

What's New in DB2 for z/OS, Version 8 and Beyond

Roger Miller

Monday March 6, 2006

Session 1302

Roger Miller will discuss the latest news about DB2 for z/OS Version 8 (V8), which has been generally available for about two years now. Then the discussion will move to trends and directions for the next version, which we'll call Vnext. 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. The vision for 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.

DB2 for z/OS Version 8 News



- **New function**
 - **Cross loader with LOBs**
 - **Built in functions ASCII, TIMESTAMPDIFF**
 - **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)**
- **now** {
 - **New and updated books: Library refresh April 2005, Mar 2006**
 - **Messages, Codes became separate books August 2005**
 - **Redbooks: Design Guidelines for High Performance and Availability, Business Value, Performance Topics, WebSphere, MLS, Disaster Recovery, others updated ...**
 - **Customer information on the web**

Copyright IBM Corporation Author Roger Miller

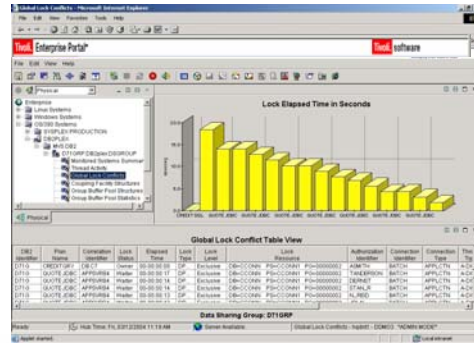
2

- **Cross loader with LOBs V7 V8: PQ90263**
- **Built in functions ASCII, TIMESTAMPDIFF 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**
- **LOAD, UNLOAD with LOBs V7 V8: PK10278 open**
- **zIIP PK18454, PQ19920, PQ19921 open, ...**
- ✓ **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/e0z2e121.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>
- ✓ **Stored Procedures** <http://www.redbooks.ibm.com/redbooks/pdfs/sg247083.pdf> **updated Feb 2006**
- ✓ **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**
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246465.pdf> **updated Feb 2006**

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 exploitation
 - DB2 Connect reporting/monitoring
 - Performance warehouse (historical data mining)
 - DB2 to CICS transaction linking
 - History monitoring
 - Event exceptions
 - Threshold checking



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

Copyright IBM Corporation. Author Roger Miller

3

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.

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
- ✓ ...



Copyright IBM Corporation Author Roger Miller

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

4

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.

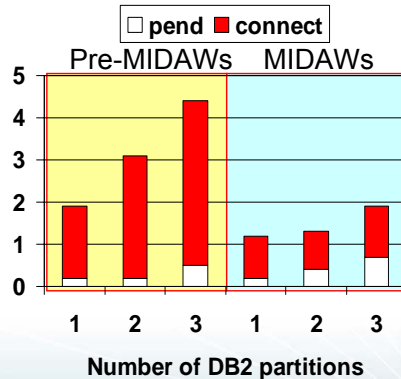
DB2 benefits from large real memory support, faster processors, and better hardware compression. DB2 uses Parallel Access Volume and Multiple Allegiance features of the IBM DS8000 and Enterprise Storage Server™. FlashCopy® can be used for DB2 backup and restore. DB2 makes unique use of the z/Architecture™ 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)





I/O Response Time (sec)



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



Copyright IBM Corporation Author Roger Miller

5

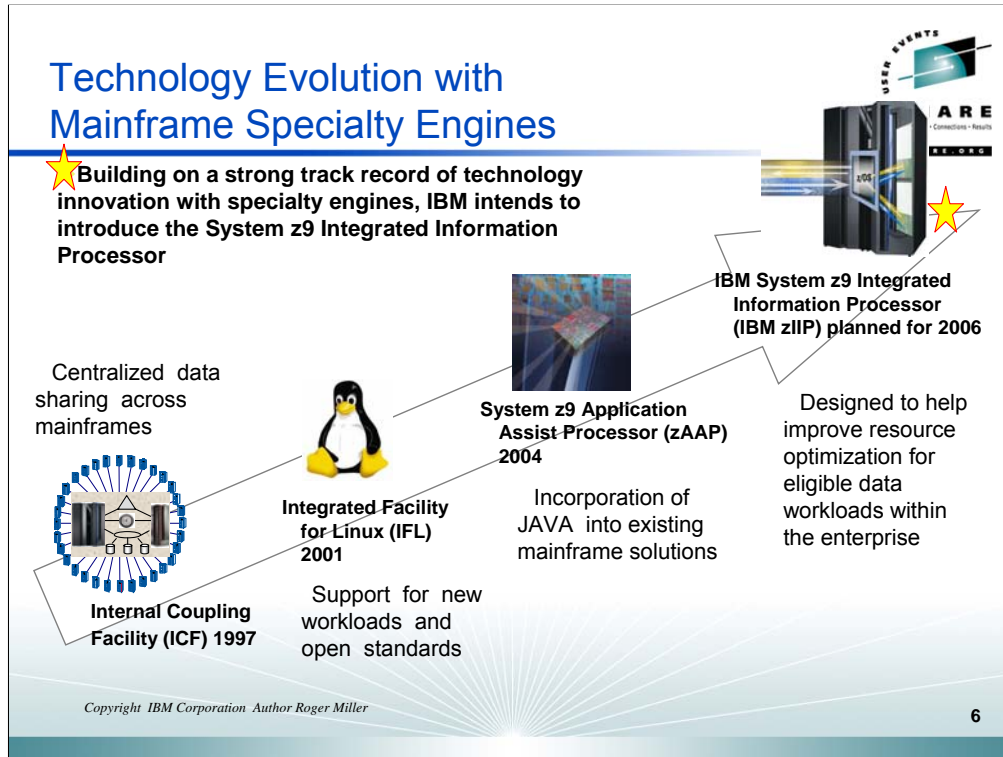
These are some preliminary measurements with new processors, new software, new channel configurations and new disks. Note the sustained scan rates of 100 megabytes per second or 170 MB / sec. for parallel access on one channel.

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.



System z9 zIIP web site

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

This site includes a zIIP FAQ:

http://www.ibm.com/systems/z/faq/pdf/ziip_faq.pdf and a press release with later information:

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

Webcast on IBM System z9, zIIP, DB2 Vnext

<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>

Article in ESJ:

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

Article in ComputerWorld

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

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

Article in NetworkWorld: <http://www.networkworld.com/news/2006/020606-ibm-db2.html?nettx=020606netflash&code=nlflash21301>

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 z9-109 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 IBM user of the zIIP with
 - System z9 109
 - z/OS 1.6 or later
 - DB2 for z/OS V8
- Webcast replay ibm.com/servers/systems/z/2006/



Copyright IBM Corporation. Author Roger Miller

7

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



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.

Copyright IBM Corporation Author Roger Miller

8

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. 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.

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:

- Select query processing of BI, ERP or CRM network-connected applications;
- 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>

- 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.

DB2 SQL

z/OS V7

common

LUW Linux, Unix & Windows V8.2

z { Range partitioning

c { 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

o {

n { 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

L

U

W

Copyright IBM Corporation Author Roger Miller

10

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.

DB2 SQL

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 { 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

o {

m {

n {

L { 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

U {

W {

Copyright IBM Corporation Author Roger Miller

11

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.

DB2 SQL

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

U

W Copyright IBM Corporation Author Roger Miller

12

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.



- XML
- MERGE
- SELECT FROM UPDATE, DELETE, MERGE
- TRUNCATE
- INSTEAD OF TRIGGER
- BIGINT, VARBINARY, DECIMAL FLOAT
- Native SQL Procedure Language
- 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

Copyright IBM Corporation Author Roger Miller

13

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



Copyright IBM Corporation Author Roger Miller

14

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.

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

Copyright IBM Corporation Author Roger Miller

16

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, and include an XQuery interface to the data.

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 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

Copyright IBM Corporation Author Roger Miller

17

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 character 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")
```

Copyright IBM Corporation Author Roger Miller

18

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.

The SQL/XML publishing functions provided in DB2 V8 are enhanced, and many new functions and predicates are added.

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
 - Tables to store XML schemas
 - Stored procedures to register XML schemas
- DSN_XMLVALIDATE() SQL/XML function
 - Test XML values for validity against XML schema
 - Obtain default values and schema normalized values from XML schema
- XML decomposition using annotated XML schema

Copyright IBM Corporation Author Roger Miller

19

The utilities are extended to support XML. You can LOAD, UNLOAD, CHECK DATA or INDEX, COPY, REBUILD, RECOVER and REORG the XML data.

The XML schema support handles a schema repository and uses tables to store XML schema information. Stored procedures are provided to register XML schemas.

Decimal Floating Point



- New datatype DECFLOAT
 - 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

Copyright IBM Corporation Author Roger Miller

20

New data type support is provided. Support will be added for:

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

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

Copyright IBM Corporation Author Roger Miller

21

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

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**

Copyright IBM Corporation Author Roger Miller

23

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

Copyright IBM Corporation Author Roger Miller

24

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.

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

Copyright IBM Corporation Author Roger Miller

26

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.

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
- APPEND option for table space

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

Copyright IBM Corporation Author Roger Miller

28

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
 - Log latch contention relief
- LOB performance, function, scalability
- 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**

Copyright IBM Corporation Author Roger Miller

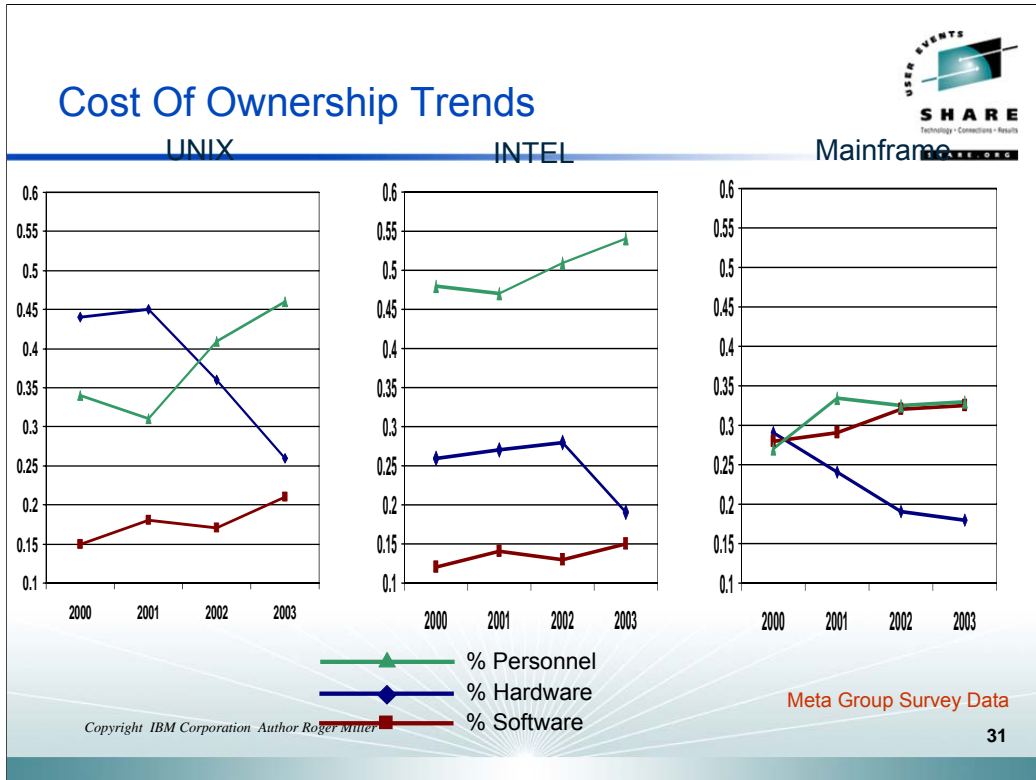
30

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.



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.

Some key implementations

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

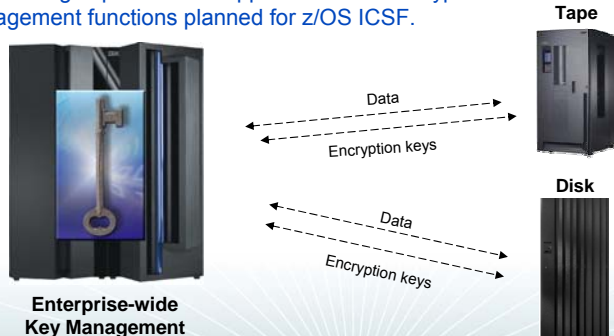


While DB2 for z/OS V8 provides many enhancements for security, there are still many more needs and much more work to do. Roles are used to provide a more flexible technique than groups or users in assigning and controlling authorization, while improving consistency with the industry. A network trusted 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.

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
Copyright IBM Corporation Author Roger Miller

MFE_140

33

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);
```

Copyright IBM Corporation Author Roger Miller

34

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.

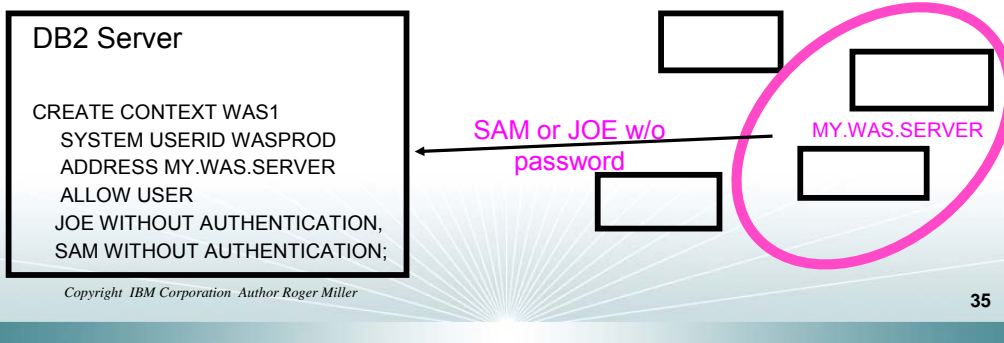
The role can be assigned and removed from individuals via the trusted authorization context as needed. This allows a DBA to perform object maintenance during a change control window on a Saturday night, for example. But when Monday arrives, they do not have the authority to do this same work.

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*.

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.

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**

Copyright IBM Corporation Author Roger Miller

37

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



- SQL enhancements: INTERSECT, EXCEPT, RANK, caseless comparisons, cultural sort, ...
- 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

Copyright IBM Corporation Author Roger Miller

38

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 for z/OS Vnext



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



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

Copyright IBM Corporation Author Roger Miller

39

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, zArchitecture, 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

Copyright IBM Corporation Author Roger Miller

40

1

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.