements that take very little time to execute, re-optimization using different input host variable values at each execution time is expensive, and it may affect the overall performance of applications. REOPT(ONCE) will re-optimize the access path only once (using one set of input variable values - currently the first set) no matter how many times the same statement is executed. The access path chosen based on the set of input variable values is stored in cache and will be used for all later executions. This solution is based on the assumption that the chosen set of host variable values at run time are better than the default ones chosen by optimizer at bind time.

Visual Explain Enhancements: Version 7 and 8

- Significant improvements to Visual Explain tool:
 - ▶ Much more detailed information about access path
 - ▶ More statistical details for each node in the graph
 - ▶ XML document describing query access path selected
 - ▶ Easier collection of information to send for help



Part 1 20

While optimization continues to improve, it has become more and more complex. SQL statements that are 2 MB long, allowing many more tables to be joined, materialized query tables, improved use of sparse indexes for star joins and making more predicates indexable are examples.

Explain table has been rearchitected to include filter factor estimation, cardinality estimation, cost estimation, parallelism, and system configuration. This information can be printed using XML structures, allowing the information to be collected more easily. We expect this process to make problem access path resolution easier.

Visual Explain can be downloaded now from the web and can be used on Version 7 or Version 8 subsystems, although some of the new function is only provided with Version 8.

<http://www.ibm.com/software/data/db2/zos/osc/ve/index.html>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OA3.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/Z31m.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/Z27.pdf>

Enhanced Instrumentation

- Package level accounting
- Accounting rollup for DDF & RRSF
- New and improved traces, larger monitor area
 - Long-running reader
 - Lock escalation
 - Full SQL statement
 - PREPARE attributes
 - High water marks
 - Secondary authorization ids
 - Storage 64-bit, READS
 - Dynamic statement cache
 - Temporary space usage
 - Auditing for multilevel security
 - Option for EBCDIC or Unicode data



Many enhancements are made in instrumentation, helping to monitor and account for the larger and more varied workloads. Additional information is provided at a package level, if those traces are on. Accounting can roll up multiple trace records into one for DDF and RRSF. A new IFCID is provided for lock escalation. The full SQL statement (not just 5000 bytes) can be traced with a new IFCID. The PREPARE statement attributes can be traced. The statement id is added to dynamic statement cache traces. Secondary ids can be retrieved with a synchronous read in an APAR that was added to V6 and V7. Additional fields were added to storage IFCIDs 225 and 217 for 64 bit addressing. Dynamic statement cache traces were improved to be more usable. A new IFCID 0342 was added for temporary space use by agents. Auditing was added for multilevel security. See the Release Guide, Appendix F for new and changed instrumentation.

Larger CI size for larger page sizes

- New CI size equals page size by default
 - e.g. 16K CI for 16K page, DSNZPARM to set
- Enables I/O striping for 8K, 16K, 32K page
- Higher data rate for 8K, 16K, and 32K page
 - 16K page measurement with 16K instead of 4K CI
 - ✓ +36% for non EF (Extended Format) data sets
 - ✓ +70% for EF data sets
 - ✓ VSAM EF getting nearly equivalent to non EF in data rate performance

A new option in DB2 V8 allows use of larger CI sizes with 8K, 16K and 32K pages, rather than using 4 K pages on disk. There are very substantial performance improvements in disk performance for this change. Note that 32K CI size results in 16K block size, so space is not wasted on disk. The improvement in performance is better for Extended Format data sets, so that having a larger DSSIZE is more reasonable for performance.

Agenda

- **Availability:** (Roadshow Part 1)
Now not even structural changes can stop DB2
- **Integration:** (Roadshow Part 2)
Increasing reliability, security and flexibility
- **Productivity:** (Roadshow Part 3)
Faster, easier application development
- **Incremental Scalability:** (this presentation)
Capacity exactly when you need it
- ➔ • **Migration Planning:** (this presentation)
What do you need to do to get ready for V8?

Next we will talk about where to get information on migration planning. This is a very high level overview, but there are good presentations, papers and books on the web.

Four step migration process

1. Test with New Install & Migrate
2. Migrate to new release without new function (**COMPAT Mode**)
 - a. Apply SPE to all members
 - b. Start all members
 - c. Start in V8 COMPAT Mode
3. **Enable New Function Mode**
4. Then in **New Function Mode**



The process of migration will take four steps.

Use a new install on a separate subsystem and practice migration on the separate subsystem. Test the function you want to use. This will ensure that you are ready to move forward.

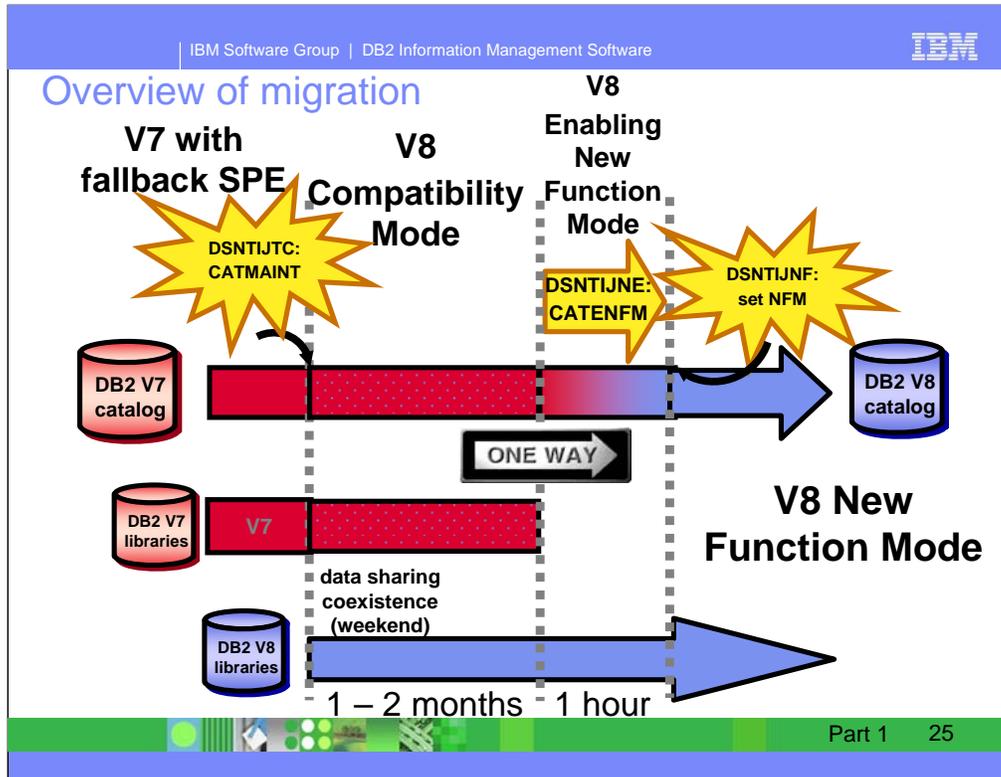
Apply the fallback SPE PQ48486 on every member of each data sharing group and on all subsystems to migrate. Applying the SPE was required only for data sharing in prior versions, but is now required for all subsystems.

Migrate to the new release in Compatibility Mode (CM). You are using the new code but almost no new function.

The third step is Enabling New Function Mode (ENFM).

Finally, after all of the members of a data sharing group are running the new version, you can switch to New Function Mode (NFM).

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/Z41.pdf>



This is an overview of the migration process. There are two steps which modify the catalog, CATMAINT to change into Compatibility Mode (CM), and also CATENFM for Enabling New Function Mode (ENFM), which contains the major catalog changes to long names and Unicode, using a technique of an online (read only or shrlevel reference) reorg. Keep the time for data sharing coexistence fairly short, such as over a weekend or one week. Compatibility mode will often be one or two months. Once you move to ENFM, you don't go back to CM or to V7. ENFM will generally be on the order of an hour. Then you turn the switch and are taking advantage of the new function.

Getting ready for V8

Planning and Education

- ✓ Get needed education, books
 - ➔ Classes, conferences, web
 - ➔ Understand dates, V8 changes, Unicode, migration process, ...
- ✓ Resolve incompatible changes
- ✓ Determine new functions to be used
- ✓ Plan & Install prerequisite hardware & software
- ✓ Perform health check



This is the outline of planning for a DB2 UDB for z/OS Version 8 migration. The most important step is being sure that you have all of the needed information. There are many good classes, ranging from a full transition class to conferences to specialized education for others at your site. There are many books, with more coming. Many of the presentations and books are on the web.

There are a substantial number of changes that can require changes in applications and databases. Get the information to the people who need it.

Reseaching the prerequisite software for the function you will be using, then installing it can take time, especially if you are not current today.

Being ready to test and to manage the inevitable surprises is very important. This is an excellent time for a health check on availability, performance, storage, commit & copy frequency, BIND currency, ...

<ftp://ftp.software.ibm.com/software/data/db2zos/S1322rm2.pdf>

Getting ready for V8



- Migrate to DB2 for z/OS V7
- Migrate to z/OS 1.4 or later (1.3 ended)
 - WLM goal mode required
 - DB2 requires z/OS Unicode services
 - Some functions need z/OS V1R5, R6, R7
- zSeries, z/Architecture 64 bit mode
- Migrate to IBM COBOL V3.3 PL/I V3.4
- IMS V9, V8, V7, CICS TS V3, V2.3
- Details: Program Directory, check vendors

Migration is only from DB2 V7, so migrate to DB2 V7.

When you migrate to z/OS 1.3 or later, WLM compatibility mode is not supported. Some of the function will require later versions, so moving to 1.4 will be better for many customers. End of service for 1.3 is March 2005, and you need to be current, so migrate to 1.4 or later, depending upon the functions you need.

This is a 64-bit virtual application, so the new bi-modal accommodation that allows z/OS on zSeries to run in 31 bit mode will not be able to run DB2 for z/OS V8.

The old COBOL compilers (OS/VS COBOL & VS COBOL II) will not be able to be used with the precompiler. Migrate to current IBM Enterprise COBOL compilers. Old COBOL load modules continue to run (under LE for support). COBOL V2 is now out of support, as is Enterprise COBOL V3R1. End of support for Enterprise COBOL and PL/I V3R2 is October 2005.

Get current software for other products, including IMS and CICS. I recommend CICS V3.

Myths, Hoaxes & Urban Legends 1: DB2 UDB for z/OS V8 Migration



- ✓ **Moving to V8 means rewriting applications.**
- ✓ **V8 requires conversion of application data to Unicode.**
- ✓ **Converting to Unicode doubles the storage.**
- ✓ **Unicode doesn't affect me.**
- ✓ **Customers should wait for a code refresh or drop of code.**

Myth 1: Moving to V8 means rewriting the applications.

V8 does have many changes, but compatibility for existing applications is preserved, with minor exceptions, as in prior versions. Reference: Installation Guide section "Make adjustments for release incompatibilities."
Reference: <http://publib.boulder.ibm.com/epubs/pdf/dsnigj10.pdf>

Myth 2: V8 requires conversion of application data to Unicode.

V8 improves handling and processing for data in Unicode format. DB2 catalog data is converted to Unicode. However, the intent is to provide a choice for customers. Most customers will keep their data in its current EBCDIC format for some time. As more customers develop the need to store international data, process data with Java, or take advantage of other improvements, they will convert some of their data to Unicode.

Myth 3: Converting to Unicode doubles the storage requirement.

A common misunderstanding exists that Unicode doubles the disk and memory space for data. Some Unicode formats store every character as two bytes. Newer formats store character and graphic data without expanding it. Although DB2 catalog columns were lengthened and data was changed to Unicode, early customers saw less than 10% growth. In some situations you will want to use data compression. The actual storage requirement for your data will depend on the usage, nature and database design choices for your data.

Myth 4: Unicode doesn't affect me.

Everyone will be impacted by the Unicode changes in V8. Unicode services are required for V8 to start. SQL is parsed in Unicode. Most of the DB2 catalog character data is stored in Unicode. Please refer to Unicode papers for details. References: <http://www.ibm.com/software/data/db2/zos/pdf/ccmstr.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/unicodep1.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/unicodep2.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OA9.pdf>

Myth 5: Customers should wait for a code refresh or another drop of code.

If you want to wait for the code refresh or recommended service level, then you can order it after July 23, 2004. Get the updated Program Directory from the web dated July 2004 for Service Level PDO0424. If you have already installed V8, then you can apply service to this same level. Being relatively current on service is always important, and especially so on a newer version.

References: <http://www.ibm.com/software/data/db2/zos/v8books.html>

<http://publibfp.boulder.ibm.com/epubs/pdf/i1085662.pdf>

Myths, Hoaxes & Urban Legends 2: DB2 UDB for z/OS V8 Migration



- ✓ **V8 is such a huge version with all this new function. Wait longer.**
- ✓ **Applications and tools are not ready.**
- ✓ **V8 works with only the latest hardware and software.**
- ✓ **Stay in compatibility mode for 6 months.**
- ✓ **We don't need to plan V8 migration carefully.**

Myth 6: V8 is such a huge version with all this new function. Wait longer.

The large amount of new function and code are not a myth. Many actions were taken to provide better quality than prior versions. A substantial amount of the "new" code has been delivered in other DB2 family products, so the delivery is more robust. The test process and early ship program were much longer, more rigorous and more robust than in prior versions. For example, the weekly regression tests began earlier in the development cycle and had a broader, more rigorous work load. Native stress tests ran using all of the power of a 16 engine zSeries z990 with more diverse work loads and configurations. The stress test work loads added heavy use of SQLJ and JDBC under WebSphere, Linux for zSeries clients and tested multiple concurrent work loads. Stress test ran for 21 months, rather than 12 months in V7. One early customer ran a successful regression test with over 4300 batch jobs on 320 GB of data.

Reference: <ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OVSP.pdf>

Myth 7: Applications and tools are not ready.

SAP and PeopleSoft have certified their large applications. We started to work with vendors and tools nearly two and a half years before general availability, by providing detailed specifications. Vendors had code roughly 18 months prior to general availability and many have been able to deliver tools that are ready. Contact your vendors for specifics.

For IBM tools, see Reference:

http://www.ibm.com/support/docview.wss?rs=434&context=SSZJXP&q1=DB2+V8&uid=swg21162152&loc=en_US&cs=utf-8&lang=en

Myth 8: V8 works with only the latest hardware and software.

DB2 V8 works with most current hardware and software. DB2 V8 does require a zSeries processor running in 64-bit virtual mode and z/OS V1.3 or later. If you are not running DB2 Version 7 now, you first need to migrate to V7. If your software is out of service, such as an old compiler, you need to migrate. For level information, see the Program Directory. Reference: <http://publibfp.boulder.ibm.com/epubs/pdf/1085662.pdf>

Myth 9: Stay in compatibility mode for 6 months to a year or more.

Most of the benefits of V8 are provided only when running in new function mode. The primary reason for running in compatibility mode is the ability to fall back. If you don't have a rigorous stress test other than production, compatibility mode is very useful to run until you know you will not need to fall back. Run compatibility mode for a business cycle, through the month end or quarter end. Then move to new-function mode to take advantage of the benefits of V8.

Myth 10: We don't need to plan V8 migration carefully.

Too little time or planning has a very high correlation with problems. You must have adequate time and resources to succeed. When we don't have enough time to read, we get to make time for undoing, reading and redoing.

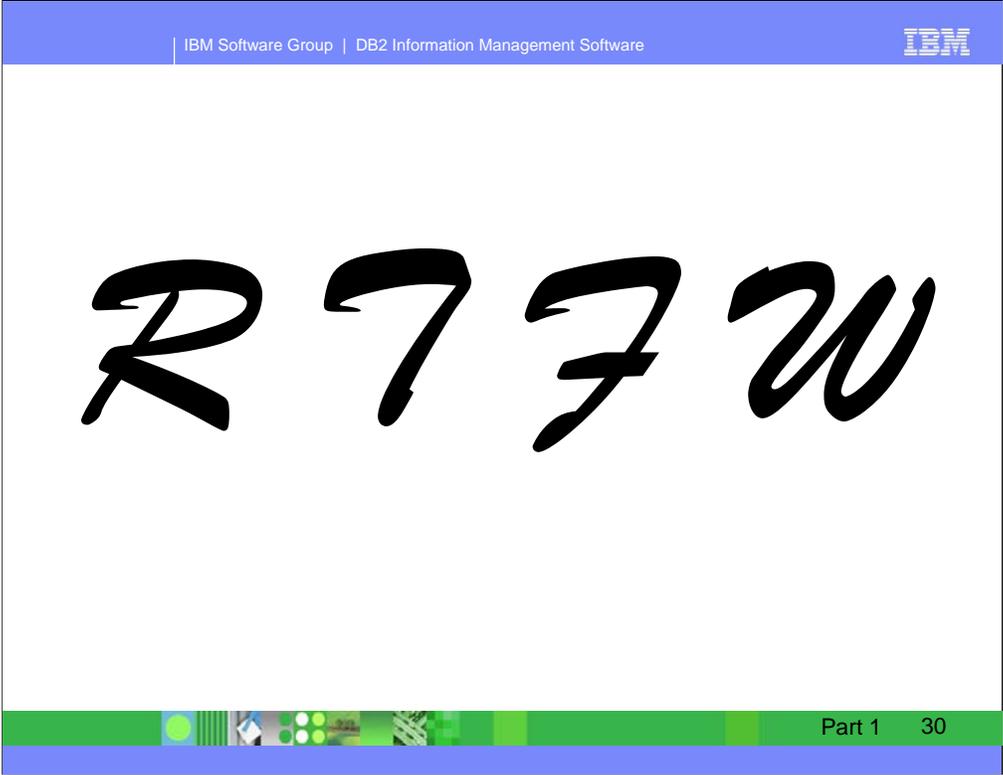
References: <http://publib.boulder.ibm.com/epubs/pdf/dsnigj10.pdf>

<http://www.redbooks.ibm.com/redbooks/pdfs/sq246079.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/V8planmigrate.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/V8migrateprocess.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/S1350.pdf>



RTFW

RTFW is the acronym for Read the Friendly Web. Let's take a short walk on the wild, wild web. The problem with the web is not too little information, but rather too much information. The experience is a bit like trying to take a drink from a fire hose. So I'd like to help a bit by narrowing the search with the twenty five cent tour of a few of my favorite DB2 web sites. A lot more information has been added in the past month or two, with many new books and web pages.

Let's start with the DB2 family. I'll generally show the short form or alias of the URL, omitting `http://`

Here are some tips for avoiding the 404. You don't need `www` in front of `ibm.com` in most situations. There is often something after `www`, such as the `-306` in `www-306` when you get the URL back from the browser. Remove the hyphen and number when you save the URL, since that number changes more quickly than the rest of the URL.

This is the main DB2 UDB for z/OS web page. You can get to the other DB2 for z/OS pages from here, so I often call this my home page. This page changes frequently, so look at the highlighted NEW items. Do you want to look in a DB2 book? Click on Library to see books on DB2 and QMF Version 8 (about 40), Version 7, 6, 5 or even 4. You can check the latest changes by looking at the Information Updates or go to the Information Center. From this page, you can look for conferences (Events), specific classes (Education), or services. If you want to see the latest on Version 8, click on the Version 8 link. If your primary concern is application development, the Developer Domain is for you. DB2 Magazine covers a broad range of topics about DB2. The latest machines z990 and z890 are on the zSeries page. Click DB2 and IMS Tools to see the wide range of help they provide.

IBM Software Group | DB2 Information Management Software

Library V8 page
ibm.com/software/db2zos/library.html

Home | Products | Services & solutions | Support & downloads | My account

Products & services > Software > DB2 Information Management > DB2 Product Family >

DB2 for z/OS and OS/390

Library

Here is your complete DB2 for z/OS and OS/390 information resource library! It includes product documentation, announcement letters, presentations, and important technical tips.

DB2 Information Management Software Information Center for z/OS Solutions

The [DB2 Information Management Software Information Center for z/OS Solutions](#) provides centralized, integrated product information for:

- DB2 UDB for z/OS Version 8 - updated!
- DB2 QMF Version 8
- IMS Version 9 **NEW**
- IMS Version 8
- Several DB2 and IMS Tools **NEW** and updated!

The information center also includes:

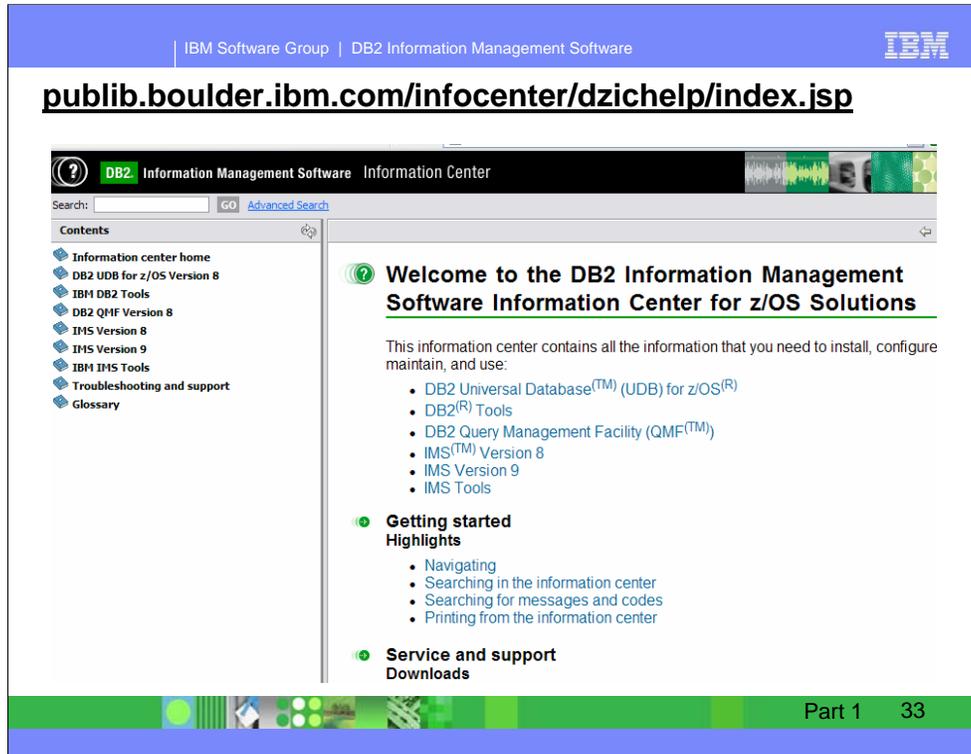
- Integrated LookAT functions for quick message search **NEW**
- Convenient troubleshooting and support information
- Common Glossary
- PDF versions of the product manuals
- Improved usability and information consistency

This information center includes all of the user documentation for the IBM products listed, except those items available to licensees only. Future releases of the information center will include additional DB2 and IMS Tools. Built upon the IBM WebSphere Help System, this information

Part 1 32

This is the Library page for Version 8. I clicked Library, and then Version 8 to get here. There are about 40 PDF books and about the same number of books in html or BookManager format. Choose the style you want. I generally download the PDF books to my computer. From this page you can get books on DB2 and QMF, including both the library and program directories. If you need to look for some tools that can be downloaded, such as Estimator or Visual Explain, click on Downloads. You can also get to the new Information Center from here.

Most of the DB2 V8 books were updated in early November 2004 and we expect to have the next update in April 2005. Watch this web page and see the dates next to each book for the latest updates.



This is the new Information Center, with a wide spectrum of information and access to books for DB2 UDB for z/OS, DB2 tools, QMF, IMS, IMS tools and more. You can get to this page from the Library page, by clicking Information Center. The Information Center provides information across the books and across multiple products.

If you click “Troubleshooting and Support”, then expand under “Searching knowledge base” and click “Web search:...”, you’ll find a helpful Web search page. From this page, you can search IBM support, DeveloperWorks, or even the whole Internet using Google.

The Official Introduction to DB2 for z/OS

Covers DB2 Universal
Database for z/OS Version 8
2nd Edition

Susan Graziano Sloan

Certification
New to DB2 UDB for z/OS



Part 1 34

This is the second edition of the popular book that discusses certification, concepts and architecture. The main part of the book discusses design, SQL and applications in various languages. The special topics include DB2 and the Internet, distributed access and data sharing. Get this book if you are interested in being certified on DB2 UDB for z/OS. This book is excellent for someone who knows about relational database and wants to understand DB2 UDB for z/OS. It became available at the end of May 2004 and can be ordered now from your online retailer, such as Barnes and Noble or Amazon or from iSource

<http://www.ibm.com/isource/cgi-bin/goto?on=IMNL604B38>

<http://vig.prenhall.com/catalog/academic/product/0,4096,0131477501,00.html>

http://www.ibm.com/certify/certs/db_apd.shtml

New DB2 Universal Database (UDB) V8 DBA Professional Certification Role and exam for z/OS

<http://www.ibm.com/isource/cgi-bin/goto?on=naa5ed0351>

<http://www.ibm.com/developerworks/db2/library/techarticle/0207fosdick/0207fosdick.html>

<http://www.ibm.com/software/data/db2/os390/techdocs/s1310.pdf>

DB2 for z/OS Version 8 DBA Certification Guide

Susan Lawson

- Prepare for V8 DBA exam
- Good for all DBAs on the z/OS platform



Part 1 35

Written by IBM Gold Consultant Susan Lawson, this new guide offers comprehensive preparation material for your IBM Certified Database Administrator-DB2 Universal Database V8 for z/OS certification. This guide presents hundreds of useful tips, practical techniques, and expert guidelines for every facet of DB2 UDB database administration, including database implementation, operation, recovery, security, auditing, performance, installation, migration, and much more.

<http://www.ibm.com/software/data/education/bookstore/>

DB2 Developer's Guide

Covers DB2 Universal Database for z/OS Version 8
5th Edition

Craig Mullins



Part 1 36

This comprehensive guide to implementing applications with DB2 was released recently. The new edition adds information on Version 7 and Version 8 to Craig's excellent book. The book has been sold at several conferences lately and has sold out. Craig Mullins is a noted authority on DB2. He writes and presents frequently to many customers. This book is written for database administrators and application programmers who want a complete outline for implementing and optimizing applications. It discusses coding, SQL, database administration and utilities.

www.craigmullins.com

IBM Software Group | DB2 Information Management Software IBM

DB2 UDB Family SQL Reference

<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>

Country/region [select] Terms of use

All of dW

Search

Home
Products
Services & solutions
Support & downloads
My account

[developerWorks](#)
[eServer](#)
[Lotus](#)
[Rational](#)
[Tivoli](#)
[WebSphere](#)
[dW Subscription](#)
[Autonomic computing](#)
[Grid computing](#)
[Java™ technology](#)
[Linux](#)
[Open source projects](#)
[Power Architecture™](#)
[SOA and Web services](#)
[Web architecture](#)
[Wireless technology](#)
[XML](#)
[Feedback](#)

[Related links](#)
[ISV resources](#)
[alphaWorks \(alpha technologies\)](#)
[Academic program](#)
[IBM eReference](#)



The SQL Reference for Cross-Platform Development

June 2002 (Updated September 22, 2004)

© 2002 International Business Machines Corporation. All rights reserved.

Great news for people who want to write portable applications! This reference makes it easy for you to develop applications using SQL that is portable across the [DB2 Universal Database Family](#), including [DB2 UDB for z/OS® Version 8](#), [DB2 UDB for iSeries™ Version 5 Release 3](#), and [DB2 UDB for Linux, UNIX®, and Windows® Version 8.1](#).

New in Version 2!

The addition of DB2 UDB for z/OS V8 and DB2 UDB for iSeries V5R3 means the following great SQL features are now included:

- Support of sequences
- A new CURRENT SCHEMA special register
- Enhanced support for scalar subselect
- New built-in functions including encryption support
- Support for common table expressions
- Support for lateral correlation
- Support of grouping by expressions
- Support for materialized query tables
- Support for altering identity columns
- Support for ITERATE statement
- Increases in several SQL limits, and enhanced long string support
- Other portability enhancements

The *SQL Reference for Cross-Platform Development* uses the familiar format of the product SQL Reference manuals and includes information on the following topics:

- Database concepts
- Built-in functions
- Statements and queries
- SQL procedural language control statements
- Limits
- Programming in C, Java™, COBOL, REXX, and external routines in enrout

Part 1 37

This is the web page for the SQL Reference for Cross-Platform Development. It has three versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions.

<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>

IBM Software Group | DB2 Information Management Software

DB2 Support Page
ibm.com/software/db2zos/support.html

Technotes (FAQs)
 Presentations
 White papers
 Redbooks
 ✓ new DR, perf.
 APARs, CST

DB2 for z/OS
 Features and benefits
 System requirements
 Library
 Success stories
 News
 How to buy
 Events
 Training and certification
 Services
 Support
 My Support
 Submit & track problems
 How to buy software support
 Help
 Site tours
 Feedback

Translate My Page
 Select language

DB2 for z/OS and OS/390

Search support for this product
 Enter search terms, phrase, error code or APAR number

Optionally, limit results
 Solve a problem (FAQs, APARs, Technotes)
 Download (Fixes, Patches)
 Learn (Manuals, White Papers, etc.)

Submit

Advanced search for this product → Search all software support

Self help
 Solve a problem
 → APARs
 → Info APARs
 → Technotes (FAQs)

Download
 → IBM Consolidated Service Test
 → DB2 Client/Server Tips & Samples

Learn
 → Online Publications (Newsletters, Periodicals)
 → Redbooks

Problem submission
 If you didn't find a solution to your problem and you have a maintenance contract:
 → Submit & track problems
 → How to buy support for your software

Flashes
 More flashes

Other resources

Part 1 38

Do you need Answers to Frequently Asked Questions? Product Information? Technical Presentations? Redbooks? White Papers? Support for all of these is on this web page. Click Support on DB2 web pages. From this page, you access roughly 900 Technotes or answers to Frequently Asked Questions, 150 technical presentations from recent conferences, a dozen white papers, and over 60 Redbooks. There are more than 50 presentations updated since May 2004 and 75 since April 2004. The Consolidated Service Test (CST) documents IBM recommended service level for a range of z/OS products. To search for the presentation you want, put a word or two into the Additional search terms or Order by date – newest first. There is a new redbook on Disaster Recovery and another being written on V8 Performance Topics.

ibm.com/software/db2zos/support.html

Web Resources: DB2 for z/OS Migration

- ibm.com/software/db2zos
Click on Support and then access the Technotes (Q&A), White Papers & Technical Presentations
- ibm.com/software/data/db2/zos/v8books.html
- ibm.com/servers/eserver/zseries/zos/bkserv/r6pdf/
- ftp sites for papers, presentations:
<ftp://ftp.software.ibm.com/software/data/db2zos/>
<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/>

See specific URLs in the notes below.

consultant, customers why migrate <http://www.ibm.com/software/os/zseries/index.html>

<http://www.ibm.com/software/os/zseries/webcast/aug31/>

<http://www.ibm.com/software/os/zseries/ondemand/#dm>

paper on why migrate <ftp://ftp.software.ibm.com/software/data/db2zos/db2v8dbv.pdf>

value of migration presentation <ftp://ftp.software.ibm.com/software/data/db2zos/Z05.pdf>

John Campbell presentations on health checking and on early customer experience.

<ftp://ftp.software.ibm.com/software/data/db2zos/A01jc.pdf>

<ftp://ftp.software.ibm.com/software/data/db2zos/G11jc.pdf>

<ftp://ftp.software.ibm.com/software/data/db2zos/G12jc.pdf>

migration planning and process <ftp://ftp.software.ibm.com/software/data/db2zos/S1322rma.pdf>

migration planning documents

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/V8planmigrate.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/V8migrateprocess.pdf>

Installation Guide <http://publib.boulder.ibm.com/epubs/pdf/dsnigj12.pdf>

Program Directories

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/v8books/l1085662.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/v8books/l1085851.pdf>

<http://publibfp.boulder.ibm.com/epubs/pdf/i1085672.pdf>

Migration presentations

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OF2.pdf>

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OA2.pdf>

What do customers say about DB2 V8?



What can you learn?

Gain their experiences.

At conferences, user groups, on the web

Univar

<http://www.ibm.com/software/os/zseries/index.html#>

<ftp://ftp.software.ibm.com/software/data/db2zos/db2v8dbv.pdf>

<http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&uid=swg27005465>

GAD

<http://www.ibm.com/software/data/db2imstools/news.html>

<http://www.ibm.com/software/os/zseries/webcast/mar1/>

Deere, system backup and restore

<ftp://ftp.software.ibm.com//software/data/db2zos/PITR18.pdf>

Hewitt Associates

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/OF2.pdf>

Summary of early customer experiences

<ftp://ftp.software.ibm.com/software/data/db2zos/A01jc.pdf>

Health checking

<ftp://ftp.software.ibm.com/software/data/db2zos/G11jc.pdf>

<ftp://ftp.software.ibm.com/software/data/db2zos/G12jc.pdf>

IBM Software Group | DB2 Information Management Software IBM

Events for more info ...

IBM Conferences

IDUG

SHARE

DB2 for z/OS and OS/390

DB2 for z/OS

Features and benefits

System requirements

Library

Success stories

News

How to buy

Events

Training and certification

Services

Support

Related software

- WebSphere
- CICS
- IMS
- DB2 Connect
- z/OS
- Linux on z/Series

Related solutions

- DB2 & IMS Tools
- Java™ on z/OS and OS/390
- DB2 for Linux

Related hardware

- Storage Systems
- z/Series

Related services

- Online Self-Service Support
- IBM Consolidated

Events

DB2 Conferences

Upcoming Conferences

- SHARE - Anaheim
27 Feb - 4 Mar 2005
Anaheim, California
- Software Development technical conference
25-29 April 2005
Nice, France
- IBM DB2 Information Management - Europe
9-13 May 2005
Vienna, Austria
- IDUG 2005 - USA
22-26 May 2005
Denver, Colorado
- SHARE - Boston
21-26 Aug 2005
Boston, Massachusetts
- IBM DB2 Information Management - USA
12-16 Sep 2005
Orlando, Florida
- IDUG Symposium
26-27 Sep 2005
Toronto, Ontario
- IDUG 2005 - Europe
24-28 October 2005
Berlin, Germany

Conference Sponsors

- IBM Learning Services: International Technical Conferences
- IBM DB2 Information Management Conferences
- IDUG (International DB2 Users Group)
- SHARE
- GSE (GUIDE SHARE Europe)

Online Support Library

- Presentations from Previous Conferences
A source for many technical presentations from IBM Developers

Part 1 41

There are many excellent conferences coming up. Use the Events page to see which one you can attend:

- IBM DB2 Information Management - Europe
9-13 May 2005 Vienna, Austria
- IDUG 2005 – USA
22-26 May 2005 Denver, Colorado
- SHARE - Boston
21-26 Aug 2005 Boston, Massachusetts
- IBM DB2 Information Management - USA
12-16 Sep 2005 Orlando, Florida
- IDUG Symposium
26-27 Sep 2005 Toronto, Ontario
- IDUG 2005 - Europe
24-28 October 2005 Berlin, Germany

<http://www.ibm.com/software/data/db2/zos/events.html>



2005 IBM TECHNICAL CONFERENCES

IBM DB2 Information Management Technical Conference

Featuring - DB2, BI, Informix, U2 and WebSphere Information Integration

September 12 - 16, 2005 Orlando, Florida

- News V8 and beyond: DB2 UDB for z/OS, Linux, UNIX, and Windows, Informix Dynamic Server, UniData 7.1, and U2.
- WebSphere Information Integration family.
- Largest assembly of IBM DB2 developers.
- Combined EXPO with a wide range of solution providers.
- Network with peers, get certified, enjoy pre-conference classes.

Enroll by July 25, 2005 and save \$150

USD off Registration Fee of \$1895

To enroll, visit : ibm.com/training/us/conf/db2

IBM, the IBM Logo, DB2, DB2 Universal Database, Informix, and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Part 1 42

DB2 Information Management Technical Conference Slide 3/24/05

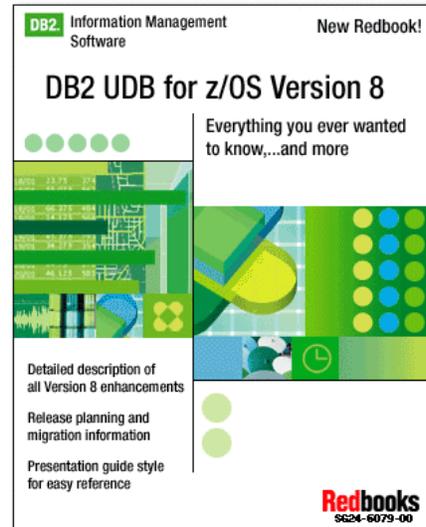
The 2005 IBM DB2 Information Management Technical Conference will give you the opportunity to experience the best of DB2, Business Intelligence, Informix, U2, and WebSphere Information Integration technologies in one dynamic event. Learn how DB2 Information Management Software simplifies and speeds development with embedded capabilities to integrate and analyze information, regardless of its format or location.

- Explore the enhancements in the DB2 information management software portfolio.
- Enhance your knowledge on DB2 UDB for z/OS V8, DB2 UDB for Linux, UNIX, and Windows, Informix Dynamic Server, UniData 7.1, and U2. Preview information about next releases of DB2 UDB including Viper.
- Master the WebSphere Information Integration family.
- Join us for the largest assembly of IBM DB2 information management product developers.
- Enjoy a combined EXPO with the solution providers for the DB2, Business Intelligence, Content Management, Informix, U2, and WebSphere Information Integration
- Network with peers, meet developers, get certified, and enjoy pre-conference classes that enhance your learning experience.

ibm.com/redbooks

New Redbooks:

- ✓ Performance Topics
SG24-6465
- ✓ MLS & DB2
SG24-6480
- ✓ Disaster Recovery,
SG24-6370
- ✓ Others updated



Part 1 43

These new IBM Redbooks are on the web now. Check the web, Redbooks and drafts. Here are the URLs for Performance Topics:

<http://www.redbooks.ibm.com/abstracts/sg246465.html?Open>

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246465.pdf>

MLS and DB2 is a draft at this point, so the URL will change:

<http://www.redbooks.ibm.com/redpieces/abstracts/sg246480.html?Open>

Disaster Recovery published November 2004

<http://www.redbooks.ibm.com/abstracts/SG246370.html?Open>

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246370.pdf>

For all recent updates, start on the Support page, click on Redbooks, sort by date – newest first

<http://www.ibm.com/software/data/db2/zos/support.html>

IBM Software Group | DB2 Information Management Software IBM

www.redbooks.ibm.com/redbooks/pdfs/sq246079.pdf

New Redbooks

- ✓ Everything: Detailed descriptions, about 1100 pages, Presentation Guide style
- ✓ WebSphere SG24-6319

Hot New Workshop!!!

DB2 UDB for z/OS Version 8: Everything you ever wanted to know ...

Featuring all V8 Enhancements and Migration to this new version of DB2 for z/OS

DB2 Information Management Software

DB2 Information Management Software New Redbook!

DB2 UDB for z/OS Version 8

Everything you ever wanted to know,...and more

Detailed description of all Version 8 enhancements

Release planning and migration information

Presentation guide style for easy reference

Redbooks
SG24-6079-00

Part 1 44

This new IBM Redbook introduces the many enhancements made available with DB2 UDB for z/OS Version 8. It will help you understand the enhancements, and provides information to help you to evaluate the applicability to your environment. This book will help you plan for the installation of DB2 V8 or the migration from DB2 V7, with the Library. The pdf version of the book has already been updated several times, with the last on April 11, 2005. Watch for more updates.

<http://www.redbooks.ibm.com/redbooks/pdfs/sq246079.pdf>

<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sq246079.html?Open>

DB2 for z/OS and WebSphere: The Perfect Couple is SG24-6319.

<http://publib-b.boulder.ibm.com/abstracts/sq246319.html?Open>

<http://www.redbooks.ibm.com/redbooks/pdfs/sq246319.pdf>



There is extensive help for porting to DB2 UDB on the web, with suggestions for the process and many resources, including a migration toolkit. The dramatically improved DB2 family compatibility in DB2 for z/OS V7 and V8 makes porting to DB2 UDB for z/OS from Unix or Windows platforms much easier.

<http://www.ibm.com/developerworks/db2/zones/porting/>

Here is the URL for the main DeveloperWorks page

<http://www.ibm.com/developerworks/db2/>

There is also a DeveloperWorks page on DB2 for z/OS.

<http://www.ibm.com/developerworks/db2/products/db2zos/>