



IBM Software Group

DB2 for z/OS V8 and DB2 V9.1 for z/OS

Roger Miller ibm.com/software/db2zos

May 23, 2006

DB2 Information Management Software

ON DEMAND BUSINESS™

We will discuss the latest news about DB2 for z/OS V8 and DB2 V9.1 for z/OS (V9). The V9 beta announcement was May 2, 2006. The beta program begins in early June 2006. This webcast will include a few other new items, such as the new IBM System z9 Integrated Information Processor (zIIP) and the new z9 processors. Data definition on demand extends the theme of online schema revolution from V8. Additional Unicode and text enhancements continue the work from V7 and V8. XML work across the DB2 family is a much larger step than in V7 or V8. SQL Procedures become more consistent across the family. While V7 and V8 removed many differences from DB2 for Linux, UNIX & Windows, V9 takes the next big step to improved productivity and consistency. Utility enhancements help with new function, more LOB and XML support, better performance and improved availability, removing the BUILD2 step from online reorg. V9 enhances DB2's ability to handle new and enterprise applications. V9 improves with XML, large objects, and many SQL and security improvements. V9 builds upon and extends DB2 traditional strengths and the ground-breaking V8 in many areas: online schema evolution, Unicode, XML, DB2 family SQL, utilities, security and 64-bit virtual storage.

Greatest Hits: DB2 for z/OS V8

- ✓ High availability
- ✓ Scalability or very large database
- ✓ Java and the web
- ✓ Queries and data warehouses
- ✓ Migrating or porting applications
- ✓ Application packages



Greatest hits are the situations where I would recommend looking at V8 soon. For these situations, V8 offers significant improvements without adding much effort. Some cases can reduce the work. Value and effort vary widely for different customers and most of the improvements fit more than one category, so this list is not ordered.

Greatest Hit 1: High availability

One of the biggest steps for database administrators in continuous availability is online schema evolution, with the ability to add partitions and make about 20 changes with ALTER. New backup and recovery utilities are useful for disaster recovery and will be the primary backup technique for some customers. Improvements in utilities include more online performance and better usability.

Greatest Hit 2: Scalability or very large databases

Separate partitioning and clustering allows two dimensional clustering with more effective IO. New index options provide more efficient access. The maximum number of partitions is raised to 4096. The availability and optimization improvements are critical for very large databases. The ability to use more memory, more effectively is key for scalability.

Greatest Hit 3: Java and the web

Improvements in the SQLJ and JDBC support, a new Java Universal Driver, enhanced Unicode support, integration with WebSphere and new XML functions make Java and web applications more robust and more productive.

Greatest Hit 4: Queries and data warehouses

Optimization changes provide the best performance improvement opportunities in V8. Faster response and reduced processing time come from improved optimization and better information for the optimizer. New database design options for indexes, clustering and materialized query tables provide more gains. Warehouses often need to have the new rotate partition capability.

Greatest Hit 5: Migrating or porting applications from other platforms

Many SQL enhancements provide better compatibility with the DB2 family and with the industry. If customers develop on Windows, Unix or Linux, and then move to z/OS, the process is much easier. Early customers reported success at porting applications.

Greatest Hit 6: Application packages: SAP, PeopleSoft, Siebel, etc. ...

About 50 improvements, including everything mentioned in the "Greatest Hits" section, are provided for most of the key vendor packages. SAP R/3 4.6 and PeopleSoft PeopleTools 8.45 are already certified for V8, less than four months after general availability, Siebel 7.8 is certified.

Continuous Availability



- **Online Schema Evolution: database changes with ALTER instead of DROP / CREATE e.g. ADD partition**
- **System-Level Log Point Recovery**
- **Data Partitioned Secondary Indexes**
- **Improved LPL Recovery**
- **Additional online zparms**

The most important change for many customers, especially database administrators, is the ability to use ALTER in many places instead of needing to drop and redefine. We call this schema evolution, and it can reduce outages by hours or days for a major structure change on an application.

The ability to have secondary indexes that are partitioned with the data can improve recovery times by an order of magnitude. It can also eliminate the outage for online reorganizing a single partition or BUILD2 phase.

We have some additional cases where subsystem parameters can be changed while the subsystem is running.

<ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/Z03.pdf>

Scalability and Very Large Database

- Add partitions
- Separate partitioning & clustering
- Index improvements
- 4096 Partitions
- Rotate partitions
- Extend columns
- Optimization improvements
- Memory and scale increased



4


Very large databases face the combined challenge of very high performance needs, continuous availability and complexity. Improvements in scale and flexibility are more important in this area. Being able to have more partitions and to add them with ALTER are a big improvement.

Often it is useful to partition by date, so that we can archive or delete an entire partition, but processing will be much more efficient with another clustering order, such as by customer. Before this change, the clustering order was the same as the partitioning. This flexibility offers many opportunities for improved performance and availability.

Some customers have an index that is used only for partitioning the data or have extra columns at the beginning of the index. Being able to avoid the extra index or columns can improve our efficiency a lot.

For these very large tables, the ability to have more partitions, to add new partitions and to be able to rotate partitions is crucial.

Java and the web



- **Application Connectivity for JAVA**
- **Improve data sharing member routing**
- **LOCATION aliases at requester & server**
- **Extended function, standards**
- **Reduced processing**
- **Stored Procedure Enhancements**
- **Java Universal Driver**

DB2 V8 also provides many improvements for Java and the web: security, availability, usability and performance. The improved data sharing member routine is more robust. Having the ability to use multiple names for a server adds flexibility. Increased levels of standards are implemented. Performance improvements will reduce cpu costs substantially. Improved granularity for stored procedures help with managing your work load, using the WLM.

The Java Universal Driver is used across the DB2 family.

Queries and data warehouses

- Optimization Improvements
 - ❑ Improved techniques
 - ❑ Enhanced data
 - ❑ Visual Explain
- Enhanced index options
- Materialized Query Tables
- New Partitioning options
- QMF improvements
- SQL enhancements



Queries and data warehousing are improved a lot in V8. Optimization improvements provide a performance boost and make the job simpler. Improved optimization techniques like ability to use indexes more, star join and scale improvements allow reduced work for computers and for people. Enhanced data helps get the best access path. Visual Explain improves the ability to analyze and resolve any problems.

The many improvements for indexes, materialized query tables and partitioning can save space and add new options for improved performance and availability, even while simplifying the process. Not padded, clustering, longer and backward scans help indexes. Being able to add, rotate and rebalance partitions improve partitioning options. QMF enhancements build upon these strengths and add new function to reporting, dash boards, and a new platform in WebSphere. SQL enhancements on this page and the next improve portability of the SQL, improve the ability to express queries, and help with performance.

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**
- **HISTORY statistics without updating main statistics**

7

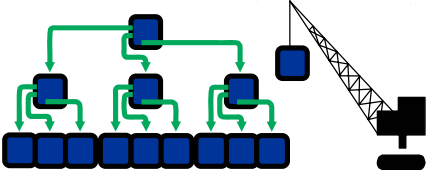
Skewed data distributions are responsible for a high proportion of performance problems with DB2 queries, especially in ad hoc queries. Symptoms can be less than optimal join sequences, too much synchronous I/O, and long response times. When there is asymmetrical distribution of data, not having distribution 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.

Index: DB2 UDB for z/OS

Index Improvements

- 
- Variable length index keys
 - Index-only access for varchar data
 - Maximum index key 2000 bytes
 - Predicates indexable for unlike types
 - Backward Index Scan
 - Partitioning separate from clustering
 - Data-partitioned secondary indexes (DPSI)
 - Create index online during select, insert
 - Add column to index

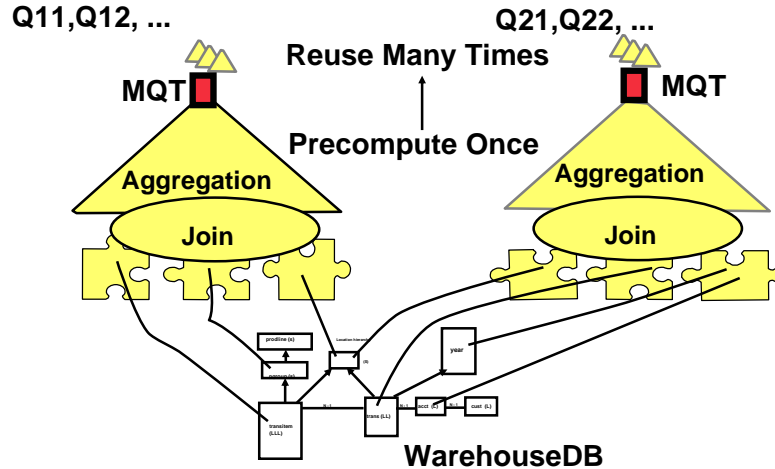
DB2 V8 provides many new opportunities for improving index processing, rebuilding the architecture for indexes.

We are able to use indexes more effectively, reducing the space in variable-length indexes, being able to have index-only access with variable-length data and being able to use the index when the predicates do not match.

In some cases, such as backward index scans or partitioning, we will be able to work as efficiently with one less index. Being able to eliminate an index will improve the insert, delete, LOAD, REORG and update processing.

We have more flexibility in indexes, with longer index keys, the ability to partition secondary indexes and the ability to have more effective clustering.

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.

Productivity, migrating & porting applications



- Multi-row INSERT, FETCH & UPDATE
- GET DIAGNOSTICS
- INSERT within SELECT
- IDENTITY Column enhancements
- SEQUENCES
- CURRENT PACKAGE PATH
- SQL Procedure Language
- Dynamic Scrollable Cursor, Common Table Expression, Recursion
- Scalar Fullselect
- Materialized Query Table
- UNICODE SQL, Multiple CCSIDs
- XML Publishing
- Long names, long statements...

Improvements in the SQL have made migrating from other platforms, such as Unix and Windows much easier. Many of the early customers moved applications.

Version 8 is a breakthrough in SQL, with too many new functions to list them all. We will discuss a few of them. Add support for volatile tables, group by expression, multiple DISTINCT clauses, and qualified names for INSERT and UPDATE and many more to the list above.

Table function improvements, trigger performance, star join sparse index improvements, cost-based parallel sort, better ability to use indexes, longer statements, longer names, longer literals and predicates, session variables, new special registers and SQL procedures improvements, provide a giant leap for SQL.

DB2 UDB Family SQL Reference



Version 2
September 2004

IBM DB2 Universal Database

SQL Reference for Cross Platform Development

z/OS OS/400 AIX HP-UX Solaris Linux Windows

SQL Reference for DB2 UDB family

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

Version 2 adds DB2 UDB for z/OS V8

DB2 UDB for iSeries V5R3

Do you want to design or write applications for the entire DB2 family, rather than for just one of the platforms? If so, you need the IBM DB2 Universal Database SQL Reference for Cross-Platform Development.

Version 2 covers these product versions:

DB2 UDB for Linux, UNIX & Windows V8




DB2 UDB for z/OS V8

DB2 UDB for iSeries V5 R3

SQL consistency across the DB2 family has improved substantially in the past few versions, while significant new common function has been added. SQL that is common to the DB2 UDB relational database products and the SQL 1999 Core standard is much more comprehensive. The 852 page book describes the rules and limits for preparing portable programs for these versions. More is coming.

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

Enterprise Applications & : DB2 for z/OS

 4.6 certified	<ul style="list-style-type: none"> -64 bit virtual storage -Unicode -Schema evolution -System-level backup and recovery 	<ul style="list-style-type: none"> -Longer Table Names & Column Names -Additional statistics -Convert Column Type -Altering CLUSTER Option
	 8.45 certified	<ul style="list-style-type: none"> -Multi-row fetch & insert -Multiple DISTINCT Clauses -Lock Contention on volatile Tables -Fast Retrieval of Most Recent Value
 7.8 certified		

Most of the key items in this version help our key enterprise application partners: SAP, PeopleSoft and Siebel, but also many other applications and customers. The 2 GB address space limit is a major problem for customers who have large dynamic statement cache, many concurrently open data sets and long running units of work. Unicode helps vendors support multinational companies and improve character conversions. These applications often insert and fetch many rows. Using a technique called array fetch & insert or multirow fetch & insert can be as much as 50% less cpu time locally, but improves remote access more. SQL flexibility improvements allow DB2 to be efficient in performance and in productivity for our partners. There are roughly 50 items, with some differences, for each partner. See the book DB2 for z/OS V8: Through the Looking Glass & What SAP Found, SG24-7088 for more detail.

SAP R/3 4.6 certified with DB2 for z/OS Version 8 in 2004. The SAP kernel releases 620 and 640 are also certified with DB2 for z/OS V8. Every current, supported SAP release is currently certified for DB2 for z/OS V8.

PeopleSoft PeopleTools 8.45 were certified for V8 in 2004. Now PeopleTools levels 8.19, 8.20, 8.21, 8.22, 8.44, 8.45, 8.46 and 8.47 of PeopleSoft Enterprise are certified with DB2 for z/OS Version 8.

Siebel 7.8 was certified with DB2 for z/OS V8 in 2005.

Enhanced Unicode Support in Version 8

- DB2 catalog now stored in Unicode
 - SQL is parsed in Unicode, SQL can contain Unicode literals & names
 - provides better integration with Java & Microsoft technologies
- Unicode is native mode for V8, traces
- Built-in functions for character length
- Important to many ISVs
 - ▶ International

The logo for Siebel, featuring the word "SIEBEL" in a bold, blue, sans-serif font.The logo for SAP, consisting of a blue square with a white diagonal line and the letters "SAP" in white below it.The logo for PeopleSoft, with "PEOPLE" in blue and "Soft" in red script below it.The logo for Unicode, with "UN" in large red letters and "ICODE" in smaller black letters below it.

With DB2 for z/OS, Unicode is now the native mode for DB2, improving our ability to be the world-wide enterprise server. All customers will have Unicode, because the DB2 catalog is Unicode. SQL is parsed in Unicode. Java and Microsoft technologies use Unicode to allow communication across the globe without errors in character conversion.

Many of the key application vendors need Unicode support for their applications.

Multirow performance

- **FETCH:**
 - Up to 50% cpu savings
 - DSNTIAUL change made
 - DSNTTEP4 instead of DSNTTEP2
 - Distributed processing
 - QMF APAR PQ99482
- **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 DSNTTEP2, then changing to DSNTTEP4 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. QMF with APAR PQ99482 implements this change. 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.

How much room will you get in DBM1?

It depends.


Any virtual pools?

Buffer pool control blocks

Castout pool

EDM DBDs pool, Sort Pool

RID Pool, compression dictionaries



Some growth in thread storage

You need to calculate, monitor and manage

Some customers get big improvements in virtual storage, but others get very little. If you use little or no compression, have already minimized sort pools, RID pools, castout (no data sharing) and have moved the dynamic statement cache and buffers to data spaces, then the usable storage can shrink. With the new sizes and options in V8, some of the space grows. On the other hand, in about 80% of the cases we looked at, there are some virtual buffer pools left or substantial savings are available from the buffer pool control blocks (estimate 5% of data space and hiperspace buffer pools), data sharing castout, DBDs, sort pool, RID pool (estimate 90%), and compression dictionaries.

You need to do the calculations, monitor and manage storage in V8.

DB2 for z/OS Version 8 is

- ✓ **SQL OLTP Leadership:** name lengths, statements, scrolling, expressions, predicates, diagnostics, ...
- ✓ **Break through limitations:** storage, partitions, log
- ✓ **Performance enhancements:** index, materialized query tables, more efficient IO, multi-row
- ✓ **Database changes without an outage:** add partition, rotate partition, backup / restore
- ✓ **Integration**

System z, z/OS & ESS platform

Middleware stack



Rational.

WebSphere.

DB2

Lotus.

Tivoli

Reengineered for e-business on demand

17


Version 8 is the twelfth and largest ever release of DB2 for z/OS. It brings new synergy with the zSeries hardware and uses the z/OS 64-bit virtual addressing capabilities. V8 improves data support, application development, and query function enhancements for e-business. It also builds on the traditional zSeries and DB2 characteristics of availability, exceptional scalability, and performance for the enterprise database management system of choice. V8 has been re-engineered for e-business on demand, with many fundamental changes in architecture and structure. Key improvements enhance scalability, application porting, security, architecture, and continuous availability. Management for very large databases is made much easier, while 64-bit virtual storage support makes management simpler and improves scalability and availability. This new version breaks through many old limitations in the definition of DB2 objects, including SQL improvements, online schema evolution, longer names for tables and columns, longer SQL statements, enhanced Java and Unicode support, enhanced utilities, more log data sets, more partitions, and many more advantages. Customers, vendors and consultants tell us that DB2 for z/OS Version 8 is exciting for them. It is a very important milestone.

Version 8 includes dozens of changes in SQL, improving family consistency in many cases, playing leapfrog in others, pushing DB2 SQL beyond current boundaries for enhanced application portability, open standards. Longer names for tables and columns mean that customers can use more meaningful names, matching standards. Longer SQL statements help with SQL that is generated or used in an SQL procedure. Here are some: multi-row INSERT, FETCH & UPDATE, GET DIAGNOSTICS, INSERT within SELECT, IDENTITY Column enhancements, SEQUENCES, CURRENT PACKAGE PATH, Dynamic Scrollable Cursors, Common Table Expressions, Scalar Fullselect, Materialized Query Tables, Unicode SQL, XML Publishing and much more. These changes improve our customer productivity, consistency across the DB2 family and ability to port applications. If you want to design or write applications for the entire DB2 family, then use the IBM DB2 Universal Database SQL Reference for Cross-Platform Development. Many barriers that limit our customers are removed: using 64 bit memory, providing consistent table and column name lengths, allowing 2 megabyte SQL statements, 4096 partitions, and three to ten times the log space. Customers encountered many limitations over the past 20 years, and lifting the limits required extensive reengineering for some. Other limits allow improved scalability and availability, such as tripling the active log size and ten times the archive logs. Allowing sixteen times the number of partitions allows a table to use one partition per day for eleven years.

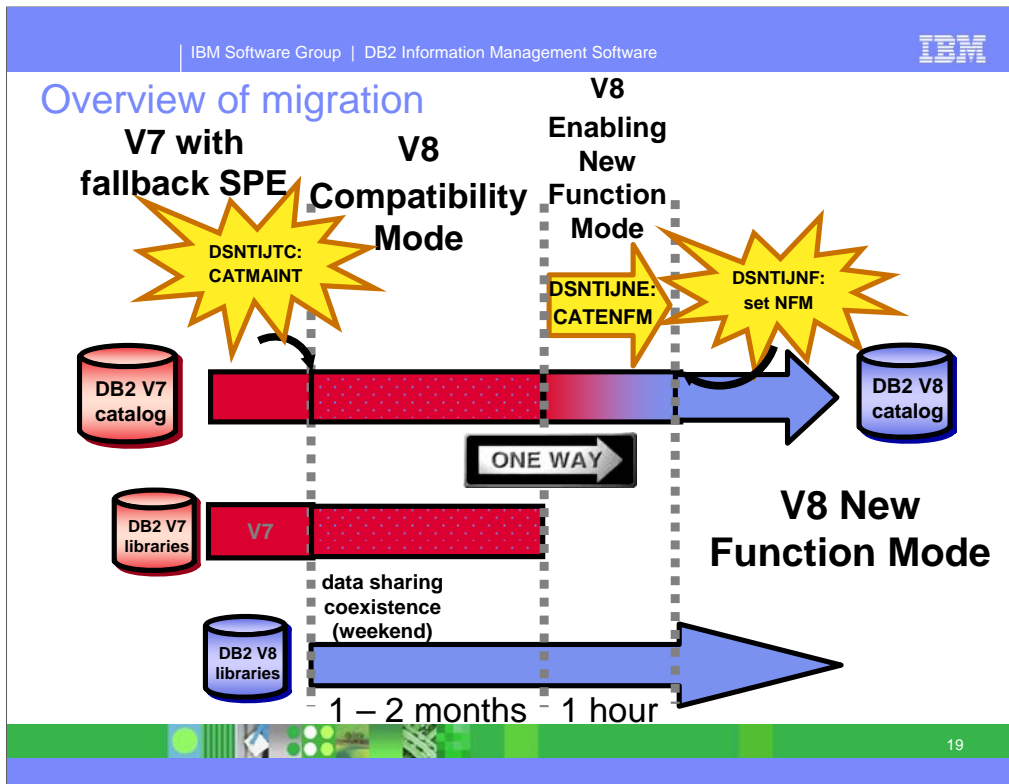
Key performance enhancements deliver better family consistency and run many times faster. Being able to make database changes without an outage, such as adding a partition, is a breakthrough for availability. Alter your table and go, no need to drop and redefine. Online schema evolution is more resilient. The most important change for many customers is the ability to use ALTER in many places instead of needing to drop and redefine. We call this online schema evolution, and it can reduce outages by hours or days for a major database structure change. Database administrators can add a partition to an existing partitioned table space or rotate the partitions. Other changes in online schema evolution allow better partitioning and improved disk access, avoiding random access with more effective database designs. Many enhancements improve our integration with zSeries, z/OS and with key vendor applications like PeopleSoft, SAP and Siebel.

DB2 for z/OS Version 8 News

- now
- Cross loader with LOBs
 - Built in functions ASCII, TIMESTAMPDIF
 - DSN1COPY with 1000 OBIDs
 - QMF with multirow fetch
 - Online Check Index
 - z/OS 1.7 up to 7257 extents
 - LOAD, UNLOAD with LOBs
- soon
- IBM System z9 Integrated Information Processor (zIIP)
 - New and updated books: Library refresh Feb. 2006
 - Messages, Codes separate books 2005
 - Redbooks: Design Guidelines for High Performance and Availability, Business Value, Performance Topics, WebSphere, MLS, Disaster Recovery, others updated ...




- Cross loader with LOBs V7 V8: PQ90263
- Built in functions ASCII, TIMESTAMPDIF V8: PQ95795
- DSN1COPY with 1000 OBIDs V7 V8: PK05758
- QMF with multirow fetch PQ99482
- Online Check Index V8: PQ92749 & PQ96956
- z/OS 1.7 up to 7257 extents V7 V8 PK07590 PK10594 OA12434
- LOAD, UNLOAD with LOBs V7 V8: PK22910 closed
- zIIP PK18454, PK19920, PK19921 open, ... <http://www.ibm.com/systems/z/feature012406/>
- ✓ DB2 for z/OS and WebSphere: The Perfect Couple, SG24-6319
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246319.pdf>
- ✓ Planning for Multilevel Security & Common Criteria (GA22-7509)
<http://publibz.boulder.ibm.com/epubs/pdf/e0z2e122.pdf>
- ✓ Multilevel Security & DB2 Row-Level Security Revealed, SG24-6480
<http://www.redbooks.ibm.com/redpieces/pdfs/sg246480.pdf>
- ✓ Ready for Java <http://www.redbooks.ibm.com/redbooks/pdfs/sg246435.pdf>
- ✓ Large Objects <http://www.redbooks.ibm.com/redbooks/pdfs/sg246571.pdf> update coming
- ✓ Stored Procedure <http://www.redbooks.ibm.com/redbooks/pdfs/sg247083.pdf> upd Feb 2006
- ✓ Cross-Platform SQL Reference V2 Sept 2004
<http://ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>
- ✓ Disaster Recovery with DB2 for z/OS , SG24-6370
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246370.pdf>
- ✓ DB2 for z/OS V8 publications update on web new in February 2006
<http://www.ibm.com/software/data/db2/zos/v8books.html>
- ✓ DB2 for z/OS Performance Topics SG24-6465
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246465.pdf> update Feb 2006
- ✓ Design Guidelines for High Performance and Availability, SG24-7134
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247134.pdf> update Feb 2006



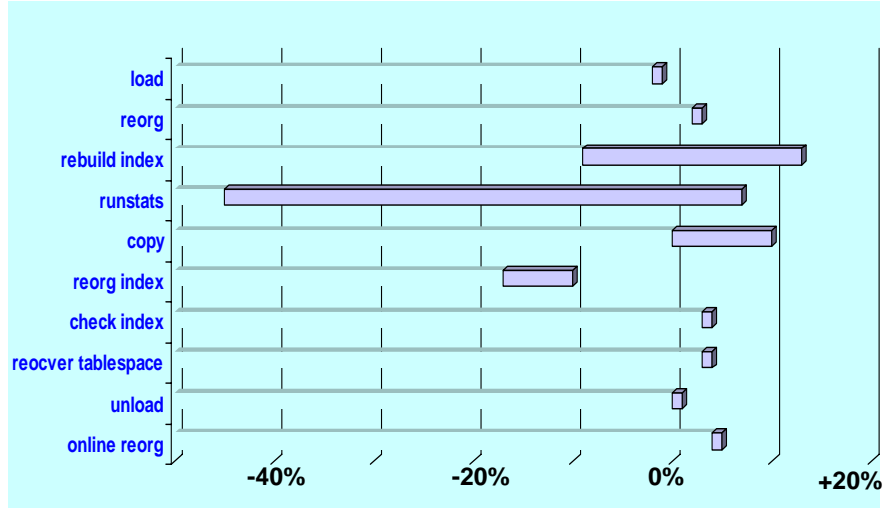
This is an overview of the migration process. There are two steps which modify the catalog, CATMAINT to change into Compatibility Mode (CM), and also CATENFM for Enabling New Function Mode (ENFM), which contains the major catalog changes to long names and Unicode, using a technique of an online (read 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.

CPU change - continued



Utility CPU - continued



Performance Regression

Expect some cpu performance regression

- ✓ 64 bit DBM1 virtual storage constraint relief
- ✓ Long names, long index keys
- ✓ Longer and more complex SQL statements
- ✓ Unicode catalog & SQL

No change in I/O time

→ **Plan to take advantage of V8 improvements**

Page fix, REBIND, multirow fetch & insert, ...

For applications not taking any advantage of V8 performance enhancements, some CPU time increase is unavoidable to support a dramatic improvement in user productivity, availability, scalability, portability, family consistency.

Taking advantage of the V8 performance improvements is very important to avoid regressions and to offset them. Some are relatively simple, and some have been implemented already in products like QMF and multirow fetch.

See the details in chapters 2 and 4 of SG24-6465, **DB2 UDB for z/OS Version 8 Performance Topics**.

CPU Performance Regression: Subject to change YMMV

- Performance objective: under 10% average
- Typical customer 5 to 10% higher average
- Work load variations
 - ▶ 0 to +15% online transaction
 - ▶ -5 to +10% transaction in data sharing
 - ▶ -5 to +20% batch
 - -5 to +5% insert
 - +5 to +20% fetch, select, update
 - ▶ -10 to +15% batch data sharing
 - ▶ -20 to +15% batch DRDA
 - ▶ -5 to +10% utility
 - ▶ -20 to +15% query

Performance objective is less than 10% average regression.

Typical customer workload regression is expected to be 5 to 10% higher on average, differing by work load: Here are some examples:

0 to +15% online transaction

-5 to +10% transaction in data sharing

-5 to +20% batch

-5 to +5% insert

+5 to +20% fetch, select, update

-10 to +15% batch data sharing

-20 to +15% batch DRDA

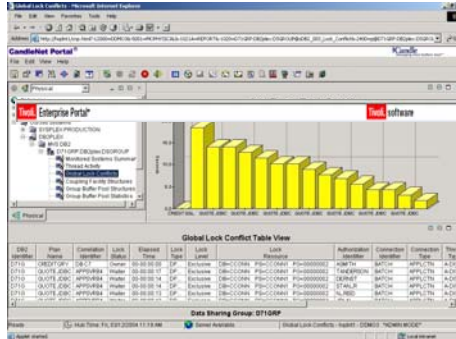
-5 to +10% utility

-20 to +15% query

■ Options with significant potential to offset an expected increase include multirow fetch, multirow insert, long term page fix and rebind.


Tivoli OMEGAMON XE for DB2 PE on z/OS 3.1.0

- Improved ability to monitor and manage mainframe based applications through a single integrated solution
- Familiar interfaces from DB2 PE and OMEGAMON XE products provides easy migration
 - DB2 z/OS V8 new function
 - DB2 Connect reporting / monitoring
 - Performance warehouse (historical data mining)
 - DB2 to CICS transaction linking
 - History monitoring
 - Event exceptions
 - Threshold checking



<http://www.redbooks.ibm.com/redpieces/abstracts/sq247224.html>

<http://www.ibm.com/software/tivoli/products/omegamon-xe-db2-peex-zos>



We merged two leading products and we also merged the development organizations. IBM Tivoli OMEGAMON XE for DB2 Performance Expert and Performance Monitor became generally available in December 2005. The Tivoli naming notes the tight integration with the rest of the suite of integrated monitors from IBM. In this new version you will see the same user interfaces that were available in the OMEGAMON XE for DB2 product as well as in the DB2 Performance Expert and DB2 Performance Monitor. This provides an easy migration path to the converged product.

You get the ability to monitor DB2 Connect, allowing you to track the communications between DB2 Connect and DB2. This is critical in this day of more and more distributed clients. It gives a complete picture of your end to end performance and resource usage. Many of you I'm sure have experienced situations where resource or performance issues have been difficult to diagnose in terms of origin. Is it the web server? Is it DB2? Is it the network? This tool can save you time and money in quick problem diagnosis

You get the value of a performance warehouse for additional long term analysis of your performance data. You can query or take action (set automated tasks) against this warehouse. You get historical reporting – the near term history function continuously captures reporting data. A good example would be using this information to review the performance of threads that have ended. You also get tremendous flexibility in setting alerts and taking action for event exceptions based upon thresholds you select. Using this tool you can identify system bottlenecks such as locking conflicts and deadlocks fast and easy.

This is a great story for our users and new users. Be sure to look at this product if you have one or both products or if you just want the best DB2 performance monitor. See the new redbook and the web page for much more detail.

<http://www.redbooks.ibm.com/redpieces/abstracts/sq247224.html>

<http://www.ibm.com/software/tivoli/products/omegamon-xe-db2-peex-zos>

System z Synergy & DB2 V9

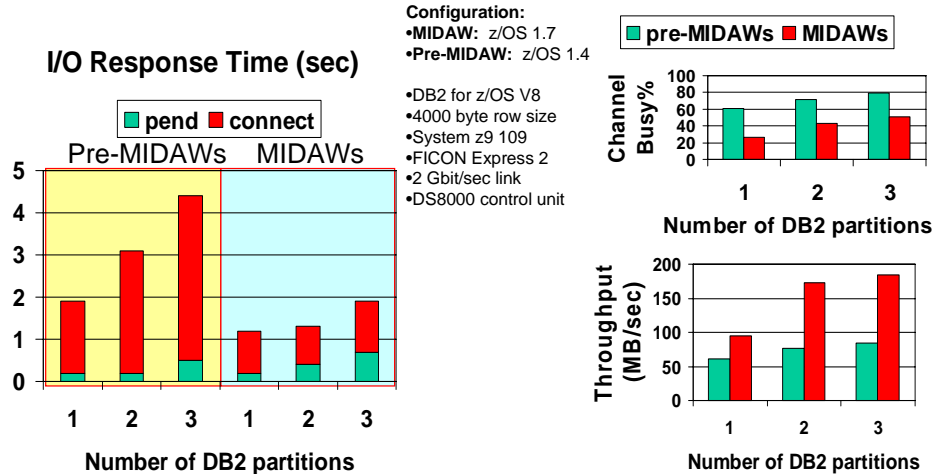
- ✓ **System z9 Integrated Information Processor (zIIP) Enterprise Class & Business Class**
- ✓ **Enhanced Cryptography**
- ✓ **Channels (4 Gb & MIDAW)**
- ✓ **Faster Processors**
- ✓ **Up to 54 Processors EC**
- ✓ **More memory, better value; 64 bit virtual storage**
- ✓ **z/Architecture new instructions**
- ✓ **Parallel Sysplex**
- ✓ IPv6
- ✓ SSL
- ✓ Java
- ✓ Decimal float
- ✓ Backup & restore
- ✓ Security
- ✓ Unicode collation
- ✓ Compression
- ✓ System z Application Assist Processor (zAAP)
- ✓ WLM enhanced ...



The latest System z9 processor improvements for DB2 are the zIIP and the new Business Class and Enterprise Class processors. DB2 V9 remote native SQL procedures are enabled for zIIP processing. V9 adds IPv6, SSL and decimal float and BIGINT data types, with enhancements for Parallel Sysplex, backup and restore, added security and encryption, more Unicode with collation, and uses the WLM in new ways. Channel enhancements (MIDAW) and improved DS8000 performance were included with the System z9 announcements. DB2 uses the latest improvements in hardware and operating system to provide better performance, improved value, more resilience and better function.

DB2 benefits from large real memory support, faster processors, specialty engines, and better hardware compression. DB2 uses Parallel Access Volume and Multiple Allegiance features of the IBM DS8000 and Enterprise Storage Server™. FlashCopy® can be used for DB2 backup and restore. DB2 makes unique use of the z/Architecture™ instruction set, and recent instructions provide improvements in reliability, performance and availability. DB2 continues to deliver synergy with hardware data compression, FICON™ (fiber connector) channels, disk storage, advanced networking function, and Workload Manager (WLM).

Parallel DB2 Table Scan, EF 4K (single channel)



26

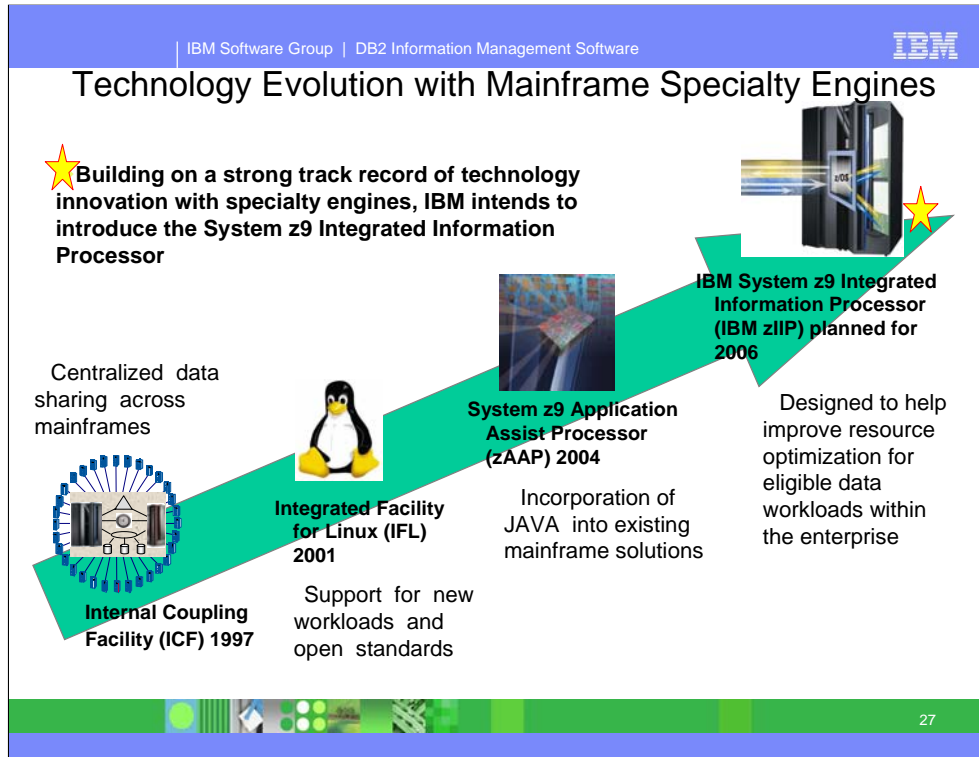
These are some measurements with relatively new processors, new software, new channel configurations and new disks. Note the sustained scan rates of 100 megabytes per second or 170 MB / sec. for parallel access on one channel.

Configuration: MIDAW: z/OS 1.7 Pre-MIDAW: z/OS 1.4

DB2 for z/OS V8 4000 byte row size System z9 109

FICON Express 2 2 Gbit/sec link DS8000 control unit

This document contains performance information. Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the numbers stated here.



System z9 zIIP web site, FAQs, press release

<http://www.ibm.com/systems/z/feature012406/>

http://www.ibm.com/systems/z/faq/pdf/ziip_faq.pdf

<http://www.ibm.com/press/us/en/pressrelease/19157.wss>

Webcast on IBM System z9, zIIP, DB2 V9

<http://www.ibm.com/servers/systems/z/2006/>

Link to the white paper, Why Data Serving on the Mainframe:

<http://www.ibm.com/systems/z/feature012406/whitepaper.html>

Articles in ESJ, ComputerWorld, ADT, NetworkWorld

<http://www.esj.com/news/article.aspx?EditorialsID=1603>

<http://www.esj.com/news/article.aspx?EditorialsID=1647>

http://www.computerworld.com/hardwaretopics/hardware/mainframes/story/0,10801,108080,00.html?source=NLT_ERP&nid=108080

<http://www.adtmag.com/article.asp?id=17854>

<http://www.networkworld.com/news/2006/020606-ibm-db2.html?nettx=020606netflash&code=nlntflash21301>

Blog discussion by Willie Favero, an IBMer on the DB2 team:

<http://blogs.ittoolbox.com/database/db2zos/archives/007533.asp>

New IBM System z9 Integrated Information Processor (IBM zIIP)

- New specialty engine for the System z9 mainframe (planned for 2006) designed to help:
 - Customers integrate data across the enterprise
 - Improve resource optimization and lower the cost of ownership for eligible data serving workloads
- z/OS manages and directs work between the general purpose processor and the zIIP
 - Number of zIIPs per z9 not to exceed number of standard processors. No changes anticipated to DB2 for z/OS V8 applications
- DB2 for z/OS V8 will be first user of the zIIP with
 - System z9
 - z/OS 1.6 or later
 - DB2 for z/OS V8
- Webcast replay ibm.com/servers/systems/z/2006/



The System z9 Integrated Information Processor (zIIP) is the latest customer inspired specialty engine planned for the IBM System z9 mainframe. The zIIP is designed to help improve resource optimization and lower the cost of portions of eligible workloads, enhancing the role of the System z9 mainframe as the data hub of the enterprise.

The zIIP's execution environment will accept eligible work from z/OS, which will manage and direct the work between the general purpose processor and the zIIP. DB2 for z/OS V8 will exploit the zIIP capability for portions of eligible workloads. The zIIP will be available on the System z9 mainframe, and its introduction can help increase the value that customers may derive from the System z9 mainframe over previous generations of the IBM mainframe.

In addition to exploiting zIIP, IBM plans for future versions of DB2 to include functional enhancements that further exploit the mainframe capabilities, including areas such as security, application development, usability and performance.

For more, see the web <http://www.ibm.com/systems/z/feature012406/>
<http://www.ibm.com/servers/systems/z/2006/>

DB2 V8 and IBM zIIP can add value to database work

- Portions of the following DB2 for z/OS V8 workloads may benefit from zIIP*:

- 1 - ERP, CRM, Business Intelligence or other enterprise applications
 - Via DRDA over a TCP/IP connection (enclave SRBs, not stored procedures or UDFs)



- 2 - Data warehousing applications*
 - Requests that utilize star schema parallel queries
- 3 - DB2 for z/OS V8 utilities LOAD, REORG & REBUILD*
 - DB2 utility functions used to maintain index maintenance structures

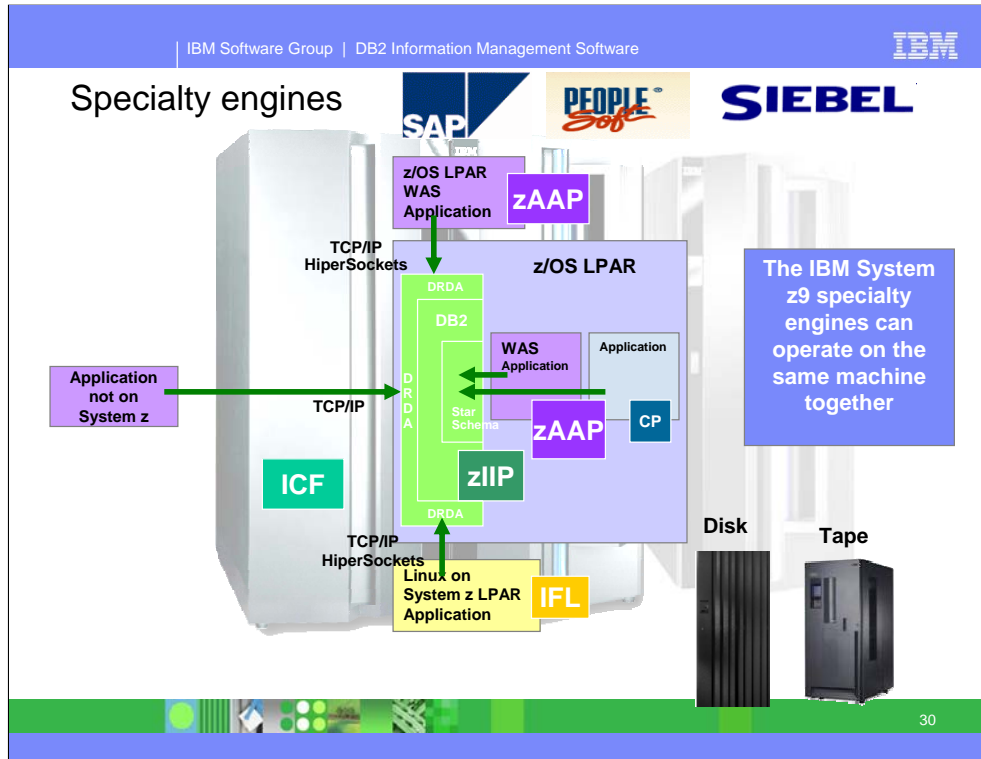
* The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP.

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Not all of this work will be run on zIIP. z/OS will direct the work between the general processor and the zIIP. The zIIP is designed so a software program can work with z/OS to dispatch workloads to the zIIP with no anticipated changes to the application – only changes in z/OS and DB2.

IBM DB2 for z/OS version 8 will be the first IBM software able to take advantage of the zIIP. Initially, the following workloads can benefit:

- SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP), CRM (Siebel), or business intelligence and are expected to provide the primary benefit to customers. Stored procedures and UDFs run under TCBs, so they are not generally eligible, except for the call, commit and result set processing. V9 remote native SQL Procedure Language is eligible for zIIP processing.
- BI application query processing utilizing DB2 star-schema parallel query capabilities; and functions of specified DB2 utilities that perform index maintenance structures.

For more, see <http://www.ibm.com/press/us/en/pressrelease/19157.wss>



The various specialty engines can be used to improve the cost of ownership, providing a low price for the hardware and not incurring software charges, since they are not general purpose processors.

A vision for System z advanced data serving

System z Enterprise Hub for Mission Critical Data

- ❖ With a strong foundation for transaction processing, built on 40+ years of technology innovation, System z servers with z/OS and DB2 can provide a premier platform for data serving, today and into the future*
- ❖ IBM plans to continue to invest in new solutions to address customers' strategic information on demand goals*



Today's Capabilities

- Industry-leading data integrity and security
- Data sharing solution for centralized view of data
- Scalability and availability for enterprise class workloads
- Comprehensive systems and data management environment



*Extension of capabilities**

- New specialty engine (zIIP) with DB2 exploitation - for mission critical ERP, CRM, and Data Warehousing workloads *
- Database support improves regulatory compliance and autonomies
- Support of encryption capability (tape subsystem) with z/OS centralized key mgmt
- Data protection to achieve highest levels of security certifications




*Future direction**

- Additional zIIP exploitation
- DB2 enhancements to help improve usability and reduce complexity and management costs.
- DB2 table scan acceleration via DS8000
- Support of encryption capability (disk subsystem) with z/OS centralized key mgmt
- Handle larger volumes of data, with improved scalability

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

DB2 V9.1 for z/OS – Addressing corporate data goals

- Improved IT Infrastructure In Support of Compliance Efforts
 - Trusted security context
 - Database roles
 - Auditing capabilities
 - Encryption improved
- Simplify development and porting
 - Many SQL improvements that simplify
 - Native SQL stored procedures
 - Default databases and table spaces
 - Automatic unique indexes to support primary keys
- Decrease Complexity and Cost
 - Fast table replacement
 - Partition by growth
 - Table append
 - Volume-based COPY/RECOVER
 - Optimization Service Center
- Evolve Your Environment & SOA
 - Integrated XML
 - WebSphere® integration



DB2 V9 z/OS Innovation: Insight

- ❑ **Business Insight**
 - ❑ Hybrid data server: XML and relational
 - ❑ Enhanced SQL 
 - ❑ Business Partner improvements
 - ❑ QMF cross-platform work station & web
- ❑ **Cost Savings through Optimization**
- ❑ **Business Resiliency**

We have grouped the innovations into three sets, those which provide insight for the business, those which provide cost savings, and those which improve business resiliency. The first set is


BUSINESS INSIGHT INNOVATIONS

- o Rich new hybrid data server support for both relational and pure XML storage, with the necessary services to support both data structures
- o Database on demand capabilities to ensure that your information design can be changed dynamically, often without database outages
- o Rich new SQL capabilities, including additional data types and built-in functions, expanded SQL commonality with DB2 family, and enriched text handling with functions like caseless comparisons, cultural sort, and index on expression
- o Enhanced support to allow key business partners, especially with ERP offerings like SAP, to continue to push the envelope with capabilities that drive business value to their customers. Specific items include new autonomic features, optimization and availability enhancements, and support for new data types
- o Completely redesigned, cross-platform workstation and web-based QMF query solution, providing on demand access to data, reports and interactive visual solutions using a rich desktop application or ordinary web browser



Capabilities Inside the Engine

Performance, Performance, Performance



Native storage Schema Index Functions Utilities

Powerful querying and transformation capabilities

Querying is merely finding a document that matches a certain criteria. For e.g., find all purchase orders with a certain order id.

More often than not, users want to transform the data into something more relevant. For example, they may want to not return all purchase orders as-is, but instead return documents that only contain the purchase order id, and the amount of the order.

DB2 already provides great support for the first two paradigms. In fact, with our XML extender, we were the first ones in the industry to support XML. We're working on completing the picture, across the DB2 family. DB2 for Linux, UNIX and Windows Viper will deliver this picture, and add an XQuery interface to the data.

SQL: Productivity, DB2 family & porting



- XML
- MERGE & TRUNCATE
- SELECT FROM UPDATE, DELETE, MERGE
- INSTEAD OF TRIGGER
- BIGINT, VARBINARY, BINARY, DECIMAL FLOAT
- Native SQL Procedure Language
- Nested compound
- Optimistic locking
- LOB File reference variable & FETCH CONTINUE
- FETCH FIRST & ORDER BY in subselect and fullselect
- INTERSECT & EXCEPT
- ROLE & trusted context
- Many new built-in functions, caseless comparisons
- Index on expression
- Improved DDL consistency
- CURRENT SCHEMA

As in Version 8, there are many improvements for SQL and for XML in V9. Improvements in the SQL have made migrating from other platforms, such as Unix and Windows much easier.

V9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on the prior slide. There are new SQL data manipulation statements in MERGE and TRUNCATE. There are new data types with DECIMAL FLOAT, BIGINT, BINARY and VARBINARY types. Improvements in LOBs provides more consistent handling and improved performance. Intersect and Except set operations make some SQL operations simpler to specify. Security is improved with ROLES and network trusted context. Data definition consistency and usability are improved. V9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

IBM Software Group | DB2 Information Management Software IBM

DB2 SQL

z z/OS V7
common
LUW Linux, Unix & Windows V8.2



z { Range partitioning

c
o
m
m
n { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, Call from trigger, statement isolation

L
U
W { Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE, DELETE & MERGE, multi-site join, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

36

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS and OS/390 Version 7, comparing a March 2001 z/OS version with an October 2004 LUW version. V7 has almost no unique function, there is a small set of common function, and a larger set of SQL unique to LUW.

The next step in the process is DB2 for z/OS Version 8. There are three sets of SQL noted above, with none that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

IBM Software Group | DB2 Information Management Software IBM

DB2 SQL

z z/OS V8
common
LUW Linux, Unix & Windows V8.2



Z	{	Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning
c o m m o n	{	Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT
L U W	{	Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE, DELETE & MERGE, multi-site join

37

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS, comparing the z/OS Version 8 from March 2004 with the LUW version from October 2004.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

If you want to improve DB2 family consistency, then DB2 for z/OS Version 8 is a big step, changing the game from one of catch up to one of leapfrog.

IBM Software Group | DB2 Information Management Software IBM

DB2 SQL

z z/OS V9
common
LUW Linux, Unix & Windows Viper



z { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE

c { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE, DELETE & MERGE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitioning

l { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, XQuery

38

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 V9 for z/OS, (V9). V9 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. At about this time we'll also have a new release of DB2 V9.1 for LUW, code named Viper. We are able to move more from the z list to the common list with Viper.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

Native SQL Procedural Language

- Eliminates generated C code and compilation
- Fully integrated into the DB2 engine
- Extensive support for versioning:
 - VERSION keyword on CREATE PROCEDURE
 - CURRENT ROUTINE VERSION special register
 - ALTER ADD VERSION
 - ALTER REPLACE VERSION
 - ALTER ACTIVATE VERSION
- BIND PACKAGE with new DEPLOY keyword



Native SQL stored procedures

Stored procedures written in SQL procedure language enhance portability and ease of use when using DB2 for z/OS as your enterprise information source. This language is an ANSI standard language. It is similar to the proprietary stored procedure languages of several competitive databases, which assists in migrating and porting to DB2 for z/OS. This is very similar to the changes in DB2 for LUW V8.2.

SQL stored procedures are supported by the DB2 Development Center tooling, providing an environment to code, test, and debug modules from your connected workstation. This language is currently converted to C when the CREATE PROCEDURE statement is executed. The C program is then automatically prepared, compiled, linked, and bound. The developer does not need to work with the C code.

SQL stored procedures code will be natively integrated into the DB2 engine, eliminating the conversion to C. Additionally, extensions to the bind command will allow for the promotion of the program and access paths between environments without needing to recreate the stored procedure.

Native SQL stored procedures are eligible to have a portion run on zIIP processors when they are invoked from a remote client.

Business Partners: Many enhancements

- SHRLEVEL(REFERENCE) for REORG of LOB table spaces
- Online RENAME COLUMN
- Online RENAME INDEX
- Online CHECK DATA and CHECK LOB
- Faster REORG by intra-REORG parallelism
- More online REORG by eliminating BUILD2 phase
- LOB Lock reduction
- Skipping locked rows option
- Online REBUILD INDEX
- Change SCHEMA & VCAT
- Tape support for BACKUP and RESTORE SYSTEM utilities
- Recovery of individual tablespaces and indexes from volume-level backups
- Enhanced STOGROUP definition
- Utility TEMPLATE switching
- Conditional restart: automatic search for appropriate checkpoint
- CLONE Table: fast replacement of one table with another
- Buffer management by WLM
- Global query optimization
- Generalizing sparse index and in-memory data caching method
- Optimization Service Center
- Autonomic reoptimization
- Logging enhancements
- LOBs Network Flow Optimization
- Faster operations for variable-length rows
- NOT LOGGED table spaces
- Index on expressions
- Universal Table spaces
- Partition-by-growth table spaces
- APPEND option at insert
- Autonomic index page split
- Index page sizes 8K, 16K, 32K
- Support for optimistic locking
- Faster and more automatic DB2 restart
- MODIFY RECOVERY enhancements
- RLF improvements for remote application servers such as SAP
- Preserving consistency when recovering individual objects to a prior point in time
- DECIMAL FLOAT, BIGINT
- VARBINARY, BINARY
- TRUNCATE TABLE statement
- MERGE statement
- FETCH CONTINUE
- ORDER BY and FETCH FIRST n ROWS in sub-select and full-select
- ORDER OF extension to ORDER BY
- Various scalar functions
- XML support in DB2 engine
- Native SQL Stored Procedures, able to use zIIP
- SELECT FROM UPDATE/DELETE/MERGE
- Enhanced CURRENT SCHEMA
- IPv6 support
- Unified Debugger
- Network Trusted Context
- Database ROLES
- Automatic creation of database objects
- Modify early code without requiring an IPL
- Utilities CPU reduction
- Temporary space consolidation
- ...

TCO Reduction

Continuous Availability

Performance

Scalability

SQL


Portability

40

Most of the key items in this version help our key enterprise application partners: SAP, PeopleSoft and Siebel, but also improve many other applications and customers. Customers working on the web and Service Oriented Architecture (SOA) see most of these benefits too. SQL flexibility improvements allow DB2 to be efficient in performance and in productivity for our partners. There are more than 50 items for each partner, improving the performance, scalability, continuous availability, SQL and portability. So the net is an improvement in the total cost of operation.

QMF

- Drag and drop querying, reporting and analytics
- Executive dashboards and data visualization
- Enhanced OLAP with DB2 Cube Views
- Reengineered cross-platform development environment
- New security model for access control & personalization
- Enabled for WebSphere Application Server



41

DB2 QMF V9 delivers: Drag and drop querying, reporting and analytics; Executive dashboards and data visualization solutions; Enhanced support for OLAP with DB2 Cube Views; Reengineered cross-platform development environment; New security model for access control and personalization; Enabled for WebSphere Application Server

Overview: DB2 Query Management Facility (QMFTM) on demand feature includes:

Support for DB2 UDB V9, including XML and enhancements to SQL

Drag and drop development of OLAP analytics, SQL queries, tabular reports, graphical reports, pivot tables and data analysis views

Drag and drop development of interactive visual solutions such as executive dashboards, capable of graphically rendering data simultaneously drawn from multiple data sources

Rich security model supports the personalization and dissemination of data and reports on a per-user/group basis

New object repository simplifies the storage, access control and distribution of QMF objects across the enterprise

Reengineered Eclipse-based desktop application extends QMF to numerous workstation operating systems including Windows, Linux, Solaris and AIX

Former QMF Visionary solution fully assimilated into core QMF product, extending the Visionary graphical objects to QMF reports and unifying the development of visual solutions within a single application

Redesigned QMF for WebSphere extends key QMF functionality to browser-based users across multiple platforms

With this release, DB2 QMF offers a completely redesigned, cross-platform workstation and web-based solution, providing on demand access to data, reports and interactive visual solutions using a rich desktop application or ordinary web browser.

DB2 V9 z/OS Innovation: Cost

- ❑ Business Insight
- ❑ Cost Savings through Optimization
 - ❑ Security and Regulatory Compliance
 - ❑ Performance improvements
 - ❑ Synergy with System z
 - ❑ Query enhancements
- ❑ Business Resiliency




The second category of innovation is

COST SAVINGS THROUGH OPTIMIZATION

- o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing
- o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance
- o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MIDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)
- o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

Regulatory Compliance Characteristics

- **Mandatory:** Agencies may require qualification and/or adherence to specific standards
- **Non-directive:** Usually do not specify exact steps needed to comply
- **Non-certifiable:** Agencies generally do not approve, recommend, or validate solutions
- **Continuously changing:** Regulations, their interpretation, and their codification into corporate policies change frequently
- **Increasing IT impact:** While few apply directly to IT, regulations increasingly affect IT systems
- **Intrinsically related to risk:** Agencies now recommend a risk-based management approach



43

Managing compliance with regulations is challenging. There are many different regulations, and compliance is often a condition of staying in business and out of jail. The regulations don't specify what you need to do, so interpreting the rules is the next challenge. Since specific solutions are not specified, recommended or validated, these tasks are added to the compliance effort. If you are in compliance, then the regulations often change, so that a new effort is needed.

These regulations are having more and more impact upon information technology, and most of the impact is indirect, compliance with business rules. Compliance is related to many areas of the business: security, privacy, government relations and audit, but the essence of the work is related to risk management.

The key demand for DB2 is increased flexibility, with robust security and audit capabilities.

Regulatory Compliance in DB2 V9 for z/OS

Key implementations

- Roles
- Network Trusted Contexts
- Improved auditing
- Data Encryption
- Secure Socket Layer
- Instead of triggers



While DB2 for z/OS V8 provides many enhancements for security, there are still many more needs and much more work to do. Roles are used to provide a more flexible technique than groups or users in assigning and controlling authorization, while improving consistency with the industry. A network trusted context provides a technique to work with other environments more easily, improving flexibility. The instead of trigger is an SQL technique that allows a trigger to be used in place of a view, consistent with DB2 for Linux, UNIX and Windows. Improved audit selectivity is needed for being able to see that security is functioning. Secure Socket Layer or SSL implementation provides encryption of data on the wire. Some additional techniques for data encryption will help protect data at rest and in backups.

Database ROLES

- ROLE is a “virtual authid”
 - Assigned via TRUSTED CONTEXT
 - Provides additional privileges only when in a trusted environment using existing primary AUTHID.
 - Can optionally be the OWNER of DB2 objects

```
CREATE ROLE PROD_DBA;  
GRANT DBADM ... TO PROD_DBA;  
  
CREATE TRUSTED CONTEXT DBA1 ...  
  DEFAULT ROLE PROD_DBA OWNER(ROLE);
```

Database role: A database role is a virtual authorization ID that is assigned to the user via the context mentioned next. DB2 privileges are assigned to the defined role.

The role exists as an object independent of its creator, so creation of the role does not produce a dependency on its creator.


This capability can allow a DBA to have privileges to create objects and manage them for a time, even though ownership is to be another id.

The role can be assigned and removed from individuals via the trusted authorization context as needed. This allows a DBA to perform object maintenance during a change control window on a Saturday night, for example. But when Monday arrives, they do not have the authority to do this same work.

Auditing trails of the work completed during the maintenance window are available for verification by a security administrator or auditor.

Trusted Security Context

- Identifies “trusted” DDF, RRS Attach, or DSN application servers
- Allows selected DB2 authids on connections without passwords
 - reduces complexity of password management
 - reduces need for an all-inclusive “system authid” in app servers
 - more visibility/auditability of which user is current running
 - enables mixed security capabilities from a single app server



Trusted security context: Today, you have the option to set a system parameter which indicates to DB2 that all connections are to be trusted. It is unlikely that all connection types, such as DRDA, RRS, TSO, and batch, from all sources will fit into this category. It is likely that only a subset of connection requests for any type and source may be trusted or that you want to restrict trusted connections to a specific server. More granular flexibility will allow for the definition of *trusted connection objects*.

Once defined, connections from specific users via defined attachments and source servers will allow trusted connections to DB2. The users defined in this context can also be defined to obtain a *database role*.

Auditing: DB2 Trace Filtering

- New filtering capabilities for `–START TRACE` that `INCLUDE` or `EXCLUDE` based on these keywords:
 - USERID -- client userid
 - WRKSTN -- client workstation name
 - APPNAME -- client application name
 - PKGLOC -- package LOCATION name
 - PKGCOL -- package COLLECTION name
 - PKGPROG -- PACKAGE name
 - CONNID -- connection ID
 - CORRID -- correlation ID
 - ROLE – end user's database ROLE

Improved trace filtering makes the jobs of auditing and of performance management easier. Many more options can be used to minimize the amount of data collected, so the overhead is reduced and the extraneous data does not need to be processed. Being able to include or exclude, based upon more qualifications reduces performance concerns and the amount of data to store and process.

Example: ROLES and Trusted Context to Secure DBA Activities


- Many customers are concerned about DBA access to sensitive customer data. DB2 V9 can help by enabling an auditable DBA process:
 1. Grant DBA privileges to a ROLE
 2. Start audit trace for that ROLE
 3. When a DBA needs to perform a system change:
 - Use Trusted Context to assign DBA ROLE to person
 - DBA is given request and performs activity
 - Revoke Trusted Context
 4. Have another person review the audit trace

The Trusted Context and ROLE support can be used to implement DBA privileges that can easily be disconnected and reconnected to individual employees. This provides function similar to shared SYSADM or DBADM userids, but avoids the audit compliance problems associated with shared userids. The ROLES have the ability to “own” DB2 objects, so that revoking a person’s ROLE does not cause the objects to be cascade deleted.

With these capabilities, customers are able to create DBA procedures that can be audited and protected so that one individual cannot violate the established rules without being detected during the audit review.

Future Directions – Extending Encryption to IBM TotalStorage

- Statement of Direction: To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.
- This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage® portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF.



Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only


Statement of Direction: To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.

This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage® portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF

DB2 V9.1 for z/OS

Performance Improvements

- CPU reductions in LOAD and REORG
- LOB performance, function, scalability
- SQL and optimization improvements
- Use zIIP in remote native SQL Procedure Language
- Improved varying length performance
- Sequential disk access



The key performance improvements in V9 are reduced cpu time in the LOAD and REORG utilities, improved LOB performance and scalability, improved optimization for SQL, the zIIP processing for remote native SQL procedures, reduced cpu time for data with varying length and better sequential access.

I'll discuss the optimization improvements on the query slide. V8 SQL procedures were not eligible to run on the zIIP, but changing to use the native SQL Procedure Language on V9 will make the work eligible for zIIP processing. Varying length data can improve substantially if there are large numbers of varying length columns. Several improvements in disk access can reduce the time for sequential disk access.

Query Enhancements

- SQL enhancements: INTERSECT, EXCEPT, RANK, caseless comparisons, cultural sort, FETCH FIRST in fullselect ...
- Index improvements: index on expression
- Improved Optimization statistics: Histogram
- Optimization techniques
 - Cross query block optimization
 - Generalize sparse index & in-memory data cache method
 - Dynamic Index ANDing for Star Schema
- Analysis: instrumentation & Optimization Support

Improving data warehousing and reporting: Today's complex applications include both transactions and reporting, so performing both well is imperative. The key improvements for reporting are optimization enhancements to improve query and reporting performance and ease of use. Improved data is provided for the optimizer, with improved algorithms and a rewritten approach to handling performance exceptions.

More queries can be expressed in SQL with new SQL enhancements. The set operators INTERSECT and EXCEPT clauses make SQL easier to write. OLAP extensions for RANK, DENSE_RANK and ROW_NUMBER add new capabilities. Other SQL statements improve consistency with the DBMS industry. V9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on a prior slide. New SQL data manipulation statements are MERGE and TRUNCATE. New data types with DECIMAL FLOAT, BIGINT, BINARY and VARBINARY. Improvements in LOBs provide new function, more consistent handling and improved performance. Security is improved with ROLES and network trusted context. Data definition consistency and usability are improved. V9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.


DB2 V9 z/OS Innovation: Resiliency

- Business Insight**
- Cost Savings through Optimization**
- Business Resiliency**


Availability: Database on Demand



Scalability



Usability



52

BUSINESS RESILIENCY INNOVATIONS

o Improvements to DB2 enhance performance, usability, and availability in data sharing environments, including fast replacement of one table with another, ability to alter table space and index logging, and improved ability to rebuild an index online as well as restart improvements.


o Scalability innovations like table space partitioning by growth, log performance, and ability to create and alter STOGROUP SMS constructs.

o Further improvements in V9 help manage business-critical information with fewer planned or unplanned outages with features like online reorganization with no BUILD2 phase, rename column and index, alter index to change page size, and ability to modify early code without requiring an IPL.

o Usability improvements keep the user at the center of the design, working with a broad range of products to allow a range of platforms, programming languages and tools.

Online Schema Evolution → Database Definition On Demand

- Online reorganization with no BUILD2 phase
- Fast replacement of one table with another
- Table space that can add partitions, for growth
- Improve ability to rebuild an index online
- Rename column and index
- Modify early code without requiring an IPL
- Alter table space and index logging
- Create & alter STOGROUP SMS constructs




One of the key initiatives of V8 was online schema evolution. Online schema evolution is expanding and changing to be data definition on demand.

Online table space reorganization for a few partitions is improved a lot, removing the BUILD2 phase for all types of secondary indexes. One of the important changes is to be able to replace one table quickly with another. Another is to be able to rename a column or an index. A new type of table space combines the attributes of segmented and partitioned, without a partitioning key. Rebuild index can be run with much less disruption. Table space and index logging can be altered. A new ability to change the DB2 early code does not require an IPL. SMS constructs MGMTCLAS, DATACLASS and STORCLAS can be defined on a STOGROUP and can be altered.

Scalability

- Insert performance APPEND INDEX LOG
INDEX on expression, 8K, 16K, 32K
Randomized index key
Log data sharing performance, archiving
Not logged table space
- Partitioned table with segmented space
- Memory improvements 64 bit address space



Performance for inserting is expected to increase substantially, through a wide range of improvements. Logging performance is improved substantially, with data sharing improvements and archiving. The newer disk and channel changes (DS8000, 4 Gb per second channels, MIDAW), improve the data rates substantially. Indexes are improved, with larger page sizes to reduce the number of page splits and also a better page split. Where performance should be optimized for inserts, rather than for later retrieval, the append option can be used. If the data need to be randomized to avoid insert hot spots, the new randomized index key is useful.

Memory improvements continue the work from V8, with shared memory above the bar for DDF and DBM1 and moving more data structures above the bar.

The segmented space structure is more efficient, so adding that space structure for the large partitioned table spaces helps DB2 scale.

Partition by Growth & Universal Table Space

- New partitioning scheme:
 - Single table tablespace, where each partition contains a segmented pageset (allows segmented to increase from 64GB to 16TB or 128 TB with 32K pages)
- Partition By Growth
 - Eliminates need to define partitioning key and assign key ranges
 - A new partition is created when a given partition reaches DSSIZE (defaults to 64G)
 - Retains benefits of Utilities and SQL parallelism optimizations for partitioned tables


Partitioned by growth and universal table space:


Partitioned tables have required key ranges to determine the target partition for row placement. When a table is partitioned, you gain the benefits of scaling objects to hold more data. You also benefit from more granular locking and parallel operations by spreading the data over more data sets.

The option to partition by growth will allow segmented tables to be partitioned as they grow, without needing key ranges. These segmented tables will gain increased table space limits and the SQL and utility parallelism, afforded for partitioned tables.

IBM Software Group | DB2 Information Management Software IBM

Usability

- Application Programming
 - Visual Studio.net
 - php
 - perl
 - Java
 - eclipse
 - Zend Core™ for IBM
 - RAILS
 - Rational™ the software development company
 - WebSphere
 - PYTHON
- Backup and restore objects, use tapes
- Optimization Support Center 
- INDEX page split
- Buffer pool management by WLM



56

Application programming is using a wider range of tools, environments and languages. The Eclipse framework is growing strongly. We need to connect the new languages and environments to the scale and value of the existing infrastructure. The Rational and WebSphere product lines provide the connection.

Database administrators need to be able to use the FlashCopy backups to restore a single object and to be able to use tape for backup more easily. If queries are not performing well, performance analysts need a full set of tools to monitor, analyze and tune. The Optimization Support Center delivers.

Most index splits today are roughly half of the page into each split. When processing is sequential, too much space is wasted, and too many page splits. A dynamic option can save space and time.

WLM management is extended with the ability to adjust buffer pool sizes.

DB2 V9 for z/OS



❖ Business Insight



❖ Cost Savings through Optimization



❖ Business Resiliency



These are primary areas for our next version, carrying on some of the key deliveries from the prior work. Migration to V9 will be only from DB2 for z/OS Version 8 and will prereq z/OS 1.7. Some items will require z/OS 1.8.


Data definition on demand extends the theme of online schema revolution from V8. Additional Unicode enhancements continue the work from V7 and V8. XML work across the DB2 family is a much larger step than in V7 or V8. SQL Procedures become more consistent across the family. While V7 and V8 removed many differences from DB2 for Linux, Unix & Windows, V9 takes the next big step to improved consistency. Utility enhancements help with new function, more LOB and XML support, better performance and improved availability, removing the BUILD2 step from online reorg.

Release Vision: The vision for DB2 V9 for z/OS is to enhance DB2's ability to handle new and enterprise applications. V9 improves the ability to handle new applications with XML, large objects, and many SQL and security improvements. V9 builds upon and extends DB2 traditional strengths and the ground-breaking Version 8 in many areas: online schema evolution, Unicode, XML, DB2 family SQL, utilities, security and 64-bit virtual storage.

Disclaimer and Trademarks

Information contained in this material has not been submitted to any formal IBM review and is distributed on "as is" basis without any warranty either expressed or implied. Measurements data have been obtained in laboratory environment. Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. The use of this information is a customer responsibility.

IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, zSeries, 



The FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC, WINDOWS 95

For additional information see ibm.com/legal/copytrade.phtml

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in the operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information may contain examples of data and reports used in daily business operations. To

illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.