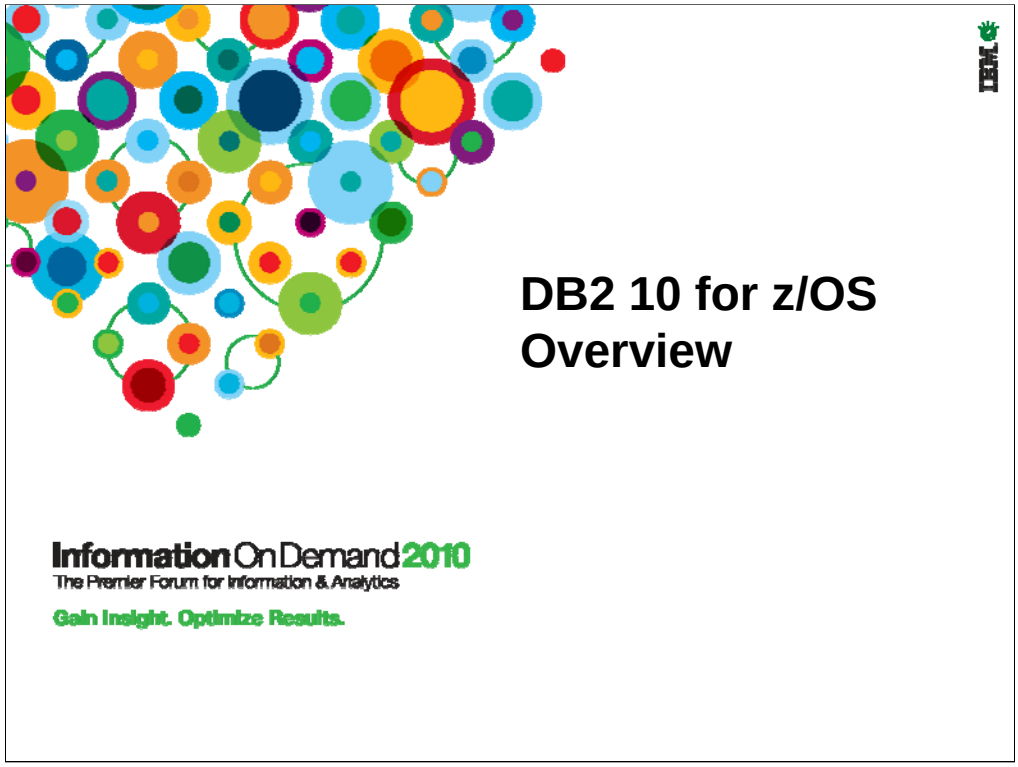

DB2 10 for z/OS – More for less

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch in 22 years. We expect most customers to reduce CPU times between 5% and 10% initially, with opportunity for more. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improves availability. SQL and pureXML improvements extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity.

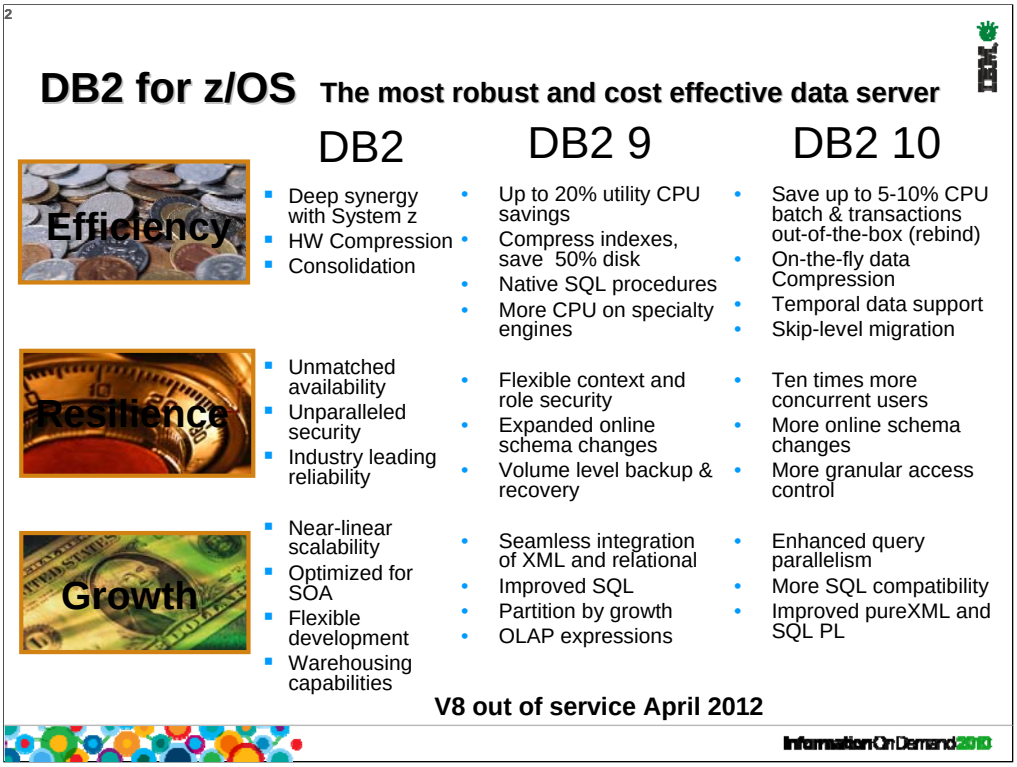

Customers have come to know DB2 as the most robust and cost effective data server. With every version of DB2, we are focused on the needs of our customers to operate efficiently, to be up and running 24x7, and to grow with their business. With DB2 9, customers get CPU and disk savings as well as a boost in application productivity with the new pureXML technology. For the next DB2 version, a lot of customers are getting excited. We are putting a lot of focus on out-of-the-box performance improvements and productivity improvements such as online schema, temporal data support, and fine-grain security controls. DB2 continues to be the choice for mission critical business data and we continue to make it easier for customers to keep data on the platform.

DB2 9: 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.

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 10: DB2 10 for z/OS provides the best reduction in CPU for transactions, queries, and batch for over 20 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% as soon as DB2 10 is out of the box, after rebinding static SQL applications. Applications which can take advantage of additional benefits, such as hash access, index include columns, inline large objects, parallel index updates, faster single row retrievals, work file in-memory, index list prefetch, 64 bit memory enhancements, use of the System z10 1 megabyte page size, buffer pools in memory, access path enhancements, member clustering for universal table spaces, efficient caching of dynamic SQL statements with literals, improved large object streaming, and SQL procedure language performance can have additional CPU and memory reductions. As always with performance, individual customer experiences will vary, and individual workloads will vary more.

Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improve availability, by using an ALTER where the only prior option was DROP and recreate. Improved concurrency to DB2 catalog access and utilities extends the scaling. Security is enhanced with better granularity for administrative privileges, masking for data, and new audit capabilities.

SQL, pureXML, and web services improvements extend usability and application portability to the System z, z/OS and DB2 for z/OS platform. Temporal or versioned data improves productivity for applications in a wide range of industries. Applications ranging from SAP to warehousing see benefits from every category and item.

The net result is productivity improvements in DB2 10 for application developers, for database administrators, and for systems administrators that are very important as data grows in scale and complexity.

DB2 Deep Synergy With System z

Key integration points include:

- Data sharing (availability and scale out)
- zIIP and other specialty engines
- Unicode conversion
- Encrypted communication & data
- Hardware data compression & encryption
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF)
- z/OS RRS integrated commit coordinator
- System z10 1 MB page size, decimal float
- Solid state disks
- zEnterprise z196, zBX, z10, …

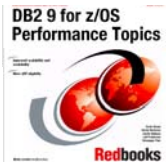

Data sharing is a prime example of deep synergy with System z. DB2 worked with the System z design team for nearly 10 years to produce a robust platform for horizontal scaling. The evolution has continued for 15 more years now after delivery.

Hardware data compression and encryption provides improved costs, easier management and robust resilience for the platform. Cross-memory and protection keys work with APF authorization and RACF for the underlying system integrity.

Specialty engines can reduce costs very substantially, reducing both hardware and software costs.

The z/OS workload manager (WLM) has changed in almost every release to improve work flow with DB2. DB2 has a dispatcher, the z/OS WLM.

Sorting, decimal arithmetic, decimal float, encryption, and Unicode conversions are examples of unique instructions in z/Architecture that DB2 uses.

DB2 has unique ways to use the z10 and zEnterprise to deliver additional value.

System zEnterprise Benefits for DB2

Taking System z synergy to the next level

- Faster CPUs, more CPUs, more memory → better DB2 performance, scalability
- Compression hardware expected to increase DB2 data compression performance
- Cache optimization, 192M L4 Cache expected to benefit DB2 work
- Hybrid architecture query performance acceleration with IBM Smart Analytics Optimizer
- Excellent synergy with DB2 10 → significant CPU reduction and scalability increase
  - CPU reductions, increased zIIP usage
  - Remove key single system scaling inhibitors: virtual storage, latching, catalog, utilities, …
  - Translation Lookaside Buffer changes expected to improve performance for 1MB page sizes
  - Buffer pool management

The zEnterprise z196 and DB2 10 take synergy to the next level. See Chris Crone's presentation for more detail. Faster CPUs, more CPUs, and more memory means better DB2 performance and scalability.

Compression hardware improvements are expected to increase DB2 data compression performance.

192M L4 Cache is expected to benefit DB2 workloads, as DB2 uses the memory. DB2 can take an advantage of cache optimization on zEnterprise.

Hybrid architecture delivers new opportunities for DB2 query performance acceleration with IBM Smart Analytics Optimizer.

Excellent synergy with DB2 10 removes many single system scaling inhibitors – virtual storage, latching, catalog concurrency, and utility concurrency.

The Translation Lookaside Buffer Changes are expected to improve DB2 10 performance for 1MB page sizes. Buffer pool improvements for large buffers will provide additional help.

Combined with DB2 10 improvements CPU reduction, buffer pool management, relief for virtual storage constraint and latch contention, DB2 applications can enjoy significant cost reduction and scalability improvement on zEnterprise.

5

Top 10 in DB2 10 for z/OS

1. CPU reductions for transactions, queries, & batch
2. Ten times more users by avoiding memory constraints
3. More concurrency for catalog, utilities, and SQL
4. More online changes for data definition, utilities and subsystems
5. Improved security with more granularity
6. Temporal or versioned data
7. SQL enhancements improve portability
8. pureXML performance and usability enhancements
9. Hash, index include columns, skip migration, …
   Pick your favorite!
10. Productivity improved for database & systems administrators, and application programmers

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for over 20 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% as soon as DB2 10 is out of the box. Applications which can take advantage of additional benefits, such as hash access, index include columns, inline large objects, parallel index updates, faster single row retrievals, work file in-memory, index list prefetch, 64 bit memory enhancements, use of the System z10 1 megabyte page size, buffer pools in memory, access path enhancements, member clustering for universal table spaces, efficient caching of dynamic SQL statements with literals, improved large object streaming, and SQL procedure language performance can have additional CPU and memory reductions. As always with performance, individual customer experiences will vary, and individual workloads will vary more.

Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improve availability, by using an ALTER where the only prior option was DROP and recreate. Improved concurrency to DB2 catalog access and utilities extends the scaling. Security is enhanced with better granularity for administrative privileges, masking for data, and new audit capabilities.

SQL, pureXML, and web services improvements extend usability and application portability to the System z, z/OS and DB2 for z/OS platform. Temporal or versioned data improves productivity for applications in a wide range of industries.

The net result is productivity improvements in DB2 10 for application developers, for database administrators, and for systems administrators that are very important as data grows in scale and complexity.

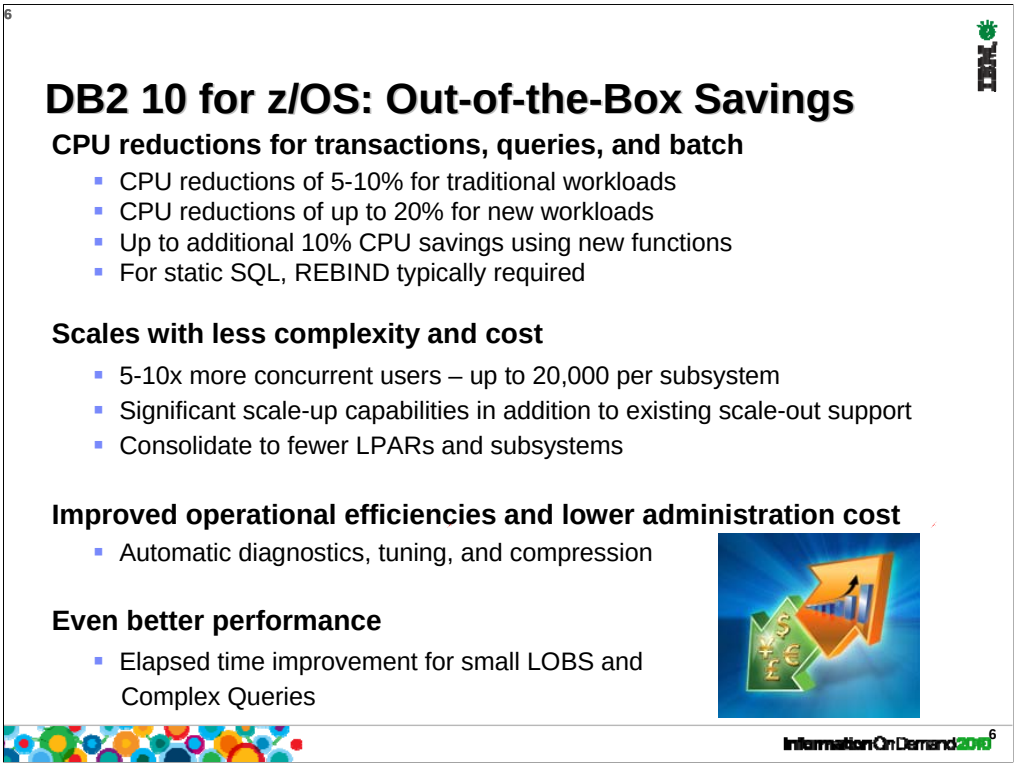

Improved operational efficiency for out-of-the-box savings Version 10 delivers great value by reducing CPU usage. Compared to previous releases of DB2 for z/OS, most customers can achieve out-of-the-box CPU savings of five to ten percent for traditional workloads and up to 20 percent for some workloads. DB2 reduces CPU usage by optimizing processor times and memory access, leveraging the latest processor improvements, larger amounts of memory, solid-state drives, and z/OS enhancements. Improved scalability and constraint relief can add to the savings. Productivity improvements for database and systems administrators can drive even more savings.

In Version 10, performance improvements focus on reducing CPU processing time without causing significant administration or application changes. Most performance improvements are implemented by simply migrating to Version 10 and rebinding. You gain significant performance improvements from distributed data facility (DDF) optimization, buffer pool enhancements, parallelism enhancements, and more.

Most customers should see 5% - 10% CPU reduction out of the box after rebinding. Some workloads and customer situations can reduce CPU time more. While versions 3, 4, 5, 6, and 7 generally increased CPU times by a small amount, less than 5%, version 8 increased CPU time by 5% to 10% for most customers. DB2 9 often reduced CPU a little or increased very little (less than 2%). New function, improved scalability, and faster hardware compensated for the increases in CPU time. Using the new function could change the increases into reductions, particularly with DB2 V8 and multi-row fetch.

Early DB2 10 performance benchmarking and customer experience has shown that most customers can expect to get 5% to 10% CPU reduction after rebinding. Some customers will get more and some less. Some situations can reduce CPU time more than that. Customers who have scalability issues, such as virtual storage constraints or latching can see higher improvements. Opportunities for tuning can take advantage of memory improvements. High volume, short-running distributed transactions can take advantage of CPU reductions, using release deallocate. Concurrent sequential insert can be reduced from 5% - 40%. Queries can be improved as much as 20% without access path change, and more for better access paths. A workload with native SQL procedures has shown up to 20% CPU reduction. For DB2 utilities, customers moving from DB2 9 should expect a small (0% to 7%) reduction in CPU times varying by utility, while customers moving from DB2 V8 will see larger CPU reductions in the range of 20%.

Productivity improvements: Improvements in SQL and XML improve productivity for those who develop new applications and for those who are porting from other platforms. Automating, reducing, or eliminating tasks, and avoiding manual invocation improves productivity and can help avoid problems. Resiliency improvements for virtual storage and availability increase productivity. DB2 10 improvements make the install, migration, and service processes faster and more reliable, including the ability to skip from V8 to DB2 10.

Innovations in Version 10 drive new value in resiliency through scalability improvements and fewer outages, whether those outages are planned or unplanned. Scalability delivers the ability to handle five to ten times more concurrent users in a single DB2 subsystem than in previous releases of DB2 for z/OS (as many as 20,000 concurrent threads). Improved availability is supported by schema evolution, or data definition on demand, and manageability enhancements for query performance

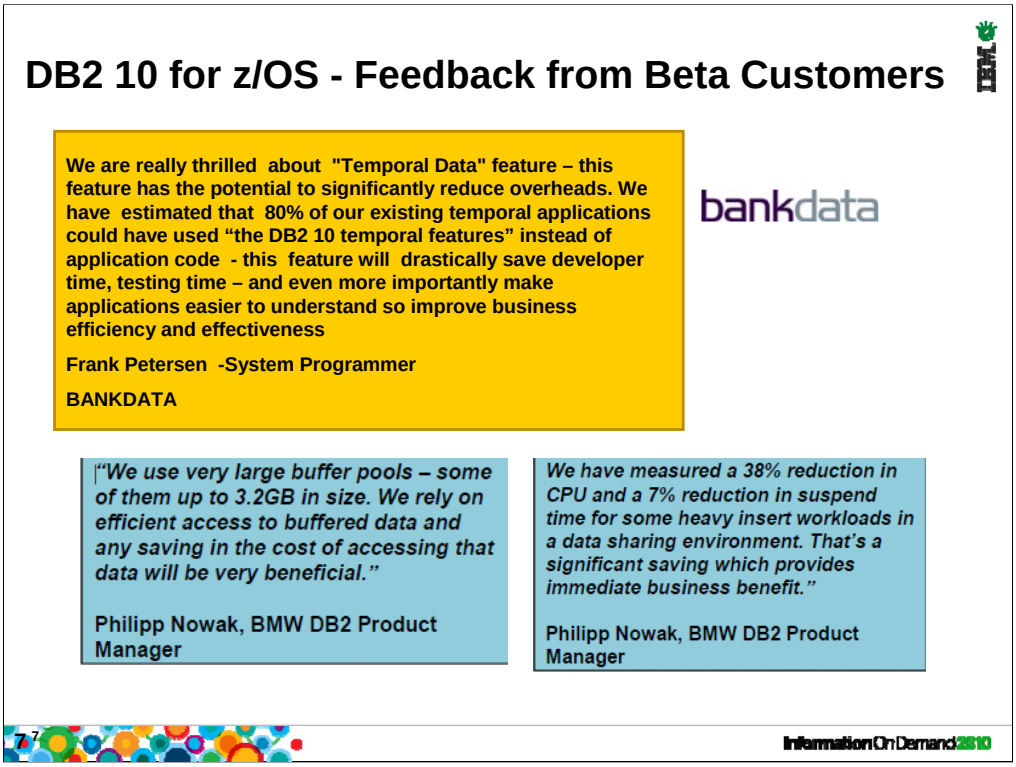

The customer quotes speak for themselves. See more on the web:

http://www.ibm.com/software/data/db2/zos/db2-10/

http://www.ibm.com/software/data/db2/zos/db2-10/testimonials.html

The customer quotes speak for themselves. See more on the web:

http://www.ibm.com/software/data/db2/zos/db2-10/

http://www.ibm.com/software/data/db2/zos/db2-10/testimonials.html

Sample of Beta Customers' Performance Feedback

Workload	Results
Customer1: Distributed Concurrent Insert	50% DB2 elapsed time reduction 15% chargeable CPU reduction after enabling high perf DBAT
Customer2: CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, Another 4% reduction with 1MB page usage
Customer3: CICS online transactions	Approx 5% CPU reduction
Customer4: Data sharing heavy concurrent insert	38% CPU reduction
Customer5: Queries	Average CPU reduction 28% from V8 to DB2 10 NFM
Customer6: Batch	Overall 28% CPU reduction after rebind packages
Customer7: DDF OLTP	40% CPU reduction for JDBC stored procedures workload, 15% CPU reduction for securities trading app

Here are some customer examples of performance results for a broad workload.

Beta Customers' Feedback on Selected New Functions

Workload	Results
Multi row insert (data sharing)	33% CPU reduction from DB2 9, 4x improvement from V8 due to LRSN spin reduction
Parallel Index Update	30-40% Elapsed time improvement with class 2 CPU time reduction
Inline LOB	SELECT LOB shows 80% CPU reduction
Include Index	17% CPU reduction in insert after using INCLUDE INDEX
Hash Access	20-30% CPU reduction in random access 16% CPU reduction comparing Hash Access and Index-data access. 5% CPU reduction comparing Hash against Index only access Further improvements delivered late in the beta program.

These are some more specific situations and performance results.

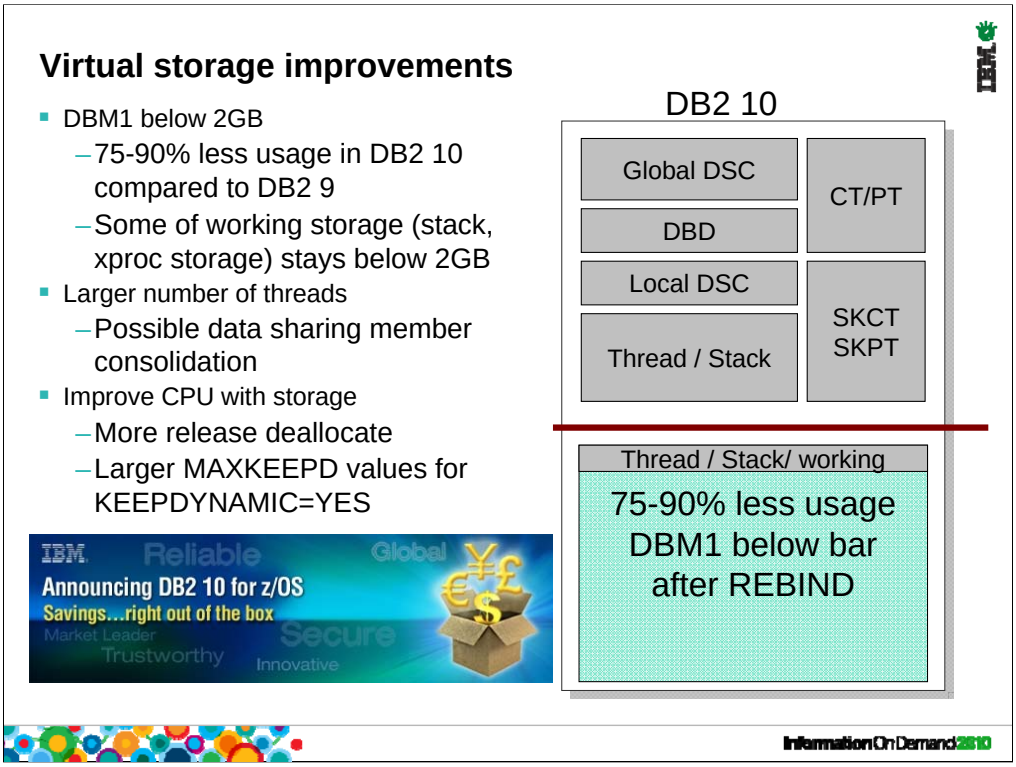

Laboratory measurements and early customer experience have shown substantial savings in the primary constrained address space, DBM1. Most measurements have shown 75% to 90% savings for the virtual storage in that address space below the 2 GB bar. Some EDMPOOL and some working storage remains below the bar.

This storage relief allows many more threads or concurrent users in a DB2 subsystem, allowing new possibilities for optimization.

Some customers will be able to consolidate data sharing members, saving on memory, CPU and administration time.

Other customers will be able to use the storage to improve service or to reduce CPU time more. Some common examples are expected to be use of RELEASE(DEALLOCATE) and larger amounts of dynamic statement cache.

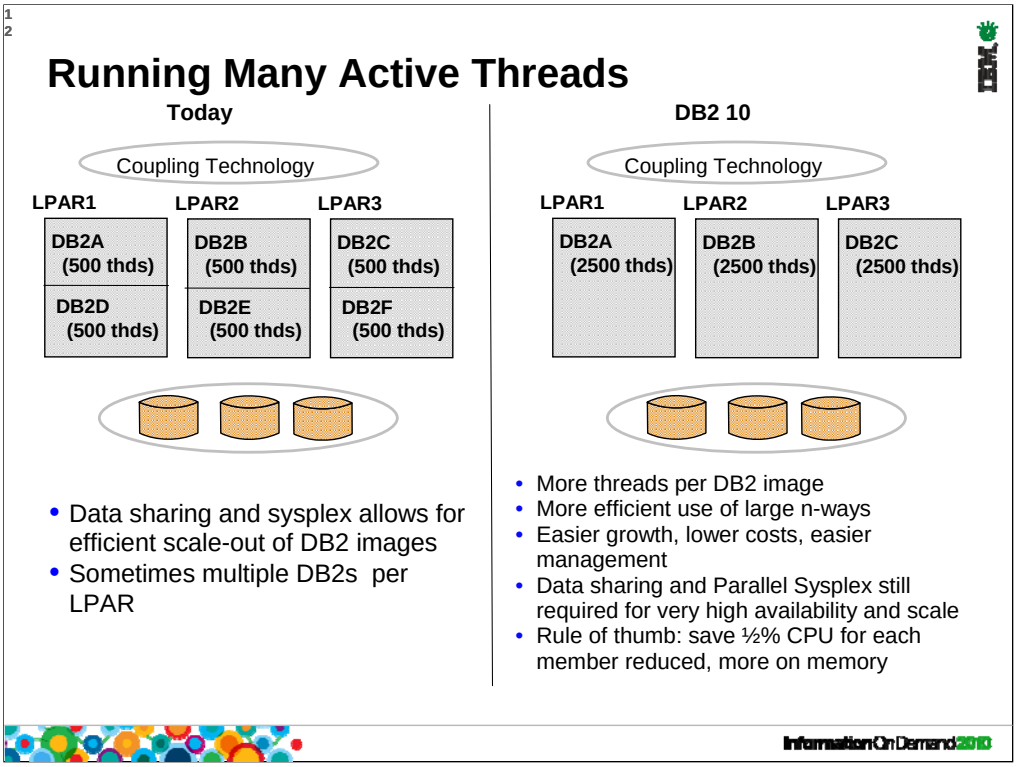

Customers are constrained by virtual memory to various degrees. This slide shows a relatively extreme situation experienced by some customers today. With a maximum of 500 threads (very dependent upon workload) in a DB2 subsystem, this customer is using two DB2 subsystems in the same data sharing group on a single LPAR. This is not efficient for memory of CPU, but avoids the memory constraints with fewer LPARs. Additional relief for virtual storage comes with IMS 11 and other products.

This example allows customers to run 10 times as many threads in a single DB2 subsystem, improving efficiency for storage and CPU. The biggest change is easier management and simpler growth. Most customers use data sharing for high availability, and that need still exists. Extreme scale continues to need data sharing, but fewer data sharing members can mean easier management and reduced resource consumption. In this example, changing from 6 members to 3 can mean a reduction of 1.5% in CPU time, as a rule of thumb.

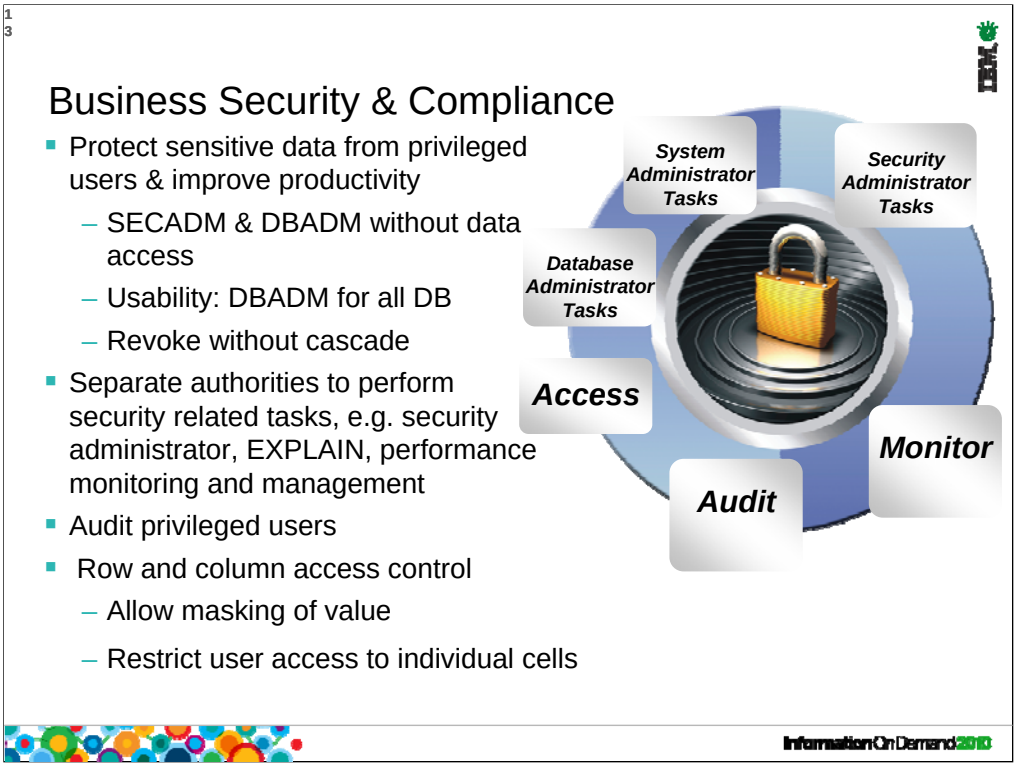

Customers are being pressed for a wide range of improved security and compliance. Data retention is a growing need. Protecting sensitive data from the privileged users and administrators is required. Separation of authority for security, access, and some common tasks, like EXPLAIN will help. Auditing for privileged users can also make compliance simpler.

Access control is refined in several ways with better granularity for the administrative privileges and with finer grained access control at the row and column level, including the ability to mask access to some fields. Auditing is also enhanced.

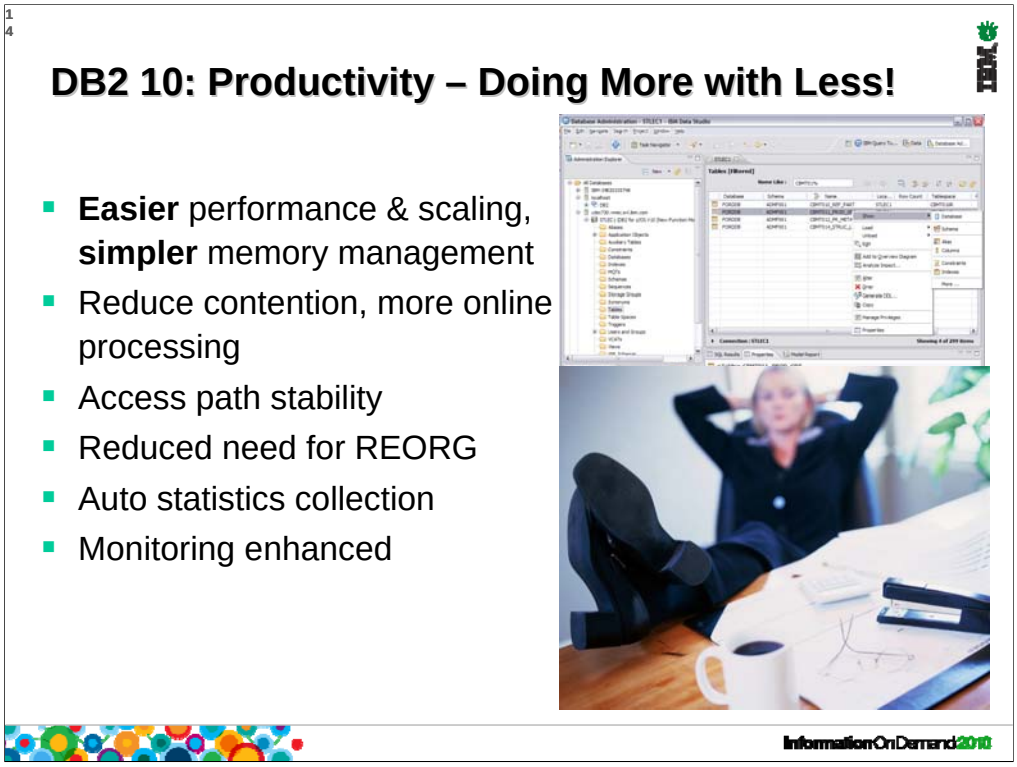

Some of the improvements come with Data Studio for application programming and administration – stronger cross-platform graphical interfaces, better integration with Java, improvements in the ability to develop and debug.

Some of the improvements come within DB2 for z/OS. Improvements in SQL and XML improve productivity for those who develop new applications and for those who are porting from other platforms. Some of the improvements remove complexity from application tasks.

DB2 has a strong focus on making DB2 easier to use by automating tasks and eliminating tasks where possible. Avoiding the manual invocations can also help avoid problems for running the function too often or not often enough. Where the task cannot be eliminated, the frequency and monitoring can be reduced, such as the need to reorganize. The improvements for virtual storage and for availability also help DBA productivity.

Allowing tailored names for DSNHDECP will permit many subsystems to share the SDSNEXIT data set.

Query Processing Enhancements

- **Performance Improvements**
  - Improved caching of dynamic SQL with literals
  - Safe Query Optimization
  - Aggressive View Merge
  - IN List Processing
  - SQL Pagination
  - Parallelism Enhancements
- **Access Path Stability**
  - Relief from package REBIND regression

Access path improvements deliver improved response time and reduced resources and simpler management for packages and queries. The performance improvements include better optimization for some common situations, improved caching for dynamic SQL statements which contain literals instead of parameter markers, and increased parallel processing. Removing some parallel restrictions provides faster response times and allows more use of zIIP.

Access path stability improvements help eliminate regression from a REBIND. The DB2 9 package management changes helped many customers reduce their fear of REBIND, and these improvements take the next step, extending and expanding the capability. New capabilities make the processing more efficient and easier to manage.

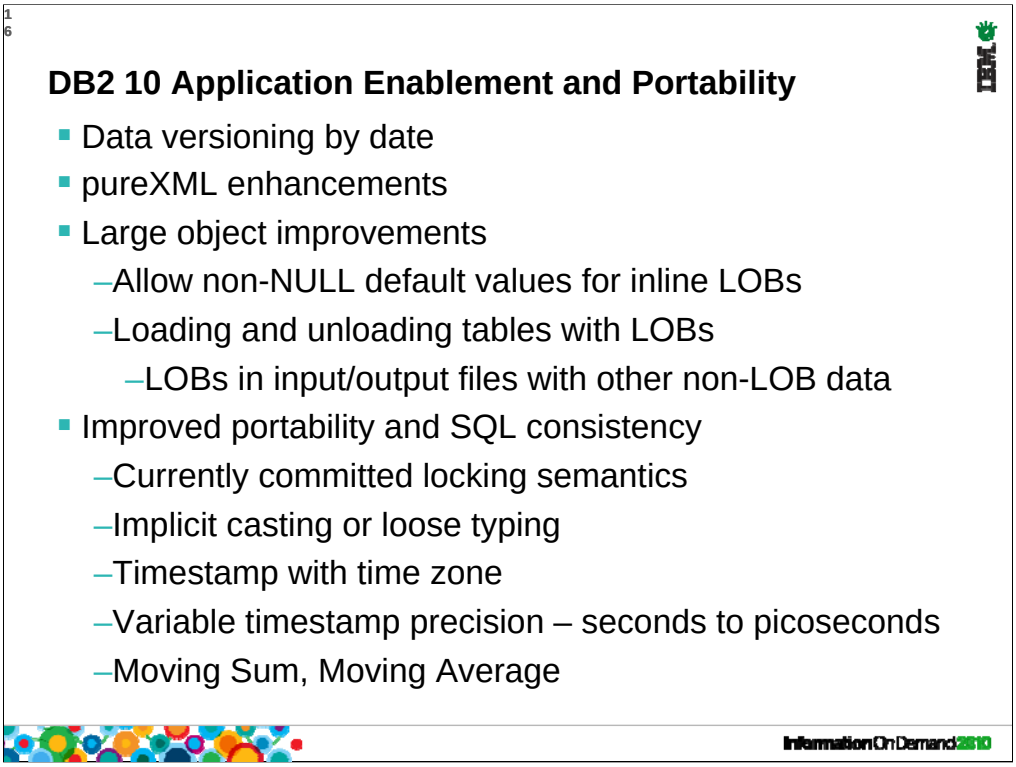

DB2 10 provides the ability to version your data by date. Using the timestamp picoseconds enhancements all the data within a table can have unique timestamps. This enhancement along timestamp time zone capabilities provides the application designer with options for tables that hold global data activities. This is an important feature for global financial and other global industry companies.

Large object are also improved by being able to put smaller LOBs that will fit on the same data page within the data row. This allows the row and the related LOB data to be on the same row and eliminate the extra I/Os to get the LOB data from its other table. This feature combined with the ability to use NULL or default values for the LOB gives additional reasons to have small LOBs in-line with their associated row data.

LOB handling is also improved as the LOB data can be included in the standard input/output files with other non-LOB data. This eliminates the hassles of the large amount of extra LOB files previously needed to support the loading and unloading of LOB data within your system.

DB2 10 also improves its compatibility and SQL consistency within the DB2 family and with other DBMS vendors. This allows any other DBMS vendor systems to be more easily ported to DB2 10 on System z eliminating availability, scalability and performance problems. This feature can quickly resolve your performance, maintenance and scalability for some of your UNIX systems, some SAP or other packaged software systems.

This compatibility also extends to the ability to implicitly cast unlike data types for easily moving or integrating data across application data types, program languages and platforms.

This portability is also reflected in the new package level parameters to control whether the application looks at only currently committed data or not. This improves application concurrency and provides flexibility within your application design for when the system should ignore rows that are in the process of being inserted and only use currently committed rows

The Timestamp with Time Zone and the pico seconds of the timestamps features are a great way to set up a fact table within a data warehouse or business intelligence database. These components along with the new SQL capabilities for calculating a moving sum or moving average are additional DB2 10 capabilities that make it easier for operational business intelligence applications.

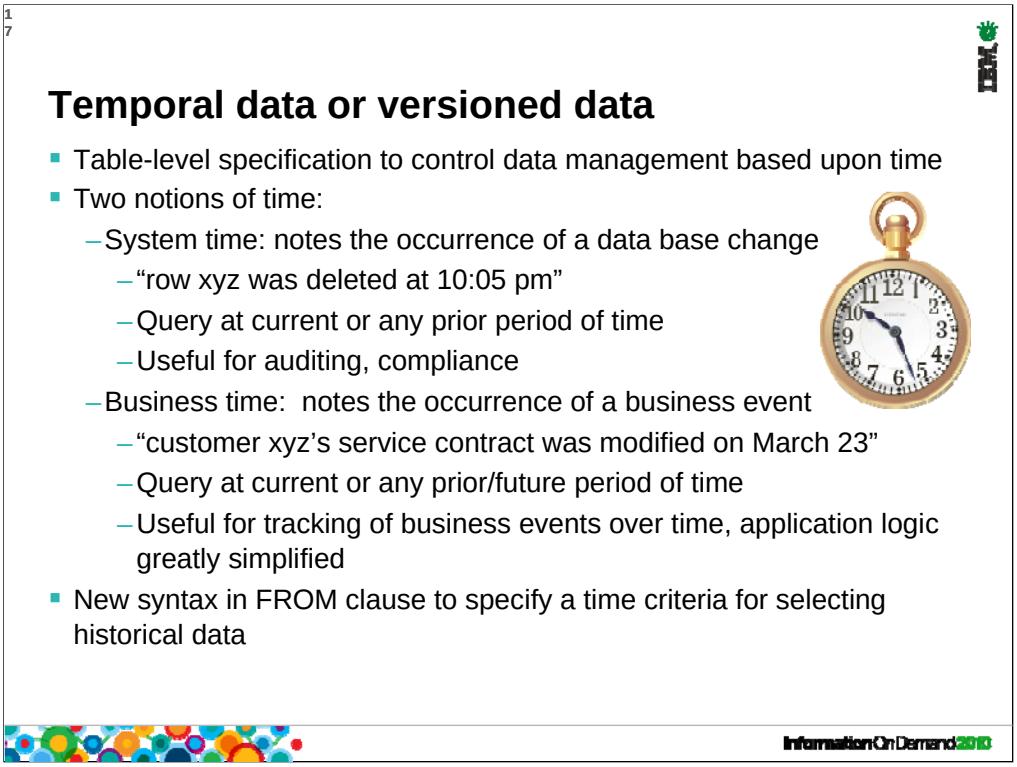

In DB2 10, you can create a temporal table, which is a base table with one or more time periods defined on it. DB2 supports two built-in types of periods, which are the system time period and the business time period. The system time period is a system-maintained period in which DB2 maintains the start and end timestamp values for a row. The business time period is a user-specified period in which you maintain the start and end values for a row.

The SYSTEM_TIME period is meaningful because of versioning. Versioning specifies that old rows are archived into another table. The table that contains the current active rows of a table is called the system-maintained temporal table. The table that contains the archived rows is called the history table. DB2 creates a history table and a table space to hold that table when you define a base table to use versioning, or when you enable versioning on an existing table. You can delete the rows from the history table when those rows are no longer needed.

Using these two built-in periods together in the same table creates a bi-temporal table. You can use a bi-temporal table to keep user-specified period information and system-based historical information. Therefore, you have a lot of flexibility in how you query data based on periods of time.

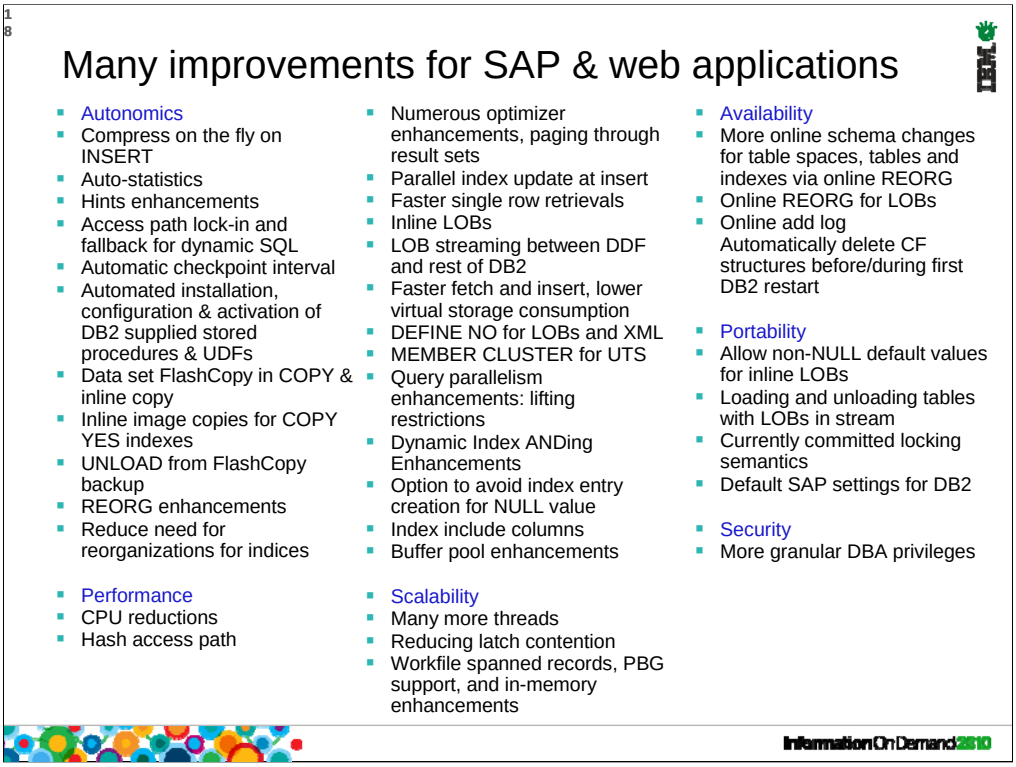

The enhancements cover many aspects of the database technology including new applications support, SQL enhancements, performance and scalability, continuous availability, data warehousing improvements as well as reducing the total cost of ownership.

DB2 10 for z/OS satisfies or partially satisfies many requirements from the worldwide user group communities such as Guide Share Europe, Japan GUIDE/SHARE, and SHARE Incorporated. In addition, this release satisfies many requirements submitted directly to IBM by customers or Business Partners.

As with recent previous releases, Enterprise Applications providers, such as SAP, many other web applications and their customers have been a very important source of the requests for new functions and features.

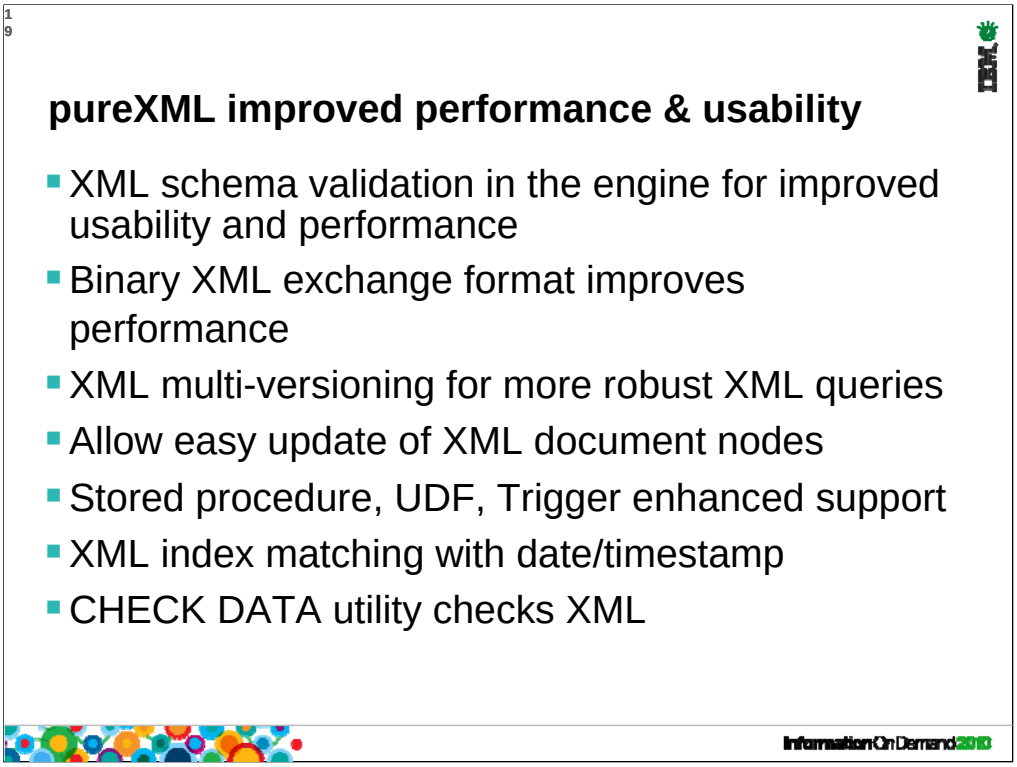

A range of XML improvements delivers a strong release 2 of the pureXML function. Customers use of DB2 9 pureXML shaped this delivery of improved performance and usability.

Multi-versioning: During the execution of a SQL statement, a row with an XML column can be kept in a work file. The row in the work file does not contain the actual XML document. Instead, the information needed for DB2 to retrieve the XML document from the XML table is cached in the work file. The problem occurs if the XML document in the XML table is deleted or updated. When the row in the work file is fetched, DB2 cannot find the expected XML document in the XML table, and the SQL statement fails with an error SQLCODE.

XML UPDATE: Applications which require parts of XML documents to be modified need to break apart the XML document into modifiable pieces, make the modification to a piece, and then construct the pieces back into an XML document.

SP/UDF/Trigger support: XML variables inside SQL PL, XML arguments, transition variables.

The CHECK DATA utility is extended to check XML data.

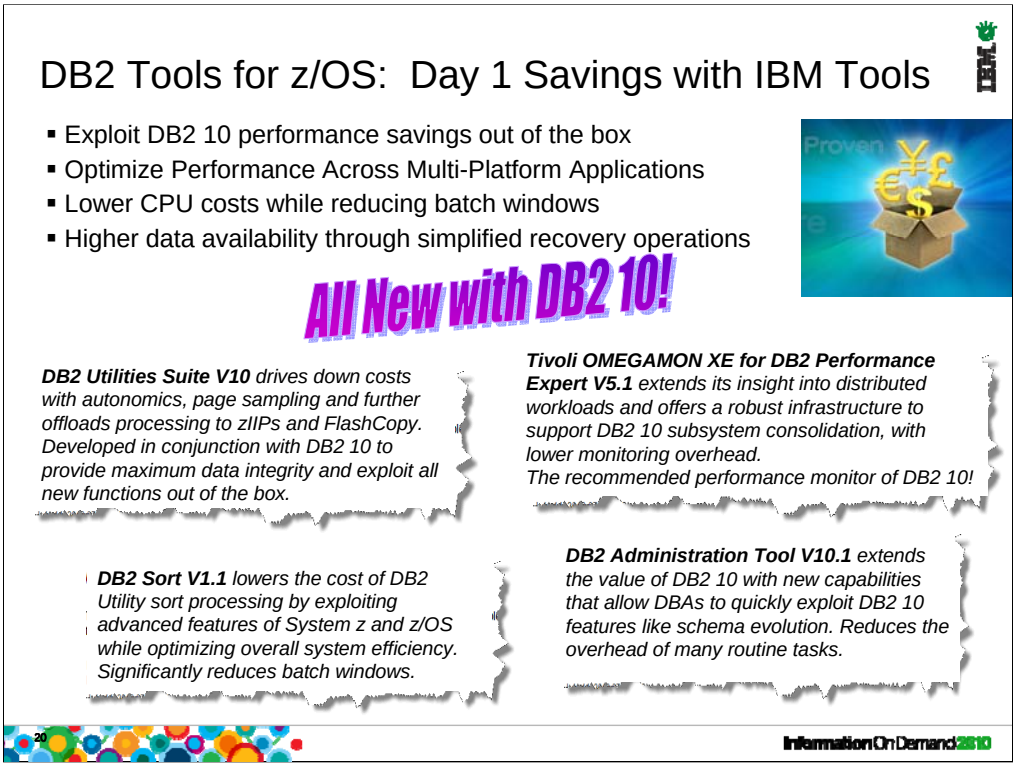

Utilities – Focus on eliminating outages, improving performance, reducing resource consumption, reduce complexity and improve automation. Day 1 utility support for DB2 10 function. Some key enhancements are:
•Flashcopy support at data-set level for COPY, RECOVER, REORG, LOAD, & REBUILD INDEX, For ex: can create an image copy data set to be a transaction-consistent image copy data set with no application outage
•Significant enhancement to REORG to reduce outages, including Support new Online Schema in DB2 10 and hash table, both before and after conversion to hash format, Improved performance for part-level REORG w/ non-partitioning indexes and REORG INDEX to reduce ET, SHRLEVEL CHANGE support for LOB table space for both LOY YES/NO w/ no mapping table required
•Major enhancements to RUNSTATS incl. zIIP support, autonomic features, page & auto-sampling rates
Admin Tool: DBA-managed performance improvements - Easily migrate existing tables to hash access, Manage new Security models and autonomic statistics collection, Rollout Application BiTemporal Data – "as of"
HPU - new internal format to UNLOAD (perf. Feature) to as well as TCPIP support via USS pipe and native XML UNLOAD support
Tivoli OMEGAMON XE for DB2 Performance Expert, v5.1 W5655-W37 introduces an end-to-end response time measurement capability surfacing DB2 for z/OS SQL metrics, making it IBM's most comprehensive DB2 application performance assessment tool. Use it to evaluate the efficiency of, and optimize performance of your DB2 for z/OS DBs across your heterogeneous application environment. It offers all the capabilities of Tiv. OMEGAMON XE for DB2 PM on z/OS and Buffer Pool Analyzer for z/OS, in addition to its own unique capabilities.
Utilities – Focus on eliminating outages, improving performance, reducing resource consumption, reduce complexity and improve automation. Day 1 utility support for DB2 10 function. Some of the key enhancements are:
•Flashcopy support at data-set level for COPY, RECOVER, REORG, LOAD, & REBUILD INDEX, For ex: can create an image copy data set to be a transaction-consistent image copy data set with no application outage
•Significant enhancement to REORG to reduce outages, including Support new Online Schema in DB2 10 and hash table, both before and after conversion to hash format, Improved performance for part-level REORG w/ non-partitioning indexes and REORG INDEX to reduce ET, SHRLEVEL CHANGE support for LOB table space for both LOY YES/NO w/ no mapping table required
•Major enhancements to RUNSTATS incl. zIIP support, autonomic features, page & auto-sampling rates
Admin Tool: DBA-managed performance improvements - Easily migrate existing tables to hash access, Manage new Security models and autonomic statistics collection, Rollout Application BiTemporal Data – "as of"
HPU - new internal format to UNLOAD (perf. Feature) to as well as TCPIP support via USS pipe and native XML UNLOAD support
Tivoli OMEGAMON XE for DB2 Performance Expert, v5.1 W5655-W37 introduces an end-to-end response time measurement capability surfacing DB2 for z/OS SQL metrics, making it IBM's most comprehensive DB2 application performance assessment tool. Use it to evaluate the efficiency of, and optimize performance of your DB2 for z/OS DBs across your heterogeneous application environment. It offers all the capabilities of Tiv. OMEGAMON XE for DB2 PM on z/OS and Buffer Pool Analyzer for z/OS, in addition to its own unique capabilities.

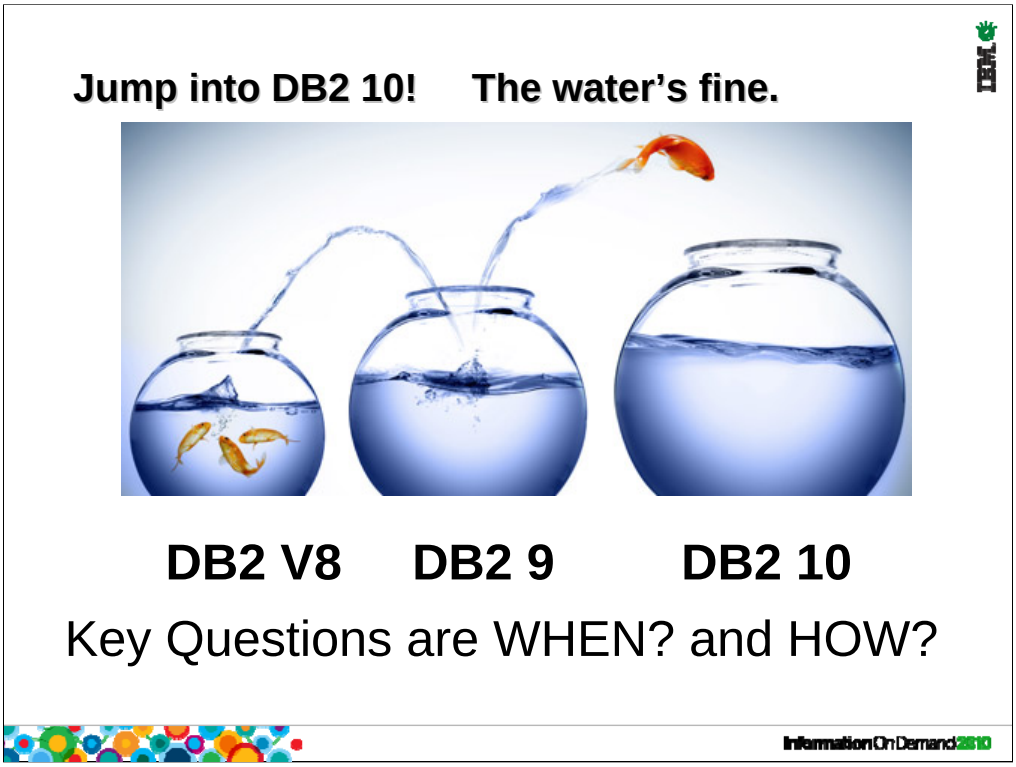

Is your current fish bowl getting constrained? What is limiting you? Is it CPU? Virtual storage? Latching? DB2 catalog and directory? Concurrent utilities? Are you currently running DB2 9? V8? V7? Should you migrate to DB2 10?

The answer is a definite Yes. The question is not so much whether to migrate as when and how to migrate. If you are running DB2 9 today, then DB2 10 is in your future, giving you more room to grow, with fewer limits, lower costs, and more for less. If you are running DB2 V8 today, then you have a choice of jumping to DB2 9 or directly to DB2 10. So the key question is, "When should I migrate to DB2 10?"

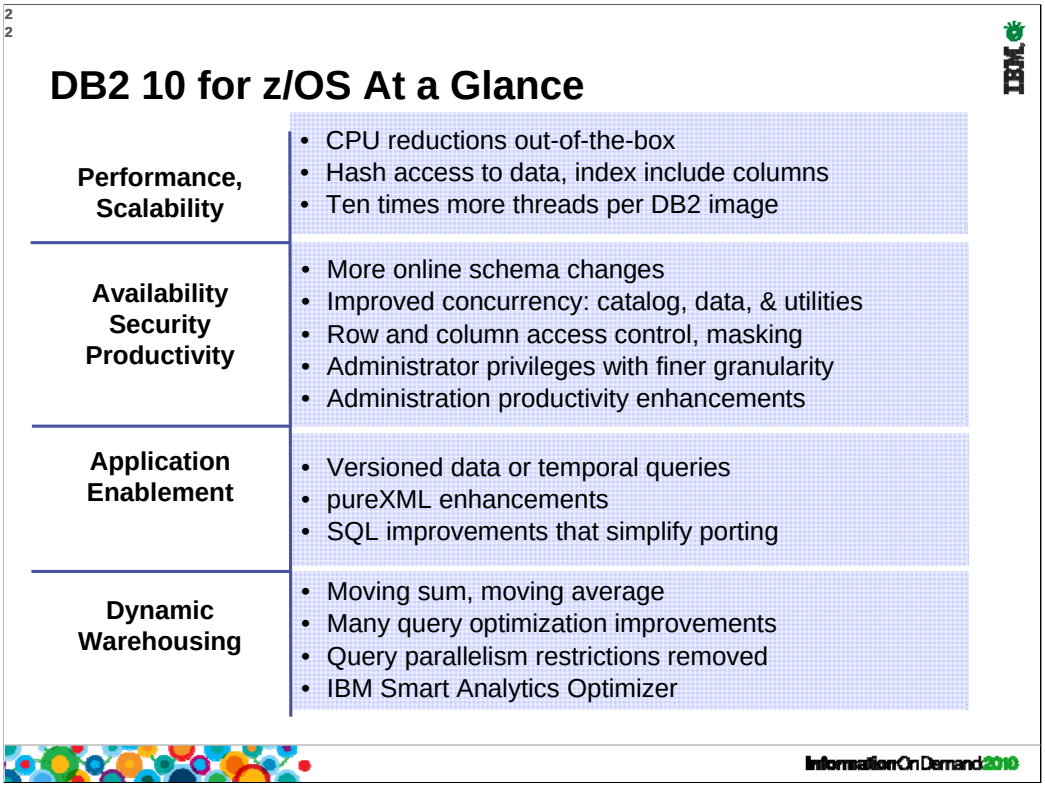

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for 22 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10%. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improves availability. SQL, pureXML, web services extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity. DBAs can avoid running statistics, some REORGs, and benefit from memory, and utilities enhancements. Warehousing continues to evolve, with improvements in SQL and XML, better optimization techniques, increased parallelism and the new IBM Smart Analytics Optimizer.

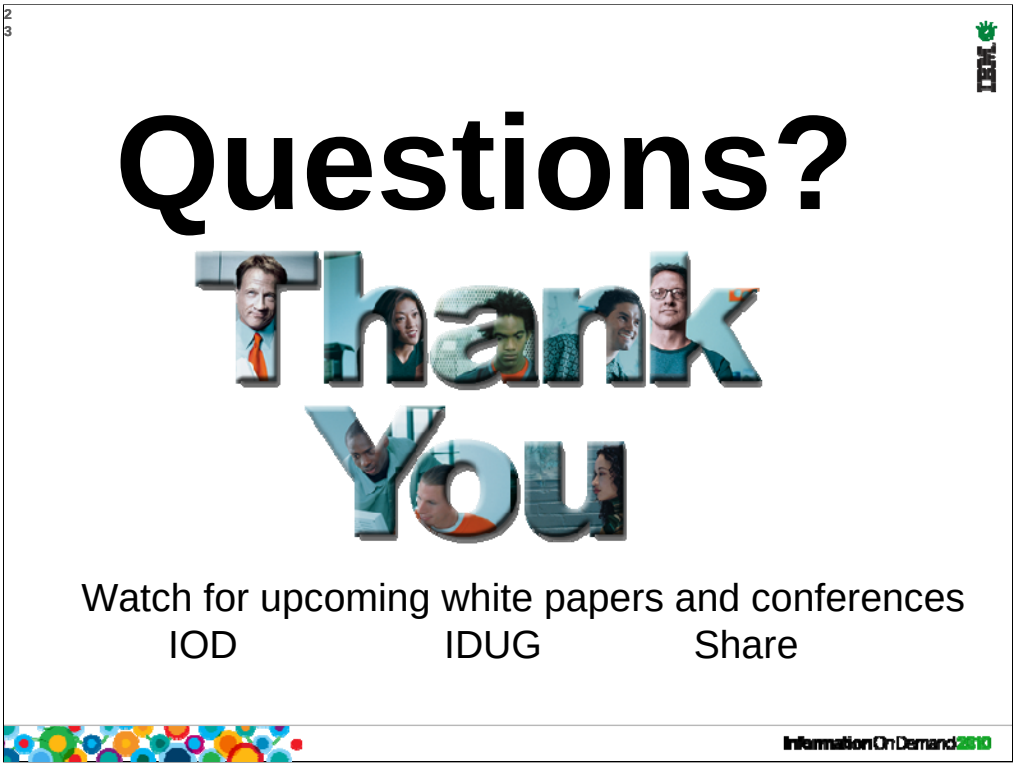

I hope we have lots of questions, but thank you in any case. We have provided many additional notes in the slides and additional slides at the end of this presentation.

We have a new white paper on DB2 10 from Dave Beulke now and expect to have another from Julian Stuhler of Triton Consulting in the future.

https://www14.software.ibm.com/webapp/iwm/web/signup.do?source=sw-infomgt&S_PKG=db2_zos_reduce_costs

DB2 10 is a hot topic at upcoming conferences, so please plan to attend IOD, IDUG or Share.

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

http://www.idug.org

http://www.share.org

http://www.ibm.com/software/uk/data/conf/

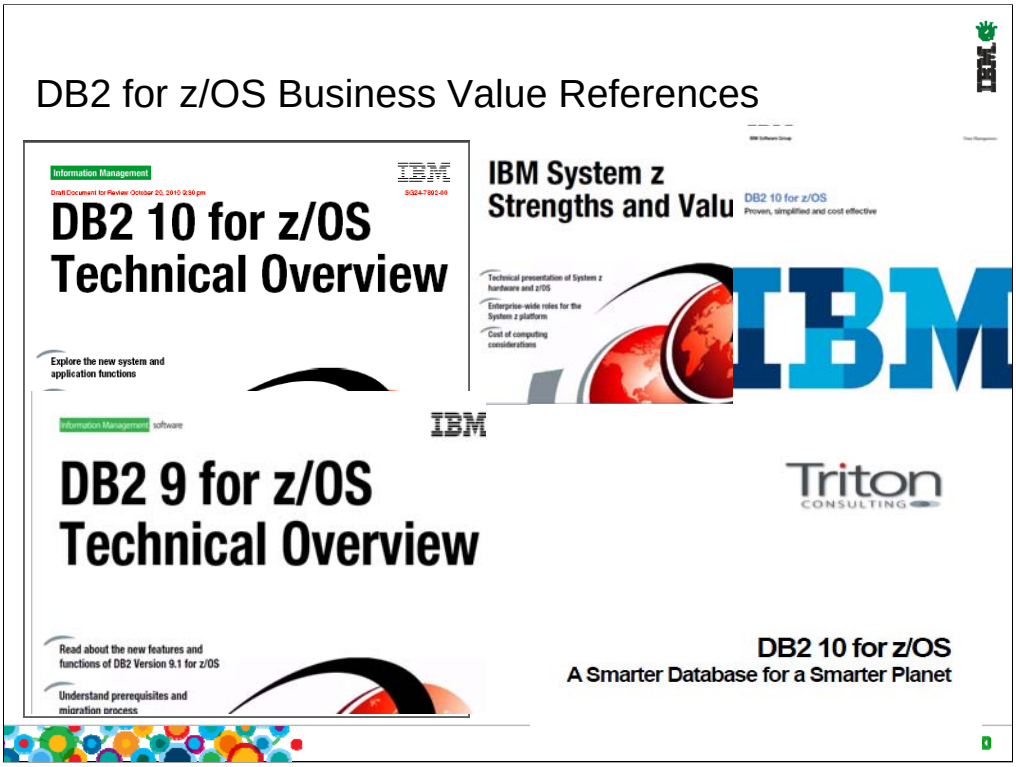

Here are some resources for a discussion of business value in DB2. White papers are useful for a quick summary, and IBM Redbooks publications provide more detailed technical discussion.

http://www.ibm.com/software/data/db2/zos/db2-10/

ftp://ftp.software.ibm.com/software/data/pubs/papers/DB2_for_zOS_V9_Business_Value_White_Paper.pdf
http://www.ibm.com/software/os/systemz/newsletter/mainstreamed11_uk.html

The Business Value of DB2 for z/OS, SG24-6763,
http://www.redbooks.ibm.com/abstracts/sg246763.html

DB2 9 for z/OS Technical Overview, SG24-7330, chapter 2 on System z synergy,
http://www.redbooks.ibm.com/abstracts/sg247330.html

DB2 9 for z/OS Performance Topics, SG24-7473, performance on z10,
http://www.redbooks.ibm.com/abstracts/sg247473.html

What DB2 Means to Business Value, Talk DB2 in business terms the suits can understand.
http://www.ibmdatabasemag.com/story/showArticle.jhtml?articleID=199201819

System z Strengths and Values, SG24-7333,
http://www.redbooks.ibm.com/abstracts/sg247333.html

http://www.idug.org/wps/portal/idug/kcxml/04_Sj9SPykssy0xPLMnMz0vM0Y_QjzKLN4o38rAESYGYnoH6kehCAQghX4_83FT9IKBUpDlQyMzDRz8qJzU9MblSP1jfWz9AvyA3NKLc29ERAP8BqUk!/delta/base64xml/L0lJSk03dWlDU1EhIS9JRGpBQU15QUJFUkVSRUlnLzRGR2dkWW5LSjBGUm9YZmcvN18yXzZITA!!?PC_7_2_6HL_WCM_CONTEXT=/wps/wcm/connect/IDUG+Site/Solutions+Journal/Solutions+Journal+Online+Magazine/Volume+15%2C+Number+2/IDUG-SolutionsJournalArticle-Content-Developer-s+Corner%3A+Time+to+Plan+for+Migration+Again

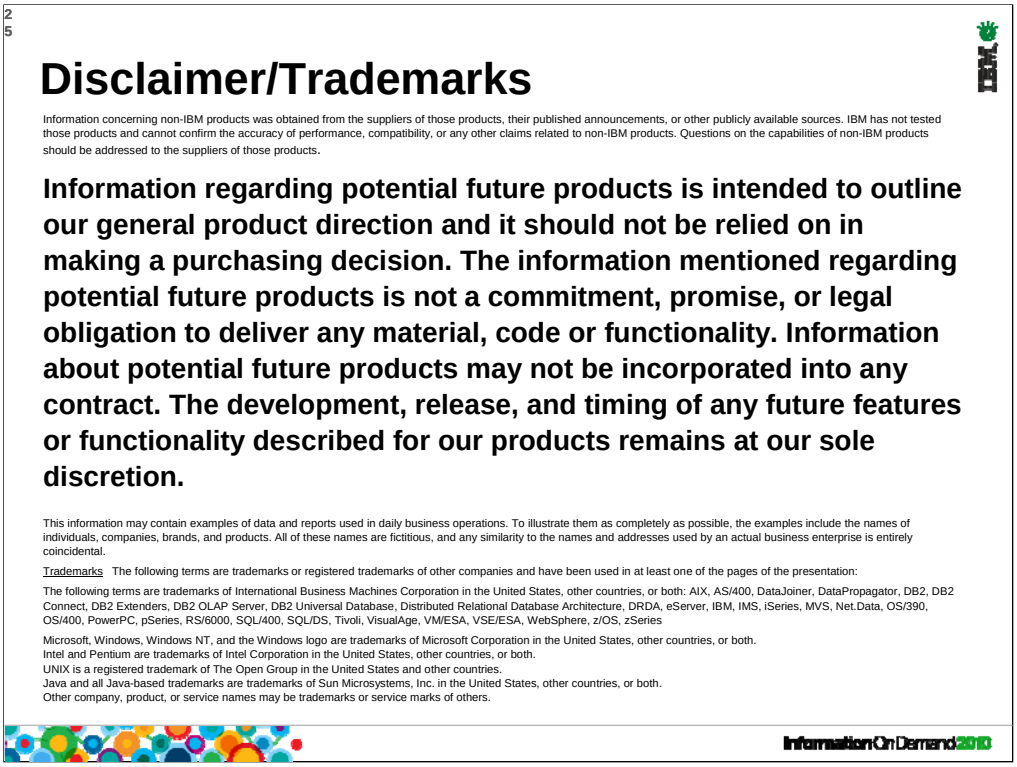

This product became generally available on October 22, 2010.

See the web for current information, the announcement, …

http://www.ibm.com/software/data/db2/zos/db2-10/

http://www.ibm.com/common/ssi/rep_ca/5/877/ENUSZP10-0015/ENUSZP10-0015.PDF

http://www.ibm.com/support/docview.wss?uid=swg27017960

http://www.ibm.com/support/docview.wss?uid=swg21006951

http://it.toolbox.com/blogs/db2zos/db2-10-for-zos-beta-announced-today-36790

http://www.ibm.com/developerworks/spaces/db2zos

http://davebeulke.com/?p=625

http://community.solutionscenter.techweb.com/community/mainframe/blog/2010/02/09/db2-10-for-zos-beta-starts-today

http://www.triton.co.uk/blog/?p=415