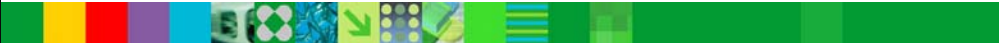


IBM Software Group

DB2 for z/OS Version 8 and Beyond

Curt Cotner, IBM Fellow
ibm.com/software/db2zos

DB2 Information Management Software



@business on demand software

This is an overview of DB2® for z/OS Version 8 (V8) and beyond. DB2 V8 is the twelfth release and delivers more function than any release of DB2 for MVS. DB2 V8 became generally available March 26, 2004, bringing extensive integration and synergy with System z9 and zSeries hardware, with middleware and with applications. Data support, application development and query enhancements are added for e-business, building upon the traditional enterprise of choice characteristics of availability, exceptional scalability, and performance. DB2 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, and continuous availability. Management for very large databases is made much easier, 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. These enhancements include SQL improvements, schema evolution, longer names for tables and columns, longer SQL statements, enhanced Java and Unicode support, enhanced utilities, more log data sets, and a lot more.

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.

Migrating and 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.

Enterprise Applications & : DB2 for z/OS

 4.6 certified	–64 bit virtual storage	–Longer Table Names & Column Names
	–Unicode	–Additional statistics
 8.45 certified	–Schema evolution	–Convert Column Type
	–System-level backup and recovery	–Altering CLUSTER Option
	–Multi-row fetch & insert	–Adding Columns to Index
 7.8 certified	–Multiple DISTINCT Clauses	–Index-only Access Path for VARCHAR
	–Lock Contention on volatile Tables	–Adding New Partitions
	–Fast Retrieval of Most Recent Value	–Separate Clustering from Partitioning
		–...

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.

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**

zSeries, z/OS & ESS platform
 Middleware stack



Rational.

WebSphere.

DB2

Lotus.

Tivoli

Reengineered for e-business on demand

9

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

- New function

- 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
 - IBM System z9 Integrated Information Processor (IBM zIIP)
- New and updated books: Library refresh April 2005, Messages, Codes August
 - Redbooks: Design Guidelines for High Performance and Availability, Business Value, Performance Topics, WebSphere, MLS, Disaster Recovery, others updated ...
 - Customer information on the web

10

✓DB2 for z/OS and WebSphere: The Perfect Couple, SG24-6319

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246319.pdf>

✓Achieving the Highest Parallel Sysplex Availability DB2, REDP-3960,

<http://www.redbooks.ibm.com/redpapers/pdfs/redp3960.pdf>

✓Planning for Multilevel Security & Common Criteria (GA22-7509)

<http://publibz.boulder.ibm.com/epubs/pdf/e0z2e111.pdf> & [e0z2e120.pdf](http://publibz.boulder.ibm.com/epubs/pdf/e0z2e120.pdf) (z/OS 1.7)

✓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>

Stored Procedures <http://www.redbooks.ibm.com/redbooks/pdfs/sg247083.pdf>

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

✓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 April 2005, one August 2005

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

✓DB2 for z/OS Performance Topics SG24-6465

- Cross loader with LOBs PQ90263

- Built in functions ASCII, TIMESTAMPDIF PQ95795

- DSN1COPY with 1000 OBIDs PK05758

- QMF with multirow fetch PQ99482

- Online Check Index PQ92749 & PQ96956

- z/OS 1.7 up to 7257 extents PK07590 PK10594

- LOAD, UNLOAD with LOBs PK10278 open

IBM System z9, z/OS & DB2 for z/OS

✓ **System z9 Integrated Information Processor (zIIP)**

- ✓ **Enhanced Cryptography**
- ✓ **Enhanced channels (MIDAW)**
- ✓ **Faster Processors**
- ✓ **Up to 54 Processors**
- ✓ **More memory, better value; 64 bit virtual storage**
- ✓ ...



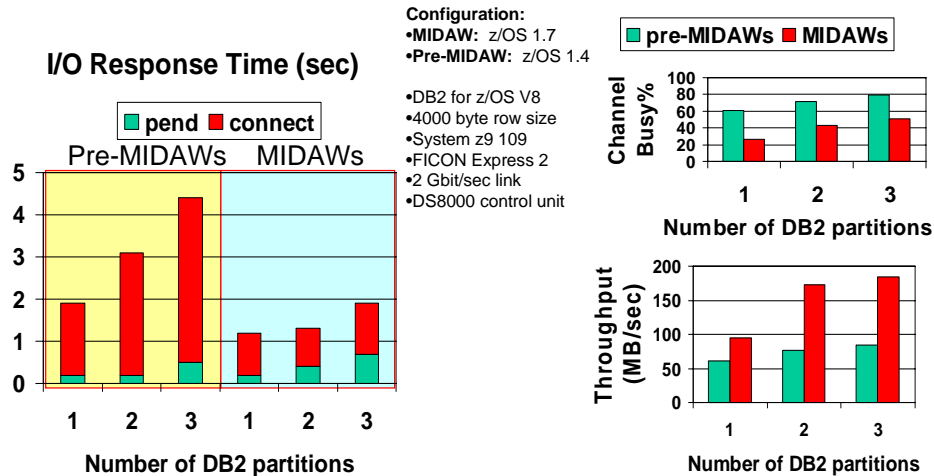
- ✓ **Backup and restore**
- ✓ **Multilevel Security**
- ✓ **Unicode conversion**
- ✓ **Compression**
- ✓ **zSeries Application Assist Processor**
- ✓ **z/Architecture new instructions**
- ✓ **WLM enhanced**
- ...

The latest System z9 processor improvement for DB2 is the zIIP. IBM announced cryptography improvements in late 2005. Channel enhancements (MIDAW) and improved DS8000 performance were included with the System z9 announcement. 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, 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).

ibm.com/software/db2zos/ Click on Support, then on Frequently Asked Questions. Qualify the search with z990 to get the answer.

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



12

These are some preliminary measurements with 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.


Technology Evolution with Mainframe Specialty Engines

★ Building on a strong track record of technology innovation with specialty engines, IBM intends to introduce the System z9 Integrated Information Processor



IBM System z9 Integrated Information Processor (IBM zIIP) planned for 2006

Centralized data sharing across mainframes



Internal Coupling Facility (ICF) 1997



Integrated Facility for Linux (IFL) 2001

Support for new workloads and open standards




System z9 Application Assist Processor (zAAP) 2004

Incorporation of JAVA into existing mainframe solutions

Designed to help improve resource optimization for eligible data workloads within the enterprise

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 z9-109 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 109
 - z/OS 1.6 or later
 - DB2 for z/OS V8



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



New Specialty Engine

- 2 - Data warehousing applications*
 - Requests that utilize star schema parallel queries
- 3 - DB2 for z/OS V8 utilities*
 - Internal 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 offloaded.

IBM Software Group | DB2 Information Management Software IBM

OMEGAMON for DB2 and DB2 Performance Monitor/Expert – Convergence: Available Now

Best of breed DB2 monitoring, analysis, and tuning solution

```

graph LR
    A[DB2 Performance Expert V2] --- B[ ]
    C[OMEGAMON XE for DB2 V300] --- B
    B --> D[OMEGAMON XE for DB2 Performance Expert on z/OS V3.1]
    E[DB2 Performance Monitor V8] --- F[ ]
    C --- F
    F --> G[OMEGAMON XE for DB2 Performance Monitor on z/OS V3.1]
    H[DB2 Buffer Pool Analyzer V2] --> I[DB2 Buffer Pool Analyzer V3.1]
  
```

16

Taking the best from both products means to combine unique functions with the existing ones of the other offering, to replace identical functions with one of both solutions and to add / implement missing functions into existing components. The target, of course, is to keep as much functions as possible either coming from OMEGAMON or coming from DB2 PM/PE. We also considered the actuality (newest DB2 version support as well the effort) to be spent to provide the functions. For simplification you might say, the new product is the DB2 PM/PE base services and functions (reporting, PWH, and DB2 API access) with the infrastructure, the integration capability and end user interface of OMEGAMON.

Although we prefer to sell OMEGAMON XE for DB2 PE, we still see a need to offer performance monitoring with and without specific buffer pool analysis, buffer pool expert advices, and buffer pool simulation because of market needs. Therefore we will continue to offer the OMEGAMON XE for DB2 PE and the OMEGAMON XE for DB2 PM as separate offerings. However, our intention is to add specific expert functions only to the OMEGAMON XE for DB2 PE and to support only base monitoring support to OMEGAMON XE for DB2 PM too, i.e. new DB2 functions and versions will be added to the PM and PE, but new specific expert functions will be added to PE only.

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

Vnext DB2 Technology Themes

- Enable high-volume transaction processing for next wave of Web applications**
 - Extend the lead in transaction processing availability, scalability and performance**
 - Reduce cost of ownership and zSeries-specific skill needs**
 - Improve data warehousing and OLTP reporting**

Themes for VNext: There are four primary themes for the VNext enhancements. The first is enabling high volume transaction processing for the next wave of web applications. This work extends DB2's ability to handle the next generation of applications by adding XML support and extending SQL and security. The SQL improvements also improve the ability to port applications to DB2 for z/OS and DB2 family compatibility.

The second theme is extending the lead in transaction processing availability, scalability and performance. The key changes in availability extend DB2's ability to make changes to the data definitions without disrupting the online processing. New structures are provided to improve scalability and performance. Bottlenecks for very high performance are removed.

The third theme is reducing the cost of ownership and the need for skills that are specific to the zSeries platform. This theme includes productivity improvements for administrators through consistency with the DB2 family and the relational database management industry. It also includes cost reductions through processing reductions.

The fourth theme is 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. More queries can be expressed in SQL with new SQL enhancements.

IBM Software Group | DB2 Information Management Software IBM

DB2 SQL

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



- z** { Range partitioning
- common** { 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
- LUW** { 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

18

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

19

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 V8 Vnext
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, TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE

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, 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

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

20


This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is the next release of DB2 for z/OS, which we call Vnext. Vnext 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 for LUW, code named Viper. We expect to be 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.

IBM Software Group | DB2 Information Management Software IBM

DB2 for z/OS Vnext SQL, DB2 family & porting



<ul style="list-style-type: none"> ▪ XML ▪ MERGE ▪ SELECT FROM UPDATE, DELETE, MERGE ▪ TRUNCATE ▪ INSTEAD OF TRIGGER ▪ BIGINT, VARBINARY, DECIMAL FLOAT ▪ Native SQL Procedure Language ▪ Optimistic locking 	<ul style="list-style-type: none"> ▪ 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
--	--

21

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

Vnext continues the progress in SQL, with many new functions statements and clauses. The biggest changes are in XML on following slides. There are new SQL data manipulation statements in MERGE and TRUNCATE. There are new data types with DECIMAL FLOAT, BIGINT and VARBINARY types. Improvements in LOBs provides more consistent handling and improved performance. Security is improved with ROLES and trusted context. Data definition consistency and usability are improved. Vnext is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

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.

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.

This is very similar to the changes in DB2 for LUW V8.2.

Optimistic Locking Support

- Built-in timestamp for each row or page
 - Automatically updated by DB2
 - Allows simple timestamp predicate to validate that row has not changed since last access
- Eliminates need for complex predicates on WebSphere CMP updates, improves performance

This change will help with migration of applications which use application consistency techniques.

XML Processing Paradigms

XML has become the “data interchange” format between B2B/B2C, inter- and intra-enterprise environments.

XML View Of Relational Data

- SQL data viewed and updated as XML
 - Done via document shredding and composition
- DTD and Schema Validation

XML Documents As Monolithic Entities

- Atomic Storage And Retrieval
- Search Capabilities

XML As A Rich Data Type

- Full Storage and Indexing
- Powerful Query Capabilities



XML is an important innovation for several reasons:

It enables business-to-business communication, making it easy to send semi-structured data across the Web so that nothing gets lost in translation, masking differences in end-point infrastructures. Provides independence from the presentation layer -- Information can be rendered appropriately for a variety of devices, because XML separates the document content from the document presentation. In this regard, it differs significantly from HTML.

It enables the use of smart agents, because XML includes a description of the data and how it relates to other pieces of data the agent may already know, enabling it to take smarter actions. It enables the use of smart searches, because XML provides a context for search arguments, for example, the name Chip rather than the food chip. From the standpoint of data management, there are three distinct paradigms for XML usage and support:

XML View Of Relational Data

This is for users who merely want to XML-ize their relational data: They want to receive XML documents and shred them into distinct relational tables. And they want to send XML documents generated out of their relational tables. There's really no trace of XML inside the data management storage. XML is used as a transient on-the-wire format.

XML Documents As Monolithic Entities


This is for users who want to: Store XML documents as-is. Retrieve them as a whole. They may also want some form of rudimentary querying capabilities on XML.

XML As A Rich Datatype

This is for users who want complete support from XML in a data management system. This implies efficient storage and indexing

XML 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 most of this picture, while DB2 for z/OS Vnext will not include XQuery.

XML Support

- Support XML data type
- Store the XML document natively
- DDL --
 - CREATE/ALTER Table with XML type column
 - Implicitly create XML Auxiliary objects (tablespace/table/index) - one per XML column
 - Index support
 - Created by users
 - uses XPath to determine which nodes in the XML document to index.

```
CREATE INDEX dependentName ON deptTable(deptDocs)
GENERATE KEY USING XMLPATTERN
'/department/empl/dependent/name' ATOMIC AS SQL
VARCHAR(20);
```
- INSERT/UPDATE/DELETE
 - INSERT with VALUES and SUBSELECT
 - No Subdocument update

Now let's discuss the wide range of new function being provided inside the DB2 for z/OS engine. This work is being done in parallel with similar changes in DB2 for Linux, Unix and Windows.

DB2 for z/OS extensibility has been implemented via *extenders*. Extenders for text, image, audio, video, and XML are delivered with DB2 and are optional installation components. They provide the tooling, user defined data types (UDT), user defined functions (UDF), and stored procedures to managed non-relational data. The XML extender provides for the storage and indexing of an XML document as a character large object (CLOB), or for the shredding of the XML document into relational columns for storage and query.

DB2 Version 8 expanded on XML support by implementing several XML publishing operations as built-in DB2 functions. This allows you to perform XML document composition from relational data with improved performance without the XML Extender.

There will be expanded support of XML in DB2 by integrating more features into the engine. This includes an XML data type, native storage of XML documents, integration of the XPath language, and catalog extensions to support definitions of XML schemas. Utilities will support creation and maintenance of XML data.

XML Support -- Query

- Enhanced V8 XML Constructors (XML Publishing Functions)
- SQL/XML Functions and Predicates
 - XMLParse - Convert a XML text to XML value
 - XMLSerialize - Converts XML to non-XML type
 - XMLQuery - executes an XPath expression against an XML value.

```
SELECT XMLQUERY ( '//item[USPrice = $price]' PASSING
PO.POrder,
T.price AS "price") FROM PurchaseOrders PO, T;
```
 - XMLCast - Cast XML to other types or other types to XML
 - XMlexists - a predicate, which returns TRUE if the XPath expression evaluates to a non-empty sequence

```
SELECT PO.pid FROM PurchaseOrders PO, T
WHERE XMLEXISTS( '//item[USPrice = $price]'
PASSING
PO.POrder, T.price AS "price")
```

XML Support (continued...)

- XPATH
 - Supported Features from XPath 2.0:
- Utility Support
 - LOAD/UNLOAD, CHECK DATA/INDEX, COPY, REBUILD, RECOVER, REORG, etc.
- XML Schema Support
 - XSR – XML Schema Repository
 - Catalog Tables to store XML schemas
 - Stored procedures to register XML schemas
- XMLVALIDATE() SQL/XML function
 - Test XML values for validity against XML schema
 - Obtain default values and schema normalized values from XML schema

TRUNCATE Statement

- Allows fast delete of all rows in a given table (simple, segmented, or partitioned)
- Very useful for nightly refresh of summary tables, warehouses, etc.

```
TRUNCATE TABLE TABLE-NAME  
  
< DROP STORAGE | REUSE STORAGE >  
  
< RESTRICT WHEN DELETE TRIGGERS |  
  IGNORE DELETE TRIGGERS >  
  
< IMMEDIATE >
```

This statement provides a fast way to delete rows with SQL, with better application portability. Truncate Table provides for the rapid removal of rows from a table. You can use this function to delete the contents of a table before applying data from a DSN1COPY or INSERT.

Decimal Floating Point

- New datatype NUMBER
 - Well suited to typical customer financial calculations
 - Similar to “calculator” mathematics
 - Eliminates rounding errors by using base 10 math
 - Has up to 34 digits of precision
 - Floating point convenience with fixed point precision!!!
 - Hardware support will be provided in the next zSeries processor generation (new IEEE standard)
 - Software emulation provided for other models



New data type support

Support will be made for:

- Decimal floating point numbers, similar to calculator mathematics and supporting the IEEE standard
- BIGINT support of double word (8 byte) integer values
- VARBINARY, providing better comparison support for binary strings

MERGE

- Array MERGE operation
- Targets OLTP applications like SAP

```
MERGE INTO account AS T
USING VALUES (:hv_id, :hv_amt) FOR 5 ROWS AS S(id,amt)
ON T.id = S.id
WHEN MATCHED THEN
  UPDATE SET balance = T.balance + S.amt
WHEN NOT MATCHED THEN
  INSERT (id, balance) VALUES (S.id, S.amt)
NOT ATOMIC CONTINUE ON SQLEXCEPTION
```

Merge

Online transaction processing (OLTP) workloads that need to place data for several rows into a table will benefit from merge SQL. This capability will allow for updating of existing rows and creation of new rows through a single SQL statement.

Multi-row insert-type capability is extended, via the MERGE statement, to take an array of column values from a program and perform insert and update operations against a table. DB2 will use the matching criteria in the merge SQL to update existing rows and perform inserts for rows that do not exist, through a single SQL statement.

SQL Improvements – Family Compatibility

- INSTEAD OF triggers
- SELECT FROM UPDATE
- SELECT FROM DELETE
- SELECT FROM MERGE
- BIGINT and VARBINARY data types
- ORDER BY and FETCH FIRST in subselect

Select from DELETE, UPDATE, and MERGE: The object-relational capabilities of DB2 allow for the incorporation of business logic into the database. This extends the power of SQL. Sometimes the application needs to know the results of this logic, when applied to the SQL issued. A subsequent SELECT for the data adds complexity and execution time to the application.

The *insert within select* feature of DB2 for z/OS Version 8 has been expanded to include the retrieval of columns from rows that are modified via DELETE, UPDATE, and MERGE SQL. One SQL call to DB2 modifies the table contents and returns the resultant changes to the application program.

When used with DELETE, the application now has the option to code a destructive read from a table. This is particularly useful when a table is used as a data queue, as with many vendor packages.

Subquery improvements: Correlated and non-correlated subqueries will benefit from improved optimization. They will provide added flexibility with the support of ORDER BY and FETCH FIRST clauses.

DDF Improvements

- 64-bit exploitation by DDF
 - Special “shared private” with xxxDBM1 to eliminate many of the data moves on SQL operations
- Support for IPv6 and SSL
- VTAM definition is now optional
- Elimination of PRIVATE protocol requester
 - Includes tools for identifying which packages need to be bound at remote servers

Distributed connections to DB2 for z/OS will benefit from z/OS V1R7 changes that DB2 will exploit. DB2's distributed communication processes (the *distributed address space*) will access data directly from the database manager address space, instead of moving the data. The distributed address space will also exploit 64-bit addressing, as the database manager and lock manager address spaces do today with Version 8.

This internal change will benefit new and existing workloads, where distributed communications are configured with another logical partition (LPAR) or to an application running on the zSeries platform.

DB2 Vnext Themes

- ❑ **Enable high-volume transaction processing for next wave of Web applications**
- ❑ **Extend the lead in transaction processing availability, scalability and performance**
- ❑ **Reduce cost of ownership and zSeries-specific skill needs**
- ❑ **Improve reporting**

Themes for VNext: There are four primary themes for the VNext enhancements. The first is enabling high volume transaction processing for the next wave of web applications. This work extends DB2's ability to handle the next generation of applications by adding XML support and extending SQL and security. The SQL improvements also improve the ability to port applications to DB2 for z/OS and DB2 family compatibility.

The second theme is extending the lead in transaction processing availability, scalability and performance. The key changes in availability extend DB2's ability to make changes to the data definitions without disrupting the online processing. New structures are provided to improve scalability and performance. Bottlenecks for very high performance are removed.

The third theme is reducing the cost of ownership and the need for skills that are specific to the zSeries platform. This theme includes productivity improvements for administrators through consistency with the DB2 family and the relational database management industry. It also includes cost reductions through processing reductions.

The fourth theme is 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. More queries can be expressed in SQL with new SQL enhancements.

Schema Evolution – Database Definition On Demand

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

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

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. Index page size can be altered. 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. Online table space reorganization for a few partitions is improved a lot, removing the BUILD2 phase for all types of secondary indexes. Table space and index logging can be altered. A new ability to change the DB2 early code does not require an IPL.

CLONE Tables

- Allows fast replacing production data without renames and rebinds
 - A capability to support online load replace
- CREATE TABLE to create a Clone Table
 - All indexes are also cloned
 - Table and Index data are not copied
 - Base and Clone tables share the same table space and index names
 - Underlying data sets are differentiated by a data set instance number

Fast replacement of a table with another

This function will allow you to generate a copy of a current table with the same attributes and same data, in the same table space. It appears to an application or user as a very fast replacement of data within a table. Web-based applications, striving for maximum availability will benefit from the option of implementing a pair of tables that are clones of each others' structure. Copies for application testing and auditing can be easily created. These *clone* tables will have the unique ability to change names quickly. Applications can quickly and almost transparently switch between dual mirror tables.

CLONE Tables...

- A clone table can only be created
 - On a single table in a table space (partitioned or non-partitioned)
 - No RI or Trigger on the base table
 - No new index can be created on a base table that has a clone
- Use insert or load to populate clone tables
- Utilities (except RUNSTATS) can operate on clone tables with a new CLONE keyword

The tables can be partitioned or non-partitioned, and are created with the CREATE TABLE syntax. The primary table's structure, including indexes, large objects (LOBs) and before triggers, will be copied. Information can then be inserted or loaded into the copy table, and the copy can have its own image copies. When the data in the copy table needs to become active to an application, an ALTER statement will switch the name, providing fast replacement of the original data

Partition by Growth

- 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)
 - 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

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.

CREATE TABLE ... APPEND(YES)

- New APPEND option:
 - Maximizes performance for “INSERT at end”
 - Avoids overhead of attempting to preserve clustering sequence

Table Append option

The Table Append option offers increased performance for inserting data into the end of a table. It reduces the instructions used to target the locations for new rows. Index maintenance techniques are also made more efficient when DB2 detects that entries are added to the beginning or ending of the index.

Relief for Sequential Key INSERT

- New page sizes: 8K and 16K for INDEX pages
 - Fewer page splits for long keys
 - More key values per page
- INSERT at the end of the key range used to result in 50% free space in each index page
 - Enhanced support dynamically adapts page split boundary to minimize wasted space in index pages

Sequential insert performance is improved by avoiding page splits with larger index page sizes and the ability to split a page more effectively. Other changes improve logging rates.

LOB Performance/Scalability

- Elimination of LOB locks – LRSN and page latching is used instead for consistency checks
- New network flows for delivering LOBs
 - JDBC, SQLJ, and CLI will let server determine whether to flow LOB values or LOCATORs based on size thresholds
 - Significant reduction in network traffic
 - Greatly reduces frequency of FREE LOCATOR statements

Large object improvements: Large objects (LOBs) were introduced in DB2 Version 6. Usage has increased substantially in the past few years, and major enhancements have been made in DB2 Version 8.

APARs on Version 8 deliver the ability to use utilities for loading and unloading large LOB data. File reference variables are used to let the large objects be accessed from data sets instead of from storage. The abilities to reorganize and to recover space are provided.

Future changes will help with improved function and usability, DB2 family compatibility, cost of ownership, performance, and scalability.

Other Performance / Availability Items

- Insert performance APPEND INDEX LOG
- INDEX on expression
- LOB performance, function, scalability
- Log latch contention relief
- CPU reductions in LOAD and REORG
- Online REBUILD INDEX
- FETCH FIRST n ROWS improvements
 - Can now be specified in a subquery or fullselect
 - ORDER BY now exploits FETCH FIRST n ROWS, so that work files are not created (less I/O)

Many other improvements help with performance, with scalability and with availability. Index on an expression can be combined with caseless comparisons to improve text search. Improved insert rates can result from improved latching of the log data. Significant reductions in cpu usage are provided with the new utilities.

Improved cpu and elapsed times can be achieved with the FETCH FIRST clause specified on a subquery.

DB2 Vnext Themes

- ❑ **Enable high-volume transaction processing for next wave of Web applications**
- ❑ **Extend the lead in transaction processing availability, scalability and performance**
- ❑ **Reduce cost of ownership and zSeries-specific skill needs**
- ❑ **Improve reporting**

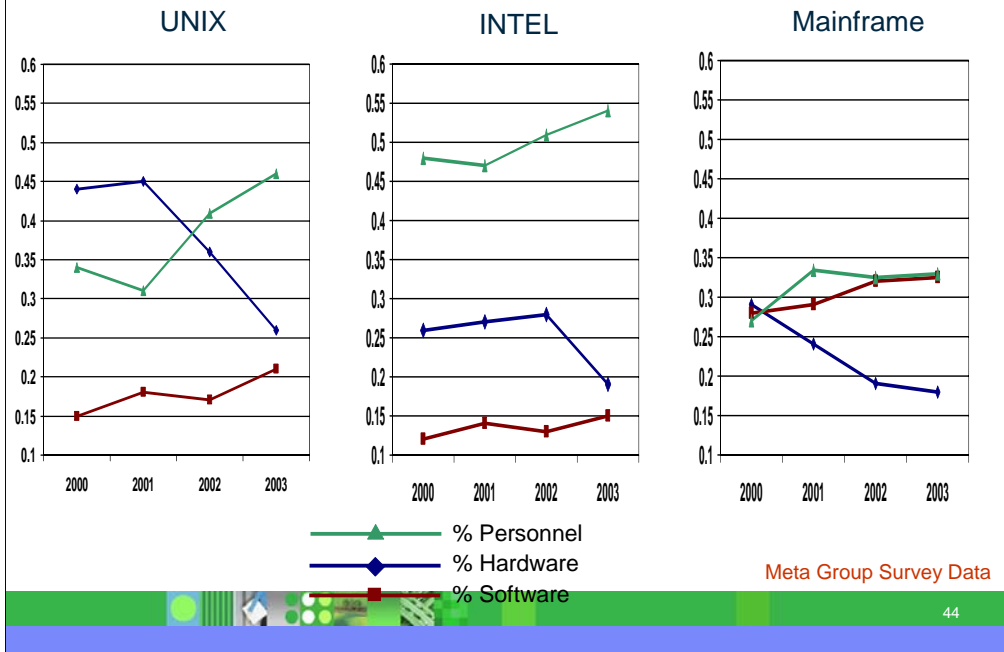
Themes for VNext: There are four primary themes for the VNext enhancements. The first is enabling high volume transaction processing for the next wave of web applications. This work extends DB2's ability to handle the next generation of applications by adding XML support and extending SQL and security. The SQL improvements also improve the ability to port applications to DB2 for z/OS and DB2 family compatibility.

The second theme is extending the lead in transaction processing availability, scalability and performance. The key changes in availability extend DB2's ability to make changes to the data definitions without disrupting the online processing. New structures are provided to improve scalability and performance. Bottlenecks for very high performance are removed.

The third theme is reducing the cost of ownership and the need for skills that are specific to the zSeries platform. This theme includes productivity improvements for administrators through consistency with the DB2 family and the relational database management industry. It also includes cost reductions through processing reductions.

The fourth theme is 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. More queries can be expressed in SQL with new SQL enhancements.

Cost Of Ownership Trends



These are trends in cost of ownership for various platforms, as noted in a Meta Group survey of customers. The percentage of costs differs significantly. Hardware costs continue to decline on all platforms, while people costs are climbing the most and software increases its percentage.

So one of the most important categories to reduce is the people cost – even as the others are contained.

Compliance/Auditing Pressure

- Regulatory compliance initiatives are impacting IT organizations in most countries/industries, and changing fast
 - Sarbanes-Oxley
 - Basel II
 - FDA: Food and Drug Administration 21 CFR Part 11
 - COPPA: Children's Online Privacy Protection Act of 2000
 - DPA: Data Protection Act (UK)
 - HIPAA: Health Insurance Portability and Accountability Act of 1996
 - PIPEDA: Personal Information Protection and Electronic Documents Act (Canada)
 - SEC Rule 17a-4: Records to be preserved by certain exchange members, brokers, dealers
 - USA Patriot Act: Uniting and Strengthening America by Providing Tools Required to Intercept and Obstruct Terrorism of 2001
- Focus is on both external threats (hackers) and internal employees

Today's computing environment is subject to increasing regulatory pressures and potentially malicious attacks. The security of the information, to which you have been entrusted, has never been more critical. DB2 for z/OS already resides on one of the most secure platforms in the industry. It provides secure access to your data through internal or external authorities and privileges, encryption, and multilevel security with row-level granularity enforcement.

Security can be difficult to manage. The plan for DB2 is to provide you with additional capability while assisting you with security management.

Security in DB2 for z/OS Vnext

Some key implementations

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



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 LUW. 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.


Protecting data on disk

- We will allow encryption for the key disk resources used by DB2:
 - Tables
 - LOBs
 - Indexes
 - Image copies
 - Logs
 - Archive logs

Encryption is a significant technology for the System z9 and zSeries platform. There are new facilities that have just been provided and more work is coming.

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

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.

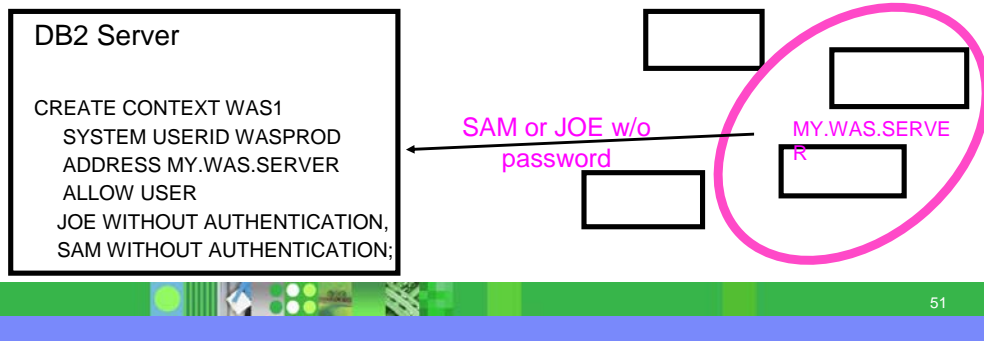
Auditing trails of the work completed during the maintenance window are available for verification by a security administrator or auditor.

Database ROLES Examples

- Dynamic SQL access to DB2 tables using JDBC or CLI, but only when running on a specific server.
- DBA can be temporarily assigned a DBA ROLE for weekend production table admin work – no table access at other times.
- DBA uses a ROLE for CREATE statements, so that the ROLE owns the objects he or she creates.
- Project librarian assigned a BIND ROLE only when running on the production code library server – can't BIND from any other server.

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*.

Trusted Security Context / ROLE WebSphere example

- WebSphere connection pool can be created with one DB2 AUTHID
- WebSphere can reuse pooled connections to DB2 with different AUTHIDs
- DB2 AUTHIDs can be given privileges that are only available when executing in WebSphere:
 - ✓ e.g. dynamic SQL access for JDBC only when using WebSphere



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 job of auditing and of performance management easier.

Volume-based COPY/RECOVER

- FlashCopy technology used to capture entire content of disk volumes
- RECOVER modified to enable object-level recovery from volume FlashCopy
- Eliminates labor associated with setting up COPY jobs for each database / table space

V8 provided a new BACKUP utility using FlashCopy technology to take very fast backups of the entire subsystem without any disruption. RECOVER is only for the whole subsystem.

With Vnext we expect to be able to recover an object, rather than the whole subsystem. This makes the job of backup and recovery simpler and easier.

Converged TEMP Space

- Single source for all temporary space in DB2, replacing: DSNDB07, temp databases, workfile database
- Access is virtualized for small amounts of data, eliminating cost of work file creation (reduced CPU and I/O)
- Supports 4K and 32K page sizes, with automatic selection of the appropriate page size

Temporary storage architecture and use

Many recent improvements have enhanced the use of temporary storage when it is needed by the DB2 engine, for example, for interim materialized result sets. The different temporary pools are converged to a single source. Use of temporary storage is eliminated when processing small amounts of data.

DDL Porting Improvements

- Automatic selection of DATABASE and TABLESPACE when DDL omits these keywords
- Automatic CREATE of UNIQUE index for PRIMARY KEY
- Deprecated simple table space, default to segmented

Additional family compatibility

Additional compatibility will be added for:

- Default databases and table spaces
- Automatic unique indexes to support defined primary keys
- The default for a table space will change to segmented, improving performance and manageability for many customers.

DB2 Vnext Themes

- Enable high-volume transaction processing for next wave of Web applications**
- Extend the lead in transaction processing availability, scalability and performance**
- Reduce cost of ownership and zSeries-specific skill needs**
- Improve data warehousing and OLTP reporting**

Themes for VNext: There are four primary themes for the VNext enhancements. The first is enabling high volume transaction processing for the next wave of web applications. This work extends DB2's ability to handle the next generation of applications by adding XML support and extending SQL and security. The SQL improvements also improve the ability to port applications to DB2 for z/OS and DB2 family compatibility.

The second theme is extending the lead in transaction processing availability, scalability and performance. The key changes in availability extend DB2's ability to make changes to the data definitions without disrupting the online processing. New structures are provided to improve scalability and performance. Bottlenecks for very high performance are removed.

The third theme is reducing the cost of ownership and the need for skills that are specific to the zSeries platform. This theme includes productivity improvements for administrators through consistency with the DB2 family and the relational database management industry. It also includes cost reductions through processing reductions.

The fourth theme is 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. More queries can be expressed in SQL with new SQL enhancements.

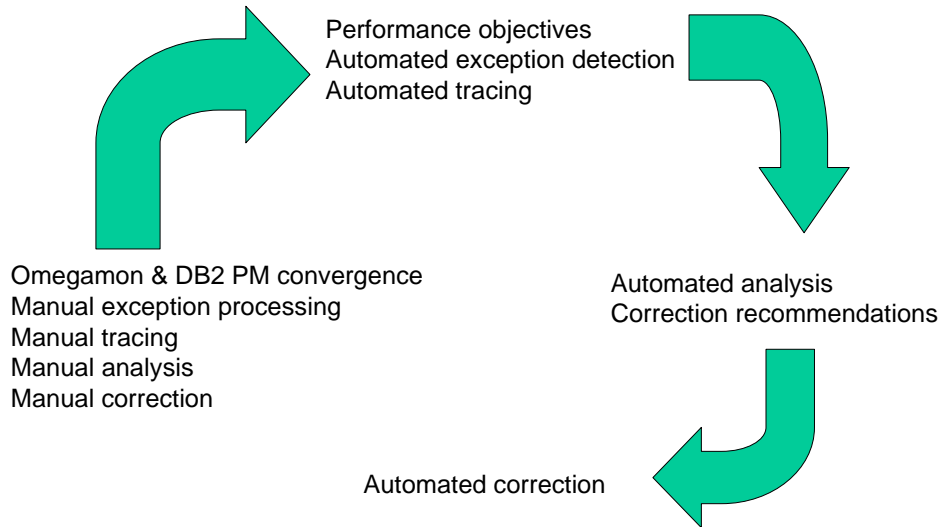
Data Warehousing, Reporting and Optimizer Improvements

- Cross query block optimization
- Histogram statistics exploitation
- Generalize sparse index and in-memory data cache method
- Dynamic Index ANDing for Star Schema
- EXCEPT and INTERSECT
- Many other SQL improvements

The fourth theme is 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. More queries can be expressed in SQL with new SQL enhancements.

Improved data is provided for the optimizer, with improved algorithms. The INTERSECT and EXCEPT clauses make SQL easier to write. Other SQL statements improve consistency with the DBMS industry.

DB2 Performance Monitoring



Work to make more of the performance monitoring and tuning work autonomic continues. One of the key steps is the convergence of our Omegamon and DB2 PM monitors, taking the best from each line.

TCO Improvements – DBA tools

- Autonomic Policy-based SQL query management/monitoring:
 - Automatic collection of performance data for long running queries
 - Automated query monitoring for the most frequent/expensive queries
 - REOPT(SMART)
 - Real time statistics exploitation by Optimizer
- Optimization Service Center (Web-based DBA admin – no 3270 screens)
 - DBA tool suite for tuning/managing SQL queries (Stats Advisor, Index Advisor, Query Rewrite Advisor, Query Workload Monitor, Resource Estimator, Query Formatter, Visual Explain, Visual Plan Hint, IBM Service Doc Generator, Partitioning/Clustering Advisor)
- Query Performance Warehouse
 - Execution history of queries
 - Identification of query patterns
 - Identification of usage patterns for tables/indexes

The optimizer service center (OSC) will initially capture meaningful information about your workload, and providing helpful serviceability features by exposing insightful optimizer information.

In a following phase it will analyze what is running and make specific workload recommendations in order to reach optimal performance, suggesting physical database design changes, like indexes, MQTs, and partitioning; statistics to be collected; and system configuration changes.

At the next level, OSC can go ahead act on your behalf, if permitted, and implement the changes autonomically for you. More of the centers and advisors will be exposed for administrator use. This will free administrators to address more important tasks, like delivering new features that further leverage your IT investment.

DB2 for z/OS Vnext

- Integration ➤ XML, Unicode, LOBs
- Availability
- Scalability ➤ SQL for DB2 family
- Productivity
- Total cost of ownership ➤ Data Definition On Demand



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


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, Vnext 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 for z/OS Vnext is to enhance DB2's ability to handle new and enterprise applications. Vnext improves the ability to handle new applications with XML, large objects, and many SQL and security improvements. Vnext 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.