IBM Information >>> On Demand 2007

IBM

DB2 9 for z/OS Overview

April 10, 2009

Act.Right.Now.

IBM INFORMATION ON DEMAND 2007
October 14 – 19, 2007
Mandalay Bay
Las Vegas, Nevada

This presentation is an overview of DB2 9 for z/OS. DB2 9 has been generally available since March 2007, after 9 months for a beta program. This presentation will include a few other new items, such as the new IBM System 10 and z9 Integrated Information Processor (zIIP) and the new z10 and z9 processors. Data definition on demand extends the theme of online schema revolution from V8. Additional Unicode and text enhancements continue the work from V7 and V8. XML work across the DB2 family is a much larger step than in V7 or V8. SQL Procedures become more consistent across the family. While V7 and V8 removed many differences from DB2 for Linux, UNIX & Windows, DB2 9 takes the next big step to improved productivity and consistency. Utility enhancements help with new function, more LOB and XML support, better performance and improved availability, removing the BUILD2 step from online reorg. DB2 9 enhances DB2's ability to handle new and enterprise applications. DB2 9 improves with XML, large objects, and many SQL and security improvements. DB2 9 builds upon and extends DB2 traditional strengths and the ground-breaking V8 in many areas: online schema evolution, Unicode, XML, DB2 family SQL, utilities, security and 64-bit virtual storage.

DB2 9 for z/OS
Addressing Corporate Data Goals

SOA Enablement	• pureXML • Optimistic locking for WebSphere • LOB performance, usability
Dynamic Warehousing	• Many SQL improvements • Dynamic index ANDing • Histogram statistics • New built-in OLAP expressions • Optimization Service Center
Simplification, Reduced TCO	• Index compression • Partition By Growth tables • Cloned tables • Volume based backup / recovery
Workload Consolidation	• More online schema changes • Online REBUILD INDEX • Trusted context and ROLEs • Parallel Sysplex clustering improvements

2

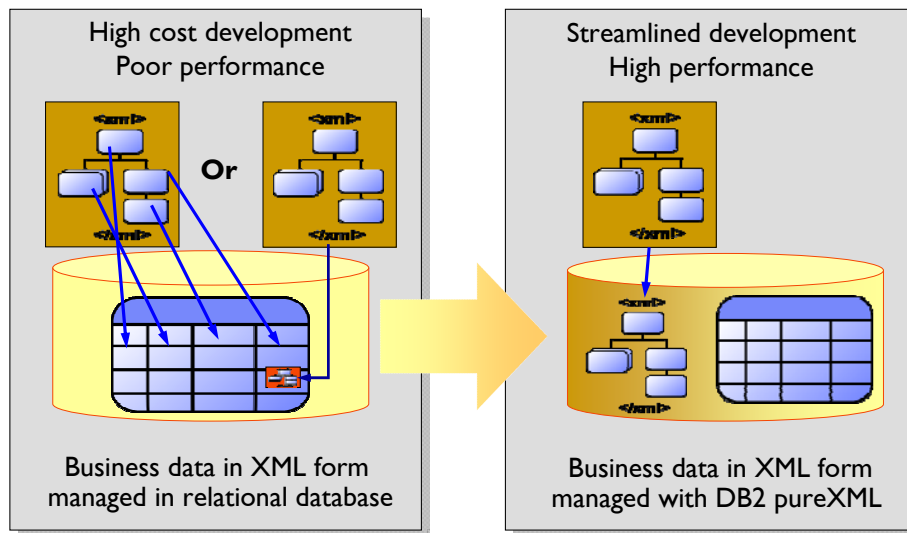*IBM INFORMATION ON DEMAND 2007*

Act Right. Now.

One of the key initiatives of V8 was online schema evolution, and that theme is expanding and changing to be data definition on demand. These are key improvements for resilience. One of the important changes is to be able to replace one table quickly with another. Another is to be able to rename a column or an index. A new type of table space combines the attributes of segmented and partitioned, without a partitioning key. Rebuild index can be run with much less disruption. 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. Changing DB2 early code does not require an IPL.

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 new utilities.

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. Improved cpu and elapsed times can be achieved with the FETCH FIRST clause specified on a subquery. The INTERSECT and EXCEPT clauses make SQL easier to write.

DB2 9: *A New Generation Hybrid Data Server*

High cost development
Poor performance

Or

Business data in XML form
managed in relational database

Streamlined development
High performance

Business data in XML form
managed with DB2 pureXML

3

Key point: The amount of business information in XML form is already as great or greater than other forms and growing faster - failure to leverage efficiently as structured data means high cost and/or missed opportunity. DB2 9 provides the best of both worlds, pureXML^{tm} for native storage and integrating XML with object-relational. Performance, integrity, protection, and scale from the proven DB2 infrastructure with the flexibility of XML/XPath and relational/SQL. This overcomes the complexity & limitations of prior models (shred, CLOB, or XML only). In 2006 IBM introduced a new generation data server with the availability of DB2 9. The explosive growth of XML based data standards in all industries means competitive advantage for those businesses that use it most effectively and efficiently. Client, policy and claims processing in Insurance; supply chain management in Retail; financial transactions and asset management in Banking; patient care in Healthcare; citizen service in Government; implementing Service Oriented Architectures (SOA) in Computing Software and Services - and many other processes across all industries - increasingly rely on information captured and exchanged in XML form. Our clients are increasingly managing XML format text documents in a content management system for proper governance and efficient use in the business process workflow.   But few are realizing the full value of all the business data they possess that are in XML format.

Early users of the pureXML feature of DB2 9 are taking advantage of the fact that data in XML format is well structured and can be queried via standard languages such as XPath and XQuery. By doing so they are bringing that data to bear in both transactional and analytic processes - with higher performance and lower development costs than previously possible with a relational database.   The difference is that DB2 9 supports both relational (tabular) and XML (hierarchical) structures in the same database so that both can be easily, efficiently and securely managed, analyzed and delivered. Unlike other relational data servers - and previous versions of DB2 - pureXML eliminates the overhead of fitting the "square peg" XML tree structure into the "round hole" row and column relational structure.

Learn more from the DB2 XML Wiki at:

http://www.ibm.com/developerworks/wikis/display/db2xml/Home

DB2 XML partners:

AgencyPort: AgencyPort is the leading provider of web-based new business, endorsement and renewal solutions designed specifically for property and casualty insurers.

Altova - XML development tools: Altova is the provider of the world's leading XML development tools, including XMLspy the industry standard XML development environment.

Coral8 - Complex Event Processing

Coral8: The Coral8 Engine enables you to build and deploy powerful Complex Event Processing (CEP) applications for any enterprise business function that can benefit from real-time awareness.

Exegenix - Document Conversion to XML: Exegenix technology converts content into XML quickly, accurately, and cost-effectively, with no re-keying, and no dependency on consistently-formatted input.

JustSystems - Rapid XML Application Development: The combination of xfy from JustSystem and DB2 9 allows users to quickly assemble XML applications to access and manage a wide range data.

Napersoft - Customer Communications Management: Napersoft provides Customer Communications Management solutions with its family of WebDemand360 Correspondence Manager for WebSphere offerings and DB2 9.

Nextance - Contract Performance Management: Nextance offers Contract Performance Management solutions that take full advantage of DB2 9 XML capabilities.

DB2 9 for z/OS pureXML *Industry Bundles*

- *Quick start for customers and partners based on popular industry standard XML*
  - ACORD for Insurance (enabled for DB2 for z/OS)
  - FIXML (enabled for DB2 for z/OS) for Banking
  - CDISC for Clinical Data
  - FpML and MISMO for Financial Services

- *Demonstrations and samples using industry standard XML formats with DB2 pureXML*

- *Data Model with registered XML schemas and XML documents stored in DB2 pureXML database*

- **Sample XML queries**

- **Example services**
  - Create, Replace, Update, Delete and Query
  - Atom feeds
  - XForms

Learn more:
http://www.alphaworks.ibm.com/tech/purexml

5

IBM INFORMATION ON DEMAND 2007

Get a fast start with XML from industry and other samples and examples.

Until DB2 9, managing XML data records with a relational data servers meant decomposing the data into columns - a process known as shredding. Or by storing the entire data record in a single cell as a character large object - known as a CLOB. The CLOB approach does not cost overhead as the data records go in. But when you query these records you pay the overhead of parsing each one at runtime which can be a significant performance impact to the application. With shredding, overhead is paid up front to turn the data into a relational record that can be queried efficiently. But overhead is also paid later if the record needs to be recreated for delivery in XML format. This process also affects the fidelity of the record itself - leading to an approach that uses both shredding and CLOB methods for applications that require both performance and fidelity. This results in even more overhead to ensure the records remain in sync.
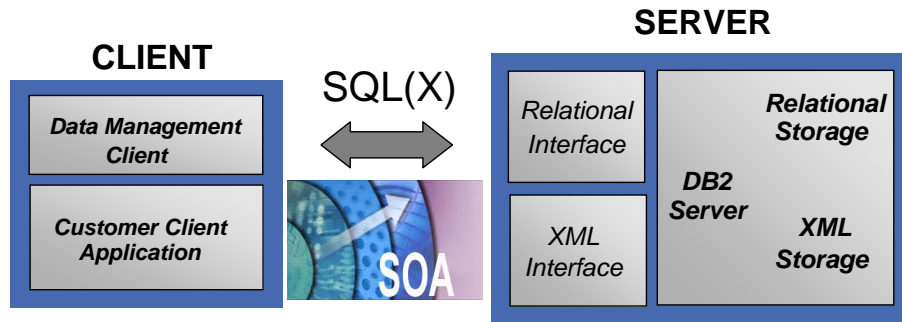
The impact of pureXML is seen by a large Banking client with a requirement to update over 500,000 XML data records per day. Attempts to use a competitors relational data server failed. Using DB2 9 with pureXML, the application was able to update more than half a million data records in less than an hour.

And a large Insurance client has seen the impact of pureXML to development time and cost with a 65% reduction in lines of code and more than 75% reduction in time required to develop services accessing XML data.

Capabilities Inside the Engine

Performance, Performance, Performance

SERVER

CLIENT

SQL(X)

Data Management Client

Customer Client Application

SOA

Relational Interface

XML Interface

DB2 Server

Relational Storage

XML Storage

Native storage   Schema Index Functions Utilities

6

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

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 9 for Linux, UNIX and Windows delivers this picture, and adds an XQuery interface to the data.

DB2 9 for z/OS Performance Improvements

- CPU reductions in utilities: LOAD, REORG, …
  - Online REORG with no BUILD2 phase
- LOB performance, function, scalability
- SQL and optimization improvements, inserts
- Use zIIP more
  - Remote native SQL Procedure Language
  - Increased parallel work
- Improved varying length performance
- Sequential disk access
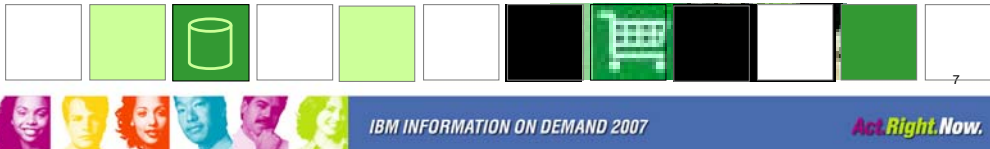
The key performance improvements in DB2 9 are reduced CPU time in the LOAD and REORG utilities, improved LOB performance and scalability, improved optimization for SQL, the zIIP processing for remote native SQL procedures, reduced CPU time for data with varying length and better sequential access.

. Improved data is provided for the optimizer, with improved algorithms and a rewritten approach to handling performance exceptions, called Optimization Support Center and Optimization Expert. Indexes can be created on expressions, rather than just on columns. Index compression is able to shrink disk space by half..
V8 SQL procedures were not eligible to run on the zIIP, but changing to use the native SQL Procedure Language on DB2 9 will make the work eligible for zIIP processing. Varying length data can improve substantially if there are large numbers of varying length columns. Several improvements in disk access can reduce the time for sequential disk access.

Most consumable DB2 9 improvements

- CM very little to no action:
  - Utility CPU reductions
  - Logging improvements
  - Improved index page split
  - Larger prefetch, write & preformat quantities
  - LOB performance
  - DDF VSCR
  - Optimization Service Center, Opt. Expert, & Data Studio

- Changed online REORG
- Improved RUNSTATS
- Optimization improvements, EDMPOOL VSCR

NFM
- LOB lock avoidance
- Reordered row format
- Index: larger page sizes, compression, index on expression

Here are some highlights for items that deliver the most quickly and easily:

Very little to no action is required for the utility CPU reductions, logging improvements, improved index page split, larger prefetch, write & preformat quantities, some LOB performance, DDF virtual storage constraint relief. The first group delivers in CM.

The next items require some work. Changed online REORG and other utility improvements require process changes and use of SHRLEVEL(CHANGE).

Improved RUNSTATS statistics needs some analysis to determine where the value is greater than the cost of gathering the new statistics.

Optimization improvements are automatic for dynamic SQL, but require work to REBIND for static SQL. In both cases, we need baselines to check for regression. REOPT(AUTO) for dynamic SQL needs analysis to be sure the improvement is working. EDMPOOL virtual storage constraint relief also requires a REBIND.

Optimization Service Center takes some learning, but should be fast for those who have used Visual Explain in the past. See the new redbook, SG24-7421, DB2 9 for z/OS: New Tools for Query Optimization.

Reordered row format requires a REORG in NFM and varying length columns.

Index improvements for larger page sizes, compression, index on expression require database design work to determine where they are applicable. ALTERs, REORGs and creation of new indexes are needed.

LOB lock avoidance requires a quiesce of all subsystems in NFM.

DB2 9 z9, z890 & z990
performance plan example scenario
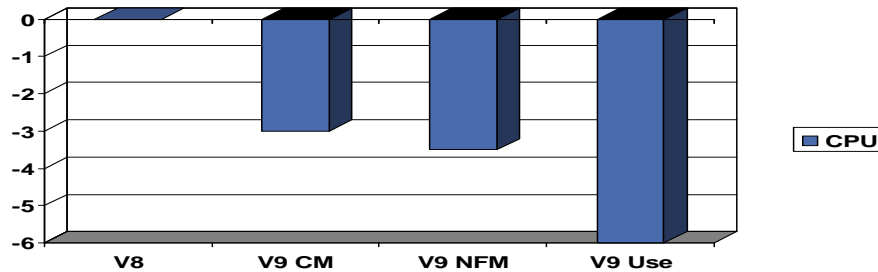
Utilities   DB design adjustments
Histogram statistics   Index improvements
REBIND   application changes
DSNZPARMS   native SQL procedures
                                    SQL adjustments
Your situation will vary. Less CPU is better.
z800 and z900 expect +5% to +10% CPU

If you have a z9, z990 or z890, this is expected to be a common shape for a DB2 9 performance plan, starting with zero for the V8 baseline. When you first move to DB2 9, total DB2 CPU time generally decreases from 0% to 5% for z9, z890 and z990 customers, shown here as a first step -3%. Utility CPU reductions help immediately. Some work will be about the same (+/-3%). Start with reorgs and collect improved histogram statistics when useful. The DB2 9 CM performance plan REBINDs the primary packages and adjusts DSNZPARMs. The REBINDs provide most of the improved access paths. On z800 or z900 the initial cpu expectation is +5 to +10% regression, more if there are many columns, so making adjustments is more important.

In moving to NFM, some additional DSNZPARMS are adjusted and all plans and packages are rebound. The DB2 9 use line takes wider advantage of DB2 9 performance improvements. Database designs start taking advantage of new indexing options, such as compression, index on expression and larger pages. After making the design changes, REORG the data and REORG or REBUILD the indexes, get the improved statistics and REBIND. Native SQL procedures, added use of zIIP, and improved SQL continue the improvements in this phase.

Scenario: Customer mix of DB2 CPU time is 30% in utilities, 70% in SQL access. With 10% improvement for the utilities, we get a -3% net, assuming that SQL is the same as before. With optimization improvements, another -½% improvement shows up in DB2 9 NFM. Then as design adjustments, reorgs and rebinds are performed, we get improvements from varchar improvements, native SQL procedures and improved SQL, another -3%.

SQL: Productivity, DB2 family & porting

Portable SQL

- XML
- MERGE & TRUNCATE
- SELECT FROM UPDATE, DELETE, MERGE
- INSTEAD OF TRIGGER
- BIGINT, VARBINARY, BINARY, DECIMAL FLOAT
- Native SQL Procedure Language
- Nested compound
- Optimistic locking

- LOB File reference variable & FETCH CONTINUE
- FETCH FIRST & ORDER BY in subselect and fullselect
- INTERSECT & EXCEPT
- ROLE & trusted context
- Many new built-in functions, caseless comparisons
- Index on expression
- Improved DDL consistency
- CURRENT SCHEMA

10

IBM INFORMATION ON DEMAND 2007   Act Right.Now.

As in Version 8, there are many improvements for SQL and for XML in DB2 9. Improvements in the SQL have made migrating from other platforms, such as Unix and Windows much easier.

DB2 9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on the prior slide. There are new SQL data manipulation statements in MERGE and TRUNCATE. There are new data types with XML, DECIMAL FLOAT, BIGINT, BINARY and VARBINARY types. Improvements in LOBs provides more consistent handling and improved performance. Intersect and Except set operations make some SQL operations simpler to specify. Security is improved with ROLEs and network trusted context. Data definition consistency and usability are improved. DB2 9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

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

common — 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

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 or DELETE, multi-site join, MDC

11

IBM INFORMATION ON DEMAND 2007   **Act Right.Now.**

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.

New cross-platform SQL Reference on web now

DB2 SQL

z z/OS 9
common
luw Linux, Unix & Windows 9

SELECT FROM T1 INSERT INTO... Portable SQL

z Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from 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 or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables**, range partitioning, compression

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

12

IBM INFORMATION ON DEMAND 2007   Act **Right**. **Now.**

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS. DB2 9 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 9 for LUW, code named Viper. We are able to move more from the z list to the common list with Viper.

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

This is the web page for the SQL Reference for Cross-Platform Development. It has four versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well. If you want to write programs that work well across the DB2 family, then this is the book to download and use.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions. The latest book is Version 3.1 for platforms DB2 9 for z/OS, DB2 for I V6.1, and DB2 9.5 for Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

• DECFLOAT data type   Additional timestamp format
• ROW CHANGE expressions   New built-in functions
• order-by-clause and fetch-first-clause in a subselect
• SELECT FROM INSERT   Full outer join
• IMPLICITLY HIDDEN columns   row-change-timestamp columns
• RESTRICT on DROP of functions and procedures
• CURRENT DECFLOAT ROUNDING MODE special register & SET statement
• ALTER FUNCTION   Other portability enhancements

Other books are for DB2 for z/OS Version 7 and Version 8.
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv31.pdf
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
http://www.ibm.com/developerworks/data/library/techarticle/db2common/?S_TACT=105AGX11&S_CMP=LP

DB2 SQL

z   z/OS 9
common
luw Linux, Unix & Windows 9.5

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, TRUNCATE, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, index compression

common — 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, 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 or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitions, data compression, **session variables, DECIMAL FLOAT, optimistic locking, ROLE**

luw — Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, **XML enhancements, array data type, global variables, Oracle syntax**

14

IBM INFORMATION ON DEMAND 2007

Act Right. Now.

This text just shows the newest delivered relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS, (DB2 9). DB2 9 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, code named Viper 2 is already generally available. We are able to move more from the unique z list to the common list with DB2 9 for luw.

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. The changes in a specific version are not consistent. As we introduce new function, sometimes it will be on one platform first, but movement from unique lists into the common list continues to be the strongest trend.

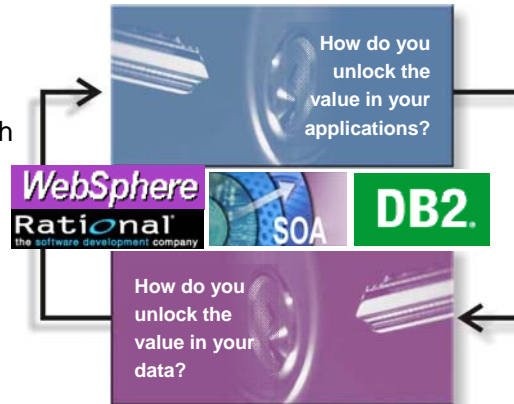## SOA and Information on demand
What the market is telling us

Huge investment in traditional programs & data
- "200 Billion lines of COBOL …"
- "more than a trillion dollars' worth of legacy mainframe code."

The key to unlocking value:
- Unstructured and structured
- Easy to find & use, tailored, cost-effective & integrated

How do you unlock the value in your applications?

How do you unlock the value in your data?

WebSphere
Rational *the software development company*
SOA
DB2.

It's time to unlock the value of these assets

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

Significant investment in traditional assets:

"200 Billion lines of COBOL code in existence" eWeek May 28[th] 2005

"Mainframe users are sitting on more than a trillion dollars' worth of legacy mainframe code."
Computerworld, 2006

"Majority of our customer data still on mainframes"
Don Greb, Mellon Financial Corp from Computerworld 2002

The key to unlocking the value is often in the WebSphere, Rational and DB2 portfolios:

85% of information is unstructured

48 disparate financial systems and 2.7 ERP systems in the average $1B company

Only 1/3 of CFOs believe that the information is easy to use, tailored, cost-effective or integrated

30% of people's time is spent searching for relevant information

Sources: IBM & Industry Studies, Customer Interviews

Business Partners: Many enhancements

- SHRLEVEL(REFERENCE) for REORG of LOB table spaces
- Online RENAME COLUMN
- Online RENAME INDEX
- Online CHECK DATA and CHECK LOB
- Faster REORG by intra-REORG parallelism
- More online REORG by eliminating BUILD2 phase
- LOB Lock reduction
- Skipping locked rows option
- Online REBUILD INDEX
- Change SCHEMA & VCAT
- Tape support for BACKUP and RESTORE SYSTEM utilities
- Recovery of individual tablespaces and indexes from volume-level backups
- Enhanced STOGROUP definition
- Utility TEMPLATE switching
- Conditional restart: automatic search for appropriate checkpoint
- CLONE Table: fast replacement of one table with another
- Buffer management by WLM

- Global query optimization
- Generalizing sparse index and in-memory data caching method
- Optimization Service Center
- Autonomic reoptimization
- Logging enhancements
- LOBs Network Flow Optimization
- Faster operations for variable-length rows
- NOT LOGGED table spaces
- Index on expressions
- Universal Table spaces
- Partition-by-growth table spaces
- APPEND option at insert
- Autonomic index page split
- Index page sizes 8K, 16K, 32K
- Support for optimistic locking
- Faster and more automatic DB2 restart
- MODIFY RECOVERY enhancements
- RLF improvements for remote application servers such as SAP
- Preserving consistency when recovering individual objects to a prior point in time

- DECIMAL FLOAT, BIGINT
- VARBINARY, BINARY
- TRUNCATE TABLE statement
- MERGE statement
- FETCH CONTINUE
- ORDER BY and FETCH FIRST n ROWS in sub-select and full-select
- ORDER OF extension to ORDER BY
- Various scalar functions
- XML support in DB2 engine
- Native SQL Stored Procedures, able to use zIIP
- SELECT FROM UPDATE/DELETE/MERGE
- Enhanced CURRENT SCHEMA
- IPv6 support
- Unified Debugger
- Network Trusted Context
- Database ROLEs
- Automatic creation of database objects
- Modify early code without requiring an IPL
- Utilities CPU reduction
- Temporary space consolidation
- . . .

TCO Reduction   Continuous Availability   Performance   Scalability   SQL   Portability   16

IBM INFORMATION ON DEMAND 2007   Act Right.Now.

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

DB2 9 Availability & Scalability Improvements

New features enhance availability & scalability of the DB2 9 environment to manage your 7x24 infrastructure easily

- New and improved DB2 capabilities
  – Improved INSERT performance
  – New universal table space
  – Partition by growth
  – Not logged table space
  – Online schema change enhancements
  – Clone tables
  – Data sharing enhancements
- More online utilities
  – Online Reorg enhancements
  – Online Check Data, Check LOB
  – Online Rebuild Index
- Memory improvements 64 bit++

Land Registry
Cymraeg

UK government Land Registry runs DB2 for z/OS to maintain the world's largest known transaction processing (OLTP) database – 23.1 TB! [1]

17

1 - Winter Corporation's "2005 Top Ten" awards - http://www.wintercorp.com/index.html

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

DB2 9 offers a broad range of improvements in both availability and scalability. INSERT, UPDATE and DELETE performance is improved substantially with faster logging, and reduced index page splitting. The Universal Table Space (UTS) provides a structure that is both partitioned and segmented for better performance. UTS includes both Partitioned by Range and Partitioned by Growth table spaces. PBR is range partitioning, but using the improved segmented page structure. Partitioned by Growth helps in situations where the data is large or might become large, but there is no good partitioning key. As the table grows, partitions are added. PBG has most of the utility benefits of partitioning. The not logged table space is for situations where there are massive, parallel inserts and recovery is not needed. Be careful, you can dig yourself a very deep hole.

DB2 9 adds to the list of attributes which can be altered, rather than requiring the table space to be dropped and recreated. Now we can rename columns and indexes, alter a column to set a default, alter logging, and change STOGROUP SMS constructs. Clone tables allows us to have two similar tables and then to exchange one for the other very quickly. Customers asked for an online LOAD REPLACE technique, but often the LOAD utility is not used, so the capability is implemented with ALTER.

Online reorganization for a few partitions is much more online, with no BUILD2 phase. Check data, Check LOB and REBUILD INDEX utilities now have an online option. More memory is moved above the 2 GB bar, giving some VSCR help. More memory is moved above the 2 GB bar, giving some VSCR help.
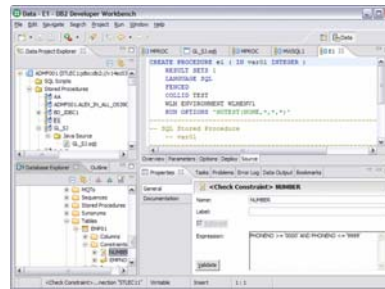
DB2 9 Easier Application Development

Simplify development process
Improve performance
Numerous SQL enhancements
  – Merge / Truncate
  – Select from update, delete & merge
  – Instead of triggers
  – BIGINT & DECIMAL FLOAT
  – New Built-in functions e.g. SOUNDEX
  – DDL porting improvements
  – Optimistic locking
- Developer Workbench
- Spatial or geographic data
- Text Search server included
- Native SQL Stored Procedures

18

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

A tremendous number of enhancements helps application developers to simplify the development process and gain performance improvements when accessing DB2 data. As the DB2 SQL is more consistent with use across the industry, porting an application to DB2 9 for z/OS is much simpler. The learning process is also much faster for those who know only some other DBMS or platform. The additional function delivered in DB2 9 means that applications can use what is provided, rather than delivering this level.

Thes SQL enhancements are noted on prior slides, but they include new SQL statements, new data types, and many new functions.

IBM DB2 9 for z/OS has been successfully certified to support SAP systems: https://www.sdn.sap.com/irj/sdn/db2?prtmode=print

▪After a formal verification process of the functionality and quality of the new DB2 database release with regard to SAP environments and its products SAP solutions are now supported on DB2 9 for z/OS.
The new database version has well passed SAP's validation tests on SAP R/3 4.6, SAP NetWeaver '04, SAP NetWeaver 7.0 and SAP ABAP stack 7.1. SAP releases DB2 9 for z/OS for all SAP solutions based on the current SAP technologies.
For migration to DB2 9 and prerequisites please refer to the SAP note 1043951 and the BestPractices document in the Knowledge Center. New IBM Redbooks are available (see IBM Information Sources page).

▪Enhancing SAP - DB2 9, SG24-7239,
http://www.redbooks.ibm.com/abstracts/SG247239.html
▪Best practices SAP BI - DB2 9, SG24-6489-01,
http://www.redbooks.ibm.com/redpieces/abstracts/sg246489.html

Dynamic Warehousing
A New Approach to Leveraging Information

Information On Demand to Optimize Real-Time Processes — *Action* — *Dynamic Warehousing*

Dynamic Warehousing Requires:

1. Real-time access – in context
2. Analytics – as part of a business process
3. Unstructured information – extracted knowledge
4. Extended capabilities – tightly integrated

OLAP to Un... Recom...

Query & to Un... What Happened — *Reporting*

20

IBM INFORMATION ON DEMAND 2007 — Act.Right.Now.

Traditional warehousing focused on query and reporting to understand what happened, and evolved to enable OLAP and data mining to understand the why those things happened and recommend future action.

Dynamic warehousing is a new approach to address the primary business challenges that organizations face today, which requires the ability to deliver the right information to the right people at the right time to more effectively leverage information and enable more effective business decisions. It's about information on demand to optimize real-time processes. I think of dynamic warehousing as the business intelligence analog of Service Oriented Architecture (SOA).   Dynamic Warehousing includes four key abilities:

1. Support for real-time access to aggregated, cleansed information, which can be delivered in the context of the activities and processes being performed;
2. Embedded analytics that can be leveraged as part of a business process;
3. The ability to incorporate knowledge from unstructured information; and
4. A complete set of integrated capabilities that extend beyond the warehouse to enable Information on Demand
The distinction between data warehousing and online transaction processing is blurring. Data warehousing and analytic applications are accessing operational or near-real-time data. Transactions have become more complex to provide better interaction and productivity for people. Dynamic warehousing has capabilities and strengths on all IBM platforms. The traditional mainframe strengths for consistency with operational data, high security, and continuous availability match well with dynamic warehousing.

Dynamic Warehousing with System z
Mission-critical analysis of operational data

Rapid and secure user-access to data analysis
- *Interactive executive dashboards & information portals*

Improved query and reporting optimizations
- *Parallel queries & SQL Procedures may run on zIIP*
- *Improved SQL & optimization*

Up to 50% reduction of storage for indexes
- *Index compression added to data compression*

Up to 50% reduction of CPU utilization
- *Across many queries & most utilities*

InfoSphere Warehouse on System z

21

IBM INFORMATION ON DEMAND 2007   Act Right. Now.

DB2 9 provides many improvements for Data Warehouse. 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. Accompanying DB2 9 is a new product, IBM DataQuant, which has the critical mass of features that allow it to be positioned as a viable BI and data analytics offering with support for both graphical reports and interactive visual dashboards. It provides a sophisticated graphics engine, supporting dozens of charts and layouts and over 100 built in functions. DataQuant provides very granular security- limit information on a per user/group, also tailor look and feel for users. It also provides a rich client or web-based development/runtime environment. IBM DataQuant is deally suited to the rapid fulfillment of 'everyday' dashboard and reporting requirements. It is simple to develop and deploy – quick turnaround at low development cost. It also allows IT groups / analysts to quickly respond to custom requirements. Where IBM DataQuant fits: Where there's a need to distribute data using straightforward graphical reports and information dashboards Where quick prototyping and rapid development is more important than complex analytical features In QMF and/or z-based environments where tracking, governing and z-based deployment are valued factors For customers that find competitor solutions too complex and costly
Index compression or also deep compression provide a significant reduction in storage. The index compression relies upon page level compression instead of row-level compression (technique used for tables).   Indexes with 32K page sizes can save up to 8x on disk space with the compression feature turned on.   Improved index compression with minimal overhead resulting in Beta customers reporting 50% or more savings in disk space
The key performance improvements in DB2 9 are reduced cpu time in the utilities, improved LOB performance and scalability, improved optimization for SQL, the zIIP processing for remote native SQL procedures, reduced cpu time for data with varying length and better sequential access.
Significant CPU time reduction in most utilities: 10% - 20% in Copy, Recover Table Space
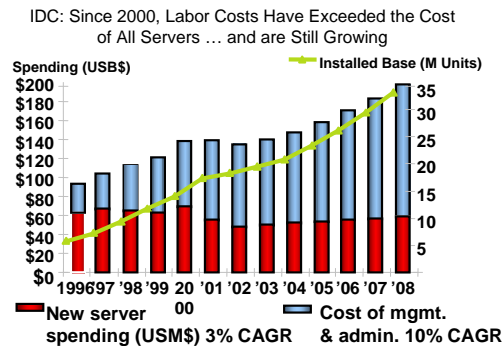   5% - 30% in Load, Reorg, Rebuild   20% - 60% in Check Index
        35% in Load Partition   40% - 50% in Reorg Index
Additional 10% to 15% improvement in virtual storage. Remote Native SQL stored procedures can take advantage of zIIP. Improved data is provided for the optimizer, with improved algorithms. New optimizer techniques and more SQL functions enhance DB2 for z/OS as a DWH platform. New OmniFind text search functions provide efficient communication interactions with DB2 for z/OS. OmniFind text indexes are persisted into DB2 tables for backup/recovery purposes.
1 - Winter Corporation's "2005 Top Ten" awards - http://www.wintercorp.com/index.html

Reduced TCO and Administration Requirements

Improvements in total cost of ownership are very broad for DB2 9. Performance improvements include ways to reduce CPU times, reducing hardware and software costs. Both software and hardware can be reduced by using zIIP processors. DB2 9 has a new category for use of zIIP – in remote native SQL procedures and in expanded ability to run in parallel.

Improved productivity is a larger cost reduction for many customers. Productivity changes help with application life cycles and with database administration. DB2 9 reduces the need to have skills and knowledge specific to System z. Applications from other platforms and other DBMS can be ported to DB2 more easily. These changes also improve DB2 family consistency.

DB2 9 for z/OS Innovation: Cost

- Business Insight

- **Cost Savings through Optimization**
  - **Security and Regulatory Compliance**
  - **Performance improvements**
  - **Synergy with System z**
 - **Query enhancements**

- Business Resiliency

23

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

COST SAVINGS THROUGH OPTIMIZATION

o Increased security and regulatory compliance through implementation of roles, network-trusted contexts, and enhanced auditing

o Performance-boosting innovations such as load and reorg CPU reductions, improved performance for varying length data, and improved logging and insert performance

o Synergy with IBM System z and z/OS in areas that include XML parsing, zIIP, MIDAW channel improvements, encryption, IPv6 and Secure Socket Layer (SSL)

o Query management enhancements to make accessing your data even faster and more accurate with indexing improvements that include index on expression, randomization, and larger index page sizes and optimization improvements that provide better data for the optimizer, improved optimization techniques, and better management with optimization services

IDC: Since 2000, Labor Costs Have Exceeded the Cost of All Servers … and are *Still* Growing
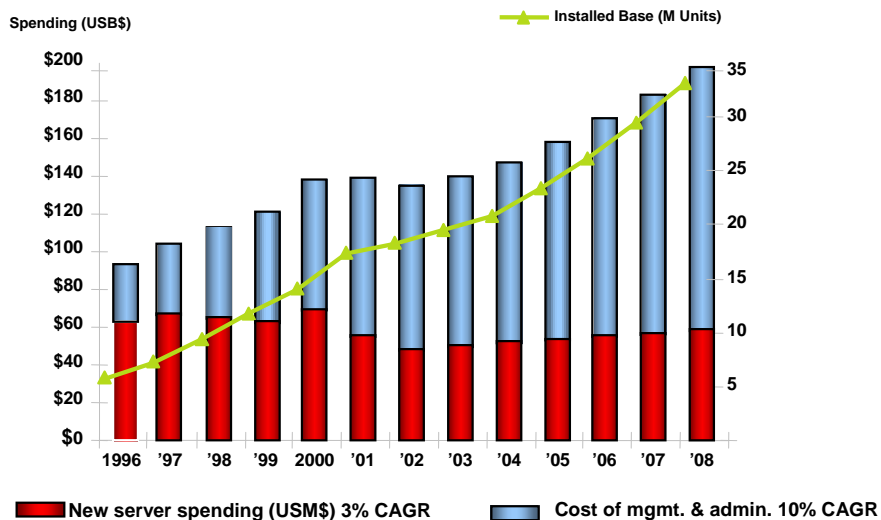
Cost of Complexity

As clients build out their infrastructures….the complexity becomes enormous. They need to optimize IT assets NOW to fuel growth, improve ROI, increase staff productivity and improve quality of service. Often, the cost of acquiring systems is now dwarfed by the cost of managing them. IDC recently came out with a report where they stated that 15 years ago, 75% of IT budgets were spent on new hardware and software; today 75% is spent on just running the infrastructure.

System vendors selling simplicity likely will win in the future, according to Steve Milunovich of Merrill Lynch. Server units are growing at a 20% rate, causing management headaches. Global server support costs could top $140 billion in 2008, nearly 2.5X the estimated hardware sales of $60 billion. Milunovich said he thinks IBM is well positioned with its on demand initiative and $1 billion of R&D spent to automate services. HP has the vision though it lacks some of the pieces and must show it can execute. Sun faces an uphill mindshare battle though its systems understanding shouldn't be underestimated. Dell doesn't have the know-how to succeed but may not need it as it leverages industry R&D and simplifies through low-cost, standard products. Microsoft and Dell do recognize that industry standard servers can be difficult to administer and have joined forces on systems management.

Our Virtualization Engine addresses this problem, enabling clients to lower the cost of their existing infrastructure, reduce the complexity of adding to that infrastructure, and build heterogeneous infrastructure across multiple datacenters that are more responsive to their business needs.

So clearly simplifying your environment is critical to your future success.

What have we done lately?

CPU

+ Memory

+ I/O and disk

+ Software

+ Energy and floor space

+ People

Total cost of ownership

Value for money

   hardware pricing, CPU saving

      specialty engines zIIP, zAAP, …

      compress disk space

   software pricing

     reductions for tiers

     parallel sysplex aggregation

     z990, z9, z10 -10% charge units

     subcapacity, zNALC, Value Unit

25

IBM INFORMATION ON DEMAND 2007   *Act.Right.Now.*

What are we doing to help you hang on to your money? Cost and value is the bottom line. Total cost of ownership or TCO is a frequent conversation with executives and managers. Each of the resources has a cost, and the most costly one for customers is people costs. As the cost of computing continues to decline and the cost for people continues to rise, the future will be dominated by people costs. Many of the important changes in DB2 are to improve people productivity.

Processing costs are often the most visible charges. Some customers charge almost all costs back by processing. Processing costs continue to be reduced for general purpose processing, as CPU power increases. Specialty engines such as zIIP and zAAP reduce the cost of processing power further. The traditional cost has been expressed in thousands of dollars per MIPS, and that figure would be roughly $140 for specialty engines today. The latest change was the z10 zIIP, with the same cost as earlier specialty engines and roughly 50% more processing power.
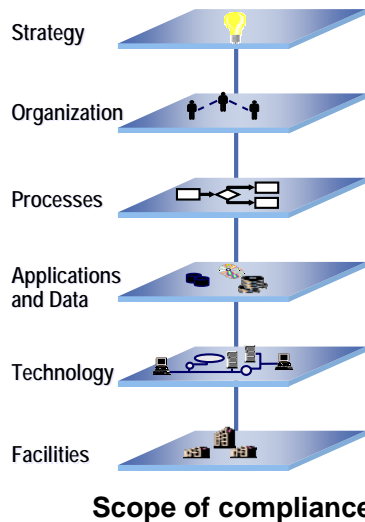
Software costs are higher than hardware costs for most customers, generally in a range of three to ten times higher. This is the reason that specialty engines are so attractive, as the software running on them does not incur software charges. Unit software costs continue to decline in many ways. Software charges have many options, but some of the basic tenets are consistent for most. System z software is charged for the processing capacity of the machine, taking into account multiprocessor effects, rather than using a linear scale. Larger amounts of processing power reduce the unit cost by more than 90%. The z990, z9 and z10 generations have each reduced the charging units by 10% over the prior generation. Changing from z900 to z10 processors would reduce the software charge units by about 28% for the same processing power. Additional options for reducing software charges include parallel sysplex charging, subcapacity pricing, zNALC pricing, and the new Value Unit Edition.

I/O costs include the disk space and the transfer to the processor. Both costs are addressed by the work to use System z effectively. Compression for data uses z/Architecture instructions to be efficient, with a solid track record since 1993 and improved performance in the latest generations. Index compression in DB2 9 adds to the savings.

The key for pricing is the value. DB2 for z/OS delivers a very high quality of service. The service is shared, so that administration time is reduced. The amounts of processing, memory, disk and people can be reduced by effective sharing, delivering the best value for the money.

Regulatory Compliance Characteristics

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

Strategy
Organization
Processes
Applications and Data
Technology
Facilities

Scope of compliance

26

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

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

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

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

Streamlined Compliance in DB2 9 for z/OS

- **DB2 9 built on security of V8 and provides even greater control, flexibility and audit**
- **DB2 9 provides more flexible authorization and control**
  - Database Roles
  - Trusted Context
  - Audit: Improved audit selectivity improves security performance, problem isolation and performance monitoring
- **Extend encryption to tape controller (now) and disk storage (future direction)**

Common Criteria security certification: DB2 for z/OS V8 evaluated at EAL3+ level.

27

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

DB2 for z/OS V8 provides many enhancements for security, DB2 9 builds on that base and improves control, flexibility and ability to audit. DB2 9 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 update, 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 – now for tapes and in the future for disks.

Common Criteria: z/OS V1.7 with the RACF optional feature has achieved EAL4+ for Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP). DB2 for z/OS Version 8 is evaluated under the Common Criteria for CAPP and LSPP with a conformance claim of EAL3+, earning the service mark above.
See:

http://www.ibm.com/systems/z/security/

http://www.ibm.com/systems/z/security/mls.html

http://www.ibm.com/systems/z/security/ccs_certification.html

http://www.ibm.com/security/standards/st_evaluations.shtml

http://www.ibm.com/software/tivoli/governance/action/08162007.html

System z & DB2 9 = Synergy

- **System z9 & z10 Integrated Information Processor (zIIP) Enterprise Class & Business Class**
- **Enhanced Cryptography**
- **Channels (4 Gb, MIDAW & AMP)**
- **Faster Processors**
- **Up to 64 Processors EC**
- **More memory, better value; 64 bit virtual storage**
- **z/Architecture, new machines, new instructions**
- **Parallel Sysplex**

- IPv6
- SSL
- Java
- Decimal float
- Backup & restore
- Security
- Unicode collation
- Compression
- System z Application Assist Processor (zAAP)
- WLM enhanced …

28

IBM INFORMATION ON DEMAND 2007

Act Right. Now.

The System z9 and z10 processor improvements for DB2 include the zIIP and increased use. DB2 9 remote native SQL procedures are enabled for zIIP processing. DB2 9 adds IPv6, SSL and decimal float and BIGINT data types, with enhancements for Parallel Sysplex, backup and restore, added security and encryption, more Unicode with collation, and uses the WLM in new ways. Channel enhancements (MIDAW & AMP) and improved DS8000 performance were included with the System z9 announcements. DB2 uses the latest improvements in hardware and operating system to provide better performance, improved value, more resilience and better function.

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

IBM z10 EC Continues the CMOS Mainframe Heritage

MHz

4000
3500
3000
2500
2000
1500
1000
500
0

DB2 9 for z/OS Performance Topics

Redbooks

4.4 GHz

3.5 GHz

1.7 GHz

1.2 GHz

300 MHz
420 MHz
550 MHz
770 MHz

| 1997 G4 | 1998 G5 | 1999 G6 | 2000 z900 | 2003 z990 | 2005 z9 EC | 2008 z10 EC |

- G4 - 1st full-custom CMOS S/390®
- G5 - IEEE-standard BFP; branch target prediction
- G6 – Copper Technology (Cu BEOL)
- z900 - Full 64-bit z/Architecture®
- z990 - Superscalar CISC pipeline
- z9 EC - System level scaling
- z10 EC – Architectural extensions

29

IBM INFORMATION ON DEMAND 2007   Act. Right. Now.

The design of the IBM System z10™ processor chip is the most extensive redesign in over 10 years, resulting in an increase in frequency from 1.7 GHz (z9 EC) to 4.4 GHz on the z10 EC. The z10 BC processors run at 3.5 GHz.

It is designed for secure data serving, yet also was enhanced to provide improvement enhances for CPU intensive workloads. The result is a platform that continues to improve upon all the mainframe strengths customers expect, yet opens a wider aperture of new applications that can all take advantage of System z10s extreme virtualization capabilities, and lowest TCO versus distributed platforms.

See section 4.3.1 z10 performance in the latest updates of DB2 9 for z/OS Performance Topics, SG24-7473 for additional detail.

z10 Unprecedented performance & capacity
Transforming the economics of the data center

- Lower software license costs
- Lower labor costs
- Lower energy and facilities costs

- **Enterprise Quad Core technology – 4.4 GHz**
- **Up to 2x performance increase in CPU intensive tasks**
- **Enhanced capacity 64-way model**
- **Up to 1.5 TB Memory, 3x more**
- **6 GBps InfiniBand® for faster I/O and Coupling Links**

70% more usable capacity

z9 EC

z990

z900

z10 EC

More Engines

50% more performance
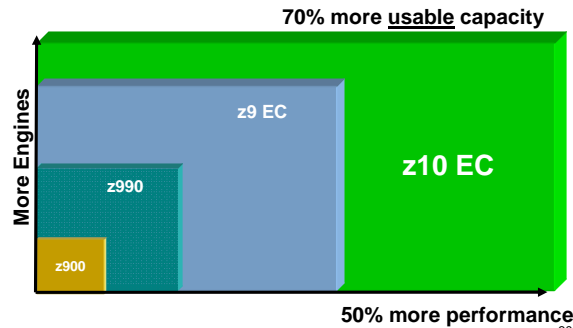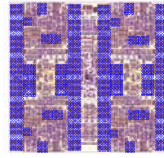
IBM INFORMATION ON DEMAND 2007

Act.Right.Now.

*IBM's next-generation, 64-processor mainframe, which uses Quad-Core technology, is built from the start to be shared, offering greater performance over virtualized x86 servers to support hundreds to hundreds of millions of users.*1*

The z10 also supports a broad range of workloads. In addition to Linux, XML, Java, WebSphere and increased workloads from Service Oriented Architecture implementations, IBM is working with Sun Microsystems and Sine Nomine Associates to pilot the Open Solaris operating system on System z, demonstrating the openness and flexibility of the mainframe.

*From a performance standpoint, the new z10 is designed to be up to 50% faster and offers up to 100% performance improvement for CPU intensive jobs compared to its predecessor, the z9, with up to 70% more capacity.*2*
http://www.ibm.com/systems/z/news/announcement/20080226_annc.html

IBM System z10 Benefits for DB2

- **Faster CPUs, more CPUs, more memory**
  - 50% more n-way performance "on average"
  - 62% more uniprocessor performance
  - 70% more server capacity (56->64 CPUs)
  - Up to 64 CPUs, z/OS 1.9 needed for 64-way in a single lpar
  - Up to 1.5 TB, z/OS 1.8 needed for >256G in a single lpar
- **Infiniband Coupling Facility links**
- **New OSA-Express3, 10 GbE for faster remote apps**
- **HiperDispatch**
- **Hardware Decimal Floating Point facility**
- **1MB page size (DB2 X plans to exploit)**
- **50+ instructions added to improve compiled code efficiency (DB2 X plans to use)**

31

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

The new z10 has faster processors and more processors.   One early measurement showed 40-50% throughput improvement or approximately 30% CPU time improvement with z10 compared to z9.

Larger memory: DB2 users can potentially see higher throughput with more memory used for DB2 buffer pools, EDM pools or SORT pools.

Improved IO: improvements in the catalog and allocation can make the large number of datasets much faster and easier to manage. Disk IO times and constraints can be reduced.

Substantial improvements in XML parsing can result from use of the zIIP and zAAP specialty engines. The z10 zIIP and zAAP engines are much faster at no additional cost. The zIIP processors can be used for XRC processing.

HiperDispatch: Only available on z10 EC. Combination of z/OS software and firmware. Minimum z/OS R1.7 + IBM zIIP Web Deliverable Support for z/OS V1.7 to enable HiperDispatch. z10 EC Driver level 73G. Single HIPERDISPATCH=YES z/OS IEAOPTxx parameter dynamically activates HiperDispatch

z/OS 1.10 Benefits for DB2

- **Extended Addressing Volumes up to 223 GB/volume**
  - VSAM data sets only at first (enabled for DB2 database objects)
  - Needs appropriate DS8000 ucode, DB2 apars on V8 / DB2 9
- **Additional XML use of zAAP and zIIP[1]**
- **z/OS XML performance enhancements**
- **WLM enhancements for lower priority lock holders to boost priority when higher priority waiters (DB2 X plans to exploit)**
- **TCP/IP performance enhancements**
- **HiperSockets Multiple Write Facility for better performance for bigger message sizes (like DRDA)**
- **Hash DSAB for better Open/Close performance**
- **Basic Hyperswap[2] to remove DASD controller as a single point of failure**
- **zIIP assisted Global Mirroring (XRC)[2]**

[1] Planned rollback to z/OS 1.7
[2] Planned rollback to z/OS 1.9

32

IBM INFORMATION ON DEMAND 2007 — Act.Right.Now.

The current maximum of the DASD space is 64K cylinders per volume. z/OS 1.10 will lift the limit and support up to 256K cylinders (higher on the future release) with EAV based on DS8000.

The Future Runs on System z

IBM INFORMATION ON DEMAND 2007

We've talked about the exciting announcements with the new z10 processors, improved disk subsystems and a preview of z/OS 1.10, with an emphasis on the zIIP. DB2 can take advantage of many new improvements immediately, and will use more in the future. A range of software has been enhanced to deliver more value.

DB2 for z/OS Version 8 has been in the field four years now, and we expect that most of you are running it. We've talked over some of the recent changes delivered in the service stream. DB2 9 will be coming soon for many of you, and migration to it is significantly easier than to DB2 V8.

As we listen to customers, we hear many concerns and continue to address them with future work on DB2 X. Total cost of ownership means reducing all kinds of resources: CPU time, disk space, software and hardware cost. Reduced risk and improved productivity for people help with the largest costs.

Helping to drive down the cost of IT
Now even more workloads can benefit from zIIP

- Integrate data across the enterprise, optimize resources and lower the cost of ownership
  - **Centralized data serving**
  - **Network encryption**
  - **Serving XML data**
  - **Remote mirror**
  - **Use by ISVs**

  Available 9/07
  new!
  new!

- zIIPs offer economics to help you
  - **PLUS** zIIP price same for z10 EC as z9 EC

**IBM System z10 Integrated Information Processor and
IBM System z9 Integrated Information Processor**

IBM INFORMATION ON DEMAND 2007   *Act.Right.Now.*

34

zIIP can help to integrate data across the enterprise by optimizing resources and lowering the cost of ownership for eligible data and transaction processing workloads

> **Centralized data serving** – First to exploit zIIP were workloads such as BI, ERP, and CRM applications running on distributed servers with remote connectivity to DB2® V8
>
> **Network encryption** – zIIP becomes an IPSec encryption engine helpful in creating highly secure connections in an enterprise
>
> **Serving XML data** – zIIP is enabled for XML parsing, first to exploit this is inserting and saving DB2 9 XML data over DRDA® . See this paper for more detail.
> DB2 9 and z/OS System Services Synergy Update,
> http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101227
>
> **Remote mirror** – zIIP becomes a data mirroring engine with zIIP assisted z/OS Global Mirror function (zGM, formerly XRC) helpful in reducing server utilization at recovery site (with z/OS V1.8 and above)
>
> Exploiting of zIIPs by ISVs

zIIPs offer economics to help you

> **PLUS** zIIP price is same for z10 EC as z9 EC and we offer no charge MES upgrades when moving to new technology

DB2 for z/OS & IBM zIIP value

Portions of DB2 V8 and DB2 9 (blue) workloads may benefit from zIIP*:

ERP, CRM, Business Intelligence or other enterprise applications via DRDA over a TCP/IP connection
- **DB2 9 for z/OS Remote native SQL procedures**
- **DB2 9 XML parsing**

Specialty Engine

Data warehousing applications & large parallel SQL queries DB2 9 higher percentage of parallel queries means more eligible for zIIP

DB2 Utilities LOAD, REORG & REBUILD maintain indexes

35

IBM INFORMATION ON DEMAND 2007   ActRight.Now.

DB2 9 uses zIIP in two new ways, remote native SQL procedures and increased use of parallelism. See IDUG Europe 2007 presentation F06 by Terry Purcell, "Tuning your SQL to get the most out of zIIPs". This session can be obtained from IDUG online Technical Library by searching for code EU07F06.

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Not all of this work will be run on zIIP. z/OS will direct the work between the general processor and the zIIP. The zIIP is designed so a software program can work with z/OS to dispatch workloads to the zIIP with no anticipated changes to the application – only changes in z/OS and DB2. IBM DB2 for z/OS version 8 was the first IBM software able to take advantage of the zIIP. Initially, the following workloads can benefit: SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP), CRM (Siebel), or business intelligence and are expected to provide the primary benefit to customers. Stored procedures and UDFs run under TCBs, so they are not generally eligible, except for the call, commit and result set processing. DB2 9 remote native SQL Procedure Language is eligible for zIIP processing. BI application query processing utilizing DB2 parallel query capabilities; and functions of specified DB2 utilities that perform index maintenance. For more, see http://www.ibm.com/systems/z/ziip/

The DB2 9 and z/OS System Services Synergy Update paper discusses recent XML benchmark measurements and performance information.
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101227
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101387

DB2 9 for z/OS remote native SQL procedures are described in this paper, showing scalability up to 3193 transactions per second for SQL procedures and redirect to zIIP of over 40%
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD104524

* zIIP allows a program working with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to zIIP. Above types of DB2 work are those running in enclave SRBs, of which portions can be sent to zIIP.

Query Enhancements

- SQL enhancements: INTERSECT, EXCEPT, cultural sort, caseless comparisons, FETCH FIRST in fullselect, OLAP specifications: RANK, ROW_NUMBER, …
- pureXML integration and text improvements
- Index improvements
  - Index on expression   Larger index pages
  - Index compression   Improved page split
- Improved Optimization statistics: Histogram
- Optimization techniques   increased parallel, more zIIP
  - Cross query block optimization
  - Generalize sparse index & in-memory data cache method
  - Dynamic Index ANDing for Star Schema
- Analysis: instrumentation & Optimization Support Center

36

IBM INFORMATION ON DEMAND 2007   Act Right. Now.

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

More queries can be expressed in SQL with new SQL enhancements. The set operators INTERSECT and EXCEPT clauses make SQL easier to write. OLAP extensions for RANK, DENSE_RANK and ROW_NUMBER add new capabilities. Other SQL statements improve consistency with the DBMS industry. DB2 9 continues the progress in SQL, with many new functions, statements and clauses. The biggest changes are in XML on a prior slide. New SQL data manipulation statements are MERGE and TRUNCATE. New data types with DECIMAL FLOAT, BIGINT, BINARY and VARBINARY. Improvements in LOBs provide new function, more consistent handling and improved performance. Security is improved with ROLEs and network trusted context. Data definition consistency and usability are improved. DB2 9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.

Usability

- Application Programming
- Backup and restore objects, use tapes
- Optimization Support Center
- INDEX page split
- Buffer pool management by WLM

37

IBM INFORMATION ON DEMAND 2007   Act Right.Now.

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

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

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

WLM management is extended with the ability to adjust buffer pool sizes.
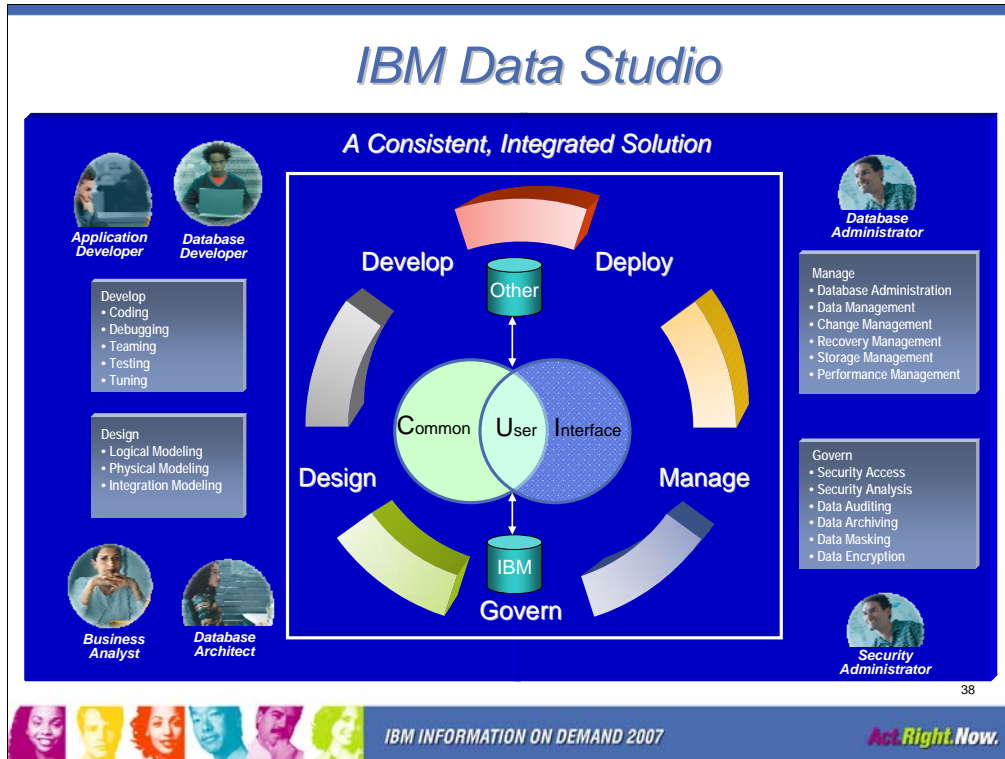
IBM Data Studio is for designers, developers, and administrators who work across the entire data life cycle and across the full set of IBM relational DBMS: DB2 for LUW, DB2 for z/OS, DB2 for i5/OS and IDS. Almost every customer will need the tools included with DB2, and many will need the additional ones..

DB2 is changing to deliver the needed interfaces, usually new stored procedures.

See the IBM Data Studio web page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast.

http://ibm.com/software/data/studio

DB2 9 for z/OS Innovation: Resiliency

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

  ▪**Availability: Database on Demand**

  ▪**Scalability**

  ▪**Usability**
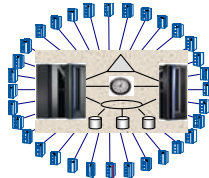
IBM INFORMATION ON DEMAND 2007   Act Right. Now.

BUSINESS RESILIENCY INNOVATIONS

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

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

o Further improvements in DB2 9 help manage business-critical information with fewer planned or unplanned outages with features like online reorganization with no BUILD2 phase, rename column and index, alter index to change page size, and ability to modify early code without requiring an IPL.

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

Online Schema Evolution → Database Definition On Demand

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

IBM INFORMATION ON DEMAND 2007   Act Right. Now.

40

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

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

Scalability

- Insert performance APPEND INDEX LOG
  - INDEX on expression, 8K, 16K, 32K pages
  - Randomized index key
  - Log Latch contention & spin relief, archiving
  - Not logged table space
- Partitioned table with segmented space
- Memory improvements 64 bit address space

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

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

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

Migration to DB2 9 for z/OS is easier

- Migration process enhancements: ENFM speed, CM*
- Much less performance regression:
 - Earlier improvements
  - Plan stability & tools for avoiding access path issues
- CCSIDs and old product issues resolved in V8
- Simpler virtual storage considerations
- Less impact from incompatible changes
- Earlier deliveries from vendors

IBM INFORMATION ON DEMAND 2007

Act.Right.Now.

42
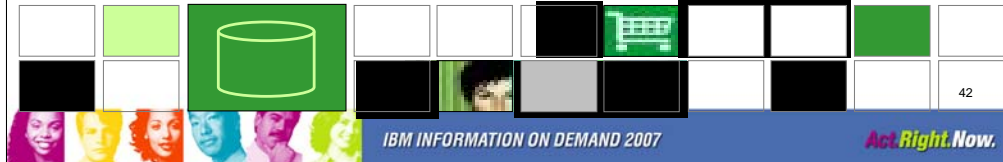
DB2 9 migration has been easier for many customers and tends to be faster than DB2 V8.

Migration process enhancements

ENFM is shorter (10 minutes versus an hour), as only two table spaces need to be reorganized in this phase.

Drop back to CM* if needed

Much less performance regression

More consumable performance improvements

Utility CPU improvements in CM

Plan stability

CCSIDs and old product issues resolved

Simpler virtual storage considerations

Less impact from incompatible changes

42

Why Migrate to DB2 9 for z/OS?

- Business needs to save money
  - Reduce CPU time & disk space
  - Improve business agility
  - Service Oriented Architecture
- Application developers need
  - PureXML for a powerful SQL and XML interface to XML data
  - Powerful new SQL enhancements
  - Portability with SQL and data definition compatibility

- Database Administrators need
  - Improve availability and performance
  - More flexible security and easier regulatory compliance
  - Better web application & data warehouse function and performance
  - LOB function, performance, usability

43

IBM INFORMATION ON DEMAND 2007

Act Right.Now.
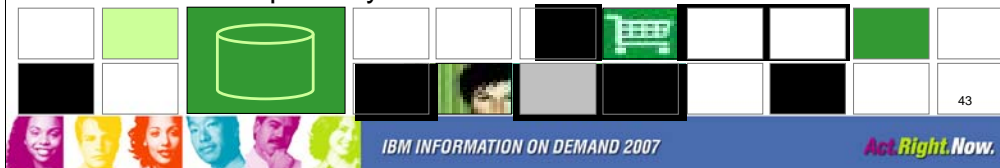
DB2 9 has a lot for everyone. Here are just a few of the strategic highlights.

The business needs include CPU cycle reductions that deliver in most utilities, disk space reduction from index compression, improved query optimization, improved business agility via faster implementation cycles, and new pureXML™ that builds a strong foundation for SOA and XML initiatives. Kevin Campbell, an Application Architect at Univar USA said it better than I can, "This is not a bolt-on or band-aid approach, DB2 9 for z/OS is XML without compromise."

Database Administrators (DBAs) need improved database availability and performance including LOBs, reorganization, backup and recovery, and partitioning enhancements. DBAs also get more flexible trusted network context and role-based security to help with regulatory compliance. A wide range of enhancements improve ERP application and data warehouse functionality and performance. Large object (LOB) function is added with file reference variables and REORG, while performance is improved.

Application developers are most excited by PureXML, which adds a powerful SQL and XML interface to access XML data stored in a native format. Application developers need powerful new SQL enhancements including MERGE and TRUNCATE statements, INTERSECT and EXCEPT set operations, and spatial support for geographical data. Text handling is improved with the XML changes, many new built-in functions, and an upcoming text server. Improved SQL and data definition compatibility with other DB2 platforms makes porting much easier.

DB2 for z/OS

Delivering Customer Value / Future Version Philosophy

OnGoing themes:
Performance Scalability
Reliability Availability Serviceability
Security Productivity
Application Development
SQL XML SOA

IBM INFORMATION ON DEMAND 2007 — Act.Right.Now.
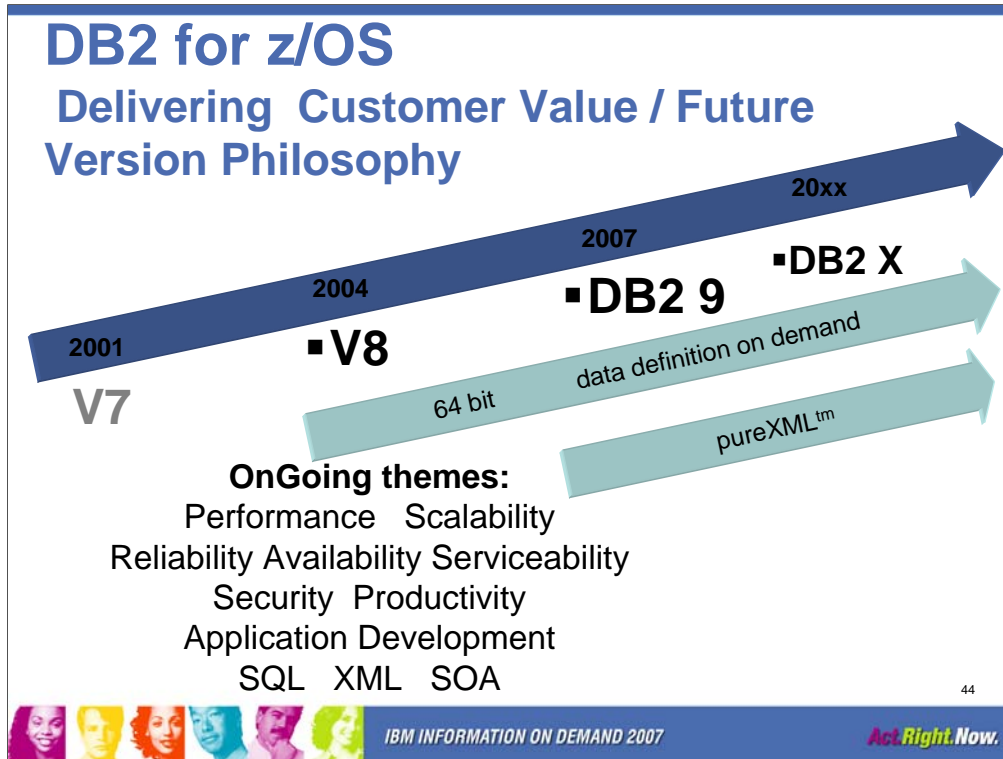
DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later. DB2 9 became generally available in March 2007, three more years. We expect the next version will be roughly 3 years from DB2 9 GA to GA of DB2 10 or DB2 X or whatever the name becomes.

The themes for future versions will continue to focus on core platform strengths of performance, scalability, reliability, stability, availability, resilience, and security. PureXML and Schema evolution or data definition on demand will be ongoing for a long time. In contrast, most of the 64 bit evolution should be completed in DB2 X.

The key interfaces for customers and vendors expand for both XML and for SQL. Information is a key leg of the SOA platform, and DB2 for z/OS provides many advantages for data management in SOA.

Standards, interoperability, portability and security along with secure access using the latest technologies are key touch points. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity.

Get updated books   **latest February 2009**
http://publib.boulder.ibm.com/infocenter/imzic/

- ✓ Program Directories
- ✓ Installation Guide
- ✓ Redbooks
- ❑ Data Sharing: Planning and Administration
- ❑ Application Programming Guide & Reference for JAVA™
- ❑ RACF Access Control Module Guide
- ❑ ODBC, Spatial, …
- ✓ Messages
- ✓ Codes
- ✓ Administration Guide
- ✓ Performance Guide
- ✓ Utility Guide and Reference
- ✓ Internationalization Guide (Unicode)

45

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

For installation, you need many books. Some are optional, for example the data sharing book is not needed if you don't use data sharing. If you use data sharing, then part of the install process is in that book. Part of the Java install process is in the Java book. Part of the ODBC install process is in the ODBC book. If you use RACF access control, then you need the RACF Access Control Module Guide. You can get most of the books from the DB2 Library web page. The books were updated in December 2007 and February, March, June, and August 2008, with some coming later, so get the latest ones. Some of the Redbooks may be helpful. You may need books from the z/OS Library as well.

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011658

http://www.ibm.com/systems/z/os/zos/bkserv/r9pdf/

Be sure to use the latest information to save time and problems. Some of the IBM Redbooks publications have always been updated and added lately (next page).

45

DB2 9 in **IBM** **Red**books Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473
3. DB2 9 Stored Procedures SG24-7604
4. Index Compression with DB2 9 for z/OS redp4345
5. SQL Reference for Cross-Platform Development
6. Enterprise Database Warehouse, SG24-7637
7. 50 TB Data Warehouse on System z, SG24-7674
8. DB2 9 Optimization Service Center SG24-7421
9. LOBs with DB2 for z/OS SG24-7270
10. Deploying SOA Solutions SG24-7663
11. Enhancing SAP - DB2 9 SG24-7239
12. SAP Application on Linux z SG24-6847
13. Best practices SAP BI - DB2 9 SG24-6489-01
14. Data Sharing in a Nutshell, SG24-7322
15. Securing DB2 & MLS z/OS SG24-6480-01
16. Data Sharing: Distributed Load Balancing & Fault Tolerant Configuration redp4449
17. Considerations on Small & Large Packages redp4424
18. Backup and Recovery Considerations redp4452
19. Powering SOA with IBM Data Servers SG24-7259
20. Packages Revisited, SG24-7688
21. Data Studio V2.1 Web Services redp4510

SQL Reference for Cross-Platform Development

LOBs with DB2 for z/OS: Stronger and Faster

Securing DB2 and Implementing MLS on z/OS

DB2 9 for z/OS Technical Overview

Enhancing SAP by using DB2 9 for z/OS

Enterprise Data Warehousing with DB2 9 for z/OS

DB2 9 for z/OS Performance Topics

DB2 9 for z/OS: Packages Revisited

Redbooks

46

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

DB2 library more information   http://www.ibm.com/software/data/db2/zos/library.html

Many IBM Redbooks publications, Redpapers and one cross-platform book on DB2 9 are published, in addition to the standard library, with more in the works. Check for updates.

1. DB2 9 Technical Overview, SG24-7330 http://www.redbooks.ibm.com/abstracts/SG247330.html
2. DB2 9 Performance Topics, SG24-7473, http://www.redbooks.ibm.com/abstracts/SG247473.html
3. DB2 9 Stored Procedures, SG24-7604, http://www.redbooks.ibm.com/abstracts/SG247604.html
4. Index Compression DB2 9, REDP4345, http://www.redbooks.ibm.com/abstracts/redp4345.html
5. Deploying SOA Solutions SG24-7663, http://www.redbooks.ibm.com/abstracts/SG247259.html
6. Cross-Platform Development Version 3, http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
7. Enterprise Data Warehousing, SG24-7637, http://www.redbooks.ibm.com/abstracts/sg247637.html
8. LOBs: Stronger & Faster SG24-7270, http://www.redbooks.ibm.com/abstracts/SG247270.html
9. Securing DB2 & MLS z/OS, SG24-6480-01, http://www.redbooks.ibm.com/abstracts/sg246480.html
10. Enhancing SAP, SG24-7239, http://www.redbooks.ibm.com/abstracts/SG247239.html
11. Best practices SAP BI, SG24-6489-01, http://www.redbooks.ibm.com/abstracts/sg246489.html
12. Optimization Service Center, SG24-7421, http://www.redbooks.ibm.com/abstracts/sg247421.html
13. Data Sharing in a Nutshell, SG24-7322, http://www.redbooks.ibm.com/abstracts/sg247421.html
14. DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration http://www.redbooks.ibm.com/abstracts/redp4449.html
15. DB2 for z/OS: Considerations on Small and Large Packages redp4424 http://www.redbooks.ibm.com/abstracts/redp4424.html
16. DB2 9 for z/OS: Backup and Recovery Considerations redp4452 http://www.redbooks.ibm.com/abstracts/redp4452.html
17. Powering SOA IBM Data Servers, SG24-7259   http://www.redbooks.ibm.com/abstracts/SG247259.html
18. DB2 9 for z/OS Packages Revisited, SG24-7688 http://www.redbooks.ibm.com/abstracts/SG247688.html
19. 50 TB Data Warehouse Benchmark on IBM System z http://www.redbooks.ibm.com/redpieces/abstracts/sg247674.html
20. SAP on DB2 9 for z/OS: Implementing Application Servers on Linux for System z http://www.redbooks.ibm.com/redpieces/abstracts/sg246847.html
21. IBM Data Studio V2.1: Getting Started with Web Services on DB2 for z/OS http://www.redbooks.ibm.com/redpieces/abstracts/redp4510.html
22. Parallel Sysplex Operational Scenarios http://www.redbooks.ibm.com/redpieces/abstracts/sg242079.html
23. Watch for titles on DB2 distributed; serialization & concurrency; utilities

Main DB2 for z/OS Web Page
ibm.com/software/db2zos

- Technical resources
- Events
- Support
- V8, DB2 9
- Developer Domain
- DB2 Magazine
- zIIP
- Tools

This is the main DB2 for z/OS web page. You can get to the other DB2 for z/OS pages from here, so I often call this my home page. This page changes frequently, so look at the highlighted NEW items. Do you want to look in a DB2 book? Click on Library to see books on DB2 and QMF Version 8 (about 40), Version 7, 6 or even 5. V6 and V5 are out of service. You can check the latest changes by looking at the Information Updates or go to the Information Center. From this page, you can look for conferences (Events), specific classes (Education), or services. If you want to see the latest on DB2 9 or DB2 Version 8, click on the DB2 9 or the V8 link. If your primary concern is application development, the Developer Domain is for you. DB2 Magazine covers a broad range of topics about DB2. The latest machines System z9, z990 and z890 are on the System z page. Click DB2 and IMS Tools to see the wide range of help we provide.

Data Studio

http://www.ibm.com/software/data/studio/

Search

Home   Solutions ▾   Services ▾   Products ▾   Support & downloads ▾   My IBM ▾

Welcome Mr. Roger Miller [Not you?] [IBM Sign in]

Software > Information Management >

IBM Data Studio

The comprehensive data management solution to design, develop, deploy, and manage data-driven applications.

IBM Data Studio
Features and benefits
System requirements
Technical resources
Trials and demos

"With IBM Data Studio, we can discover faulty assumptions and get resolution from business sponsors early in the design process. This will help us reduce development costs as it's 50 to 75 percent more expensive to fix a design problem after we implement the changes"

Kevin Campbell, Application Architect, Univar USA

We're here to help

Easy ways to get the answers you need.

📞 Call me now
✉ E-mail us

or call us at
877-426-3774
Priority code:
104CBW60

Related links
· DB2 9
· Informix
· U2
· IMS
· IBM Balanced Warehouse
· DB2 for z/OS

Why IBM Data Studio

IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

Finally! Integrated

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

See this page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast. IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio Developer: An Integrated Development Environment for creating and testing database and pureQuery applications.

IBM Data Studio pureQuery Runtime: A high-performance Java data access platform -- improves security and manageability of Java application connections to databases.

http://www.ibm.com/software/data/studio/

zIIP ibm.com/systems/z/ziip/ II14219 II12836 II10817

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

Servers > Mainframe servers > Hardware >

Mainframe servers
About IBM mainframe
Hardware
Software
Operating systems
Solutions
Customer successes
Literature
FAQs
Support
Education
News and events

zIIP – Overview

zIIP | Getting started

Overview | About | Resources |

Today's corporate data

Data is at the core of every business. But data management can be a big project. It needs to be available globally; it must be kept up in real time, and accessible 24x7. At the same time, it needs to be protected from unauthorized access and to comply with new regulations. The IBM System z™ mainframe family has a strong heritage in data serving, and its architecture is designed for massive data access, whether across the Internet, to storage devices, or to remote backup sites. Scalability, availability, security—these are all core competencies of the mainframe.

Learn More
→ What they are saying about zIIP
→ Data Serving Press release
→ Data Serving Roadmap z990
→ z/OS website
→ DB2 V9 for z/OS preview
→ System z security

We're here to help

IBM Mainframes just got even more cost competitive. Let us show you how with a customized quote.

→ Request a quote

or call us at
1-888-SHOP-IBM

Related links
• Resource Link
• Resources for IBM Business Partners
• Resources for developers
• ShopSystemZ
• Printing solutions
• ISV software support
• IBM Training
• IBM Design Centers

Exciting new features

Earlier this year, IBM extended its mainframes data serving capabilities, delivering a new roadmap for the future of data serving and information on demand, previewing new DB2 function, and introducing a new specialty engine directed toward data serving workloads.

The new specialty engine, the IBM System z9 Integrated Information Processor (IBM zIIP), is now available on the System z9 Enterprise Class (EC) and System z9 Business Class (BC) servers.

49

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

See zIIP information on the web. This web site has the most current information, & pointers to more resources. The zIIP is for customers who are concerned about costs for growth. One big cost reduction is hardware cost, which is much less than a standard processor. The biggest cost reductions are in software, as IBM does not charge for software running on the specialty processors. The zIIP fits some customers very well, but will not apply for all. As a specialty processor, not all work can use the zIIP, which will only process work running under an enclave SRB. Most applications cannot run in SRB mode. The specifics of the software charging need to be considered. Customers must be current on hardware (System z9), current on software (z/OS 1.6 or later, DB2 V8 or later) and have a work load peak using the types of work eligible for zIIP:

Remote SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business intelligence running on other platforms. Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. Stored procedures and UDFs run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing. DB2 9 remote native SQL Procedure Language is eligible for zIIP .

Parallel queries: If the work comes in remotely over DRDA using TCP/IP, then the initial work is eligible as remote work. After the initial time, the parallel processing threads are eligible and can process on the zIIP.

DB2 utility index processing: Functions of the LOAD, REORG and REBUILD DB2 utilities that perform index maintenance are eligible for zIIP. This is not a common peak capacity constraint, but could be useful in reducing CPU charges.

The best way to estimate the eligible work is to apply the needed z/OS and DB2 service, to run your work load and to take measurements. Use DB2 accounting with any product which can provide DB2 accounting reports, such as Omegamon.

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

Technotes (FAQs)

Presentations

White papers

Redbooks
✓new books

APARs, CST

Do you need Answers to Frequently Asked Questions? Product Information? Technical Presentations? Redbooks? White Papers? Support for all of these is on this web page. Click Support on DB2 web pages. From this page, you access roughly 800 Technotes or answers to Frequently Asked Questions, technical presentations from recent conferences, a dozen white papers, and over 50 Redbooks. For more presentations, see the Events page. The Consolidated Service Test (CST) documents IBM recommended service level for a range of z/OS products.

To search for the presentation you want, put a word or two into the Additional search terms or Order by date – newest first.

There are new and updated redbooks on Powering SOA with IBM Data Servers, LOBs with DB2 for z/OS: Stronger and Faster, V8 Performance Topics, Data Sharing in a Nutshell, Data Integrity, Stored Procedures: Through the Call and Beyond, Design Guidelines for High Performance and Availability, Disaster Recovery, and Security Improvements.

ibm.com/software/db2zos/support.html

DB2 9 Technical Resources page
ibm.com/support/docview.wss?rs=64&uid=swg27011656

- **Download 30 PDF books**
- **Access html books**
- **DB2 new ✓December 2008**
- **Program Directories**
- **Information Center**

DB2 for z/OS - Technical Resources

Support & downloads
DB2 for z/OS
Features and benefits
System requirements
Library
Technical resources
Success stories
News
How to buy
Events
Training and certification
Services
Support

Related software
· WebSphere
· CICS
· IMS
· DB2 Connect
· z/OS
· Linux on IBM System z

Related hardware
· IBM TotalStorage
· IBM System z

Related solutions
· DB2 and IMS Tools
· Java on OS390 and z/OS
· DB2 for Linux

Related services
· IBM Consolidated Service Test

Product documentation

Abstract
Here is your complete DB2 for z/OS information resource! Here you can find DB2 for z/OS technical information, information about related tools, and valuable technical tips.

Content
To support recent enhancements to the DB2 for z/OS Library page, the DB2 for z/OS technical information has a new address on the Web. **Please update your bookmarks!**

Information center topics
➔ Information Management Software for z/OS Solutions Information Center
Search topics across all information for DB2 for z/OS, IMS, and DB2 and IMS Tools, including QMF. The Information Management Software for z/OS Solutions Information Center is the primary source for the information about DB2 for z/OS.

Search the information center topics
[   ] [Search]

As part of IBM's ongoing effort to improve your information experience, we are optimizing the information delivered with the DB2 for z/OS product for viewing in the information center. Learn more.

Technical information by version
➔ DB2 Version 9.1 for z/OS
➔ DB2 for z/OS Version 8

51

IBM INFORMATION ON DEMAND 2007   Act Right.Now.

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

Watch this web page and see the dates next to each book for the latest updates.

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656

Installable option
http://www.ibm.com/support/docview.wss?rs=865&uid=pub1sk5t737700
http://publib.boulder.ibm.com/infocenter/imzic/

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

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

The latest change in this area is an installable Information Center, so that you can use the facility even when the internet is not accessible.
http://www.ibm.com/support/docview.wss?rs=865&uid=pub1sk5t737700

The online version is

http://publib.boulder.ibm.com/infocenter/imzic/

Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR
- ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

53

IBM INFORMATION ON DEMAND 2007   Act.Right.Now.

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.

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/400, Netfinity, RISC, RISC SYSTEM/6000, System i, System p, System x, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, System z,

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