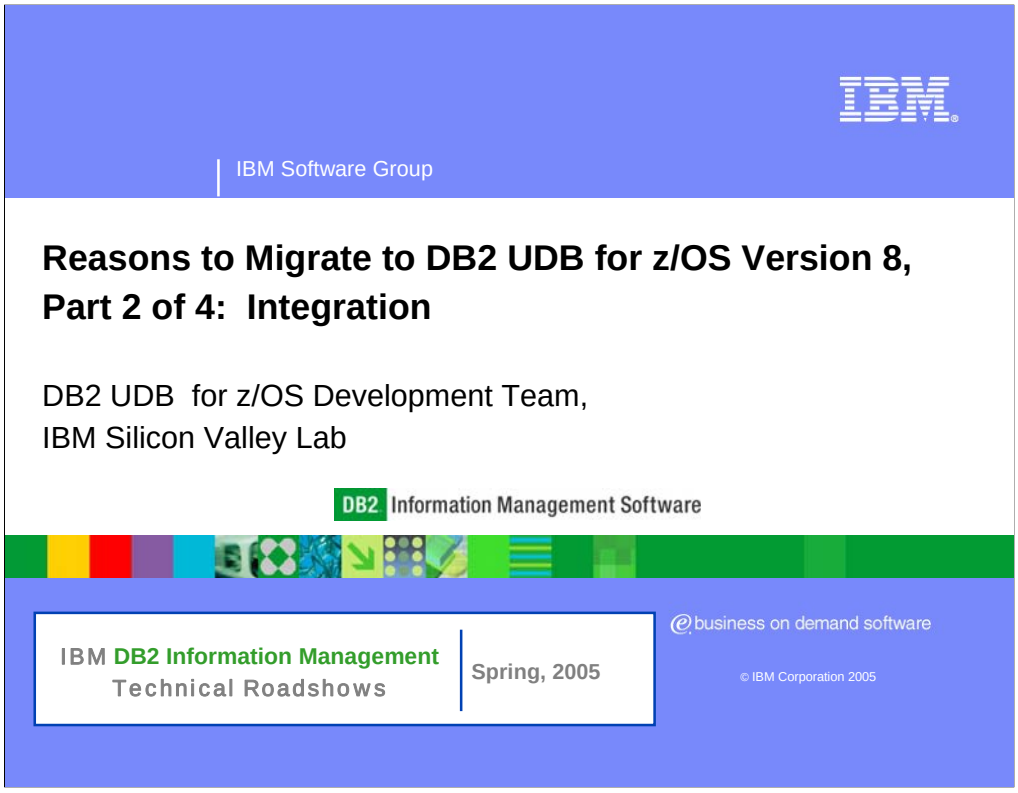

Find out how DB2® Universal Database for z/OS® Version 8 has been re-engineered, with fundamental changes in architecture and structure that will help you manage your very large databases more easily and cost effectively. "Reasons to Migrate to DB2 UDB Version 8" is a day long seminar suited to database technical managers, database administrators, applications developers and systems programmers. It focuses on four critical areas: availability, integration, application development productivity, and flexible growth and incremental scalability. Its information-packed sessions will familiarize you with the enhancements that are enabling organizations to streamline database management, respond more flexibly and quickly to business needs and ensure short and long-term growth capacity.

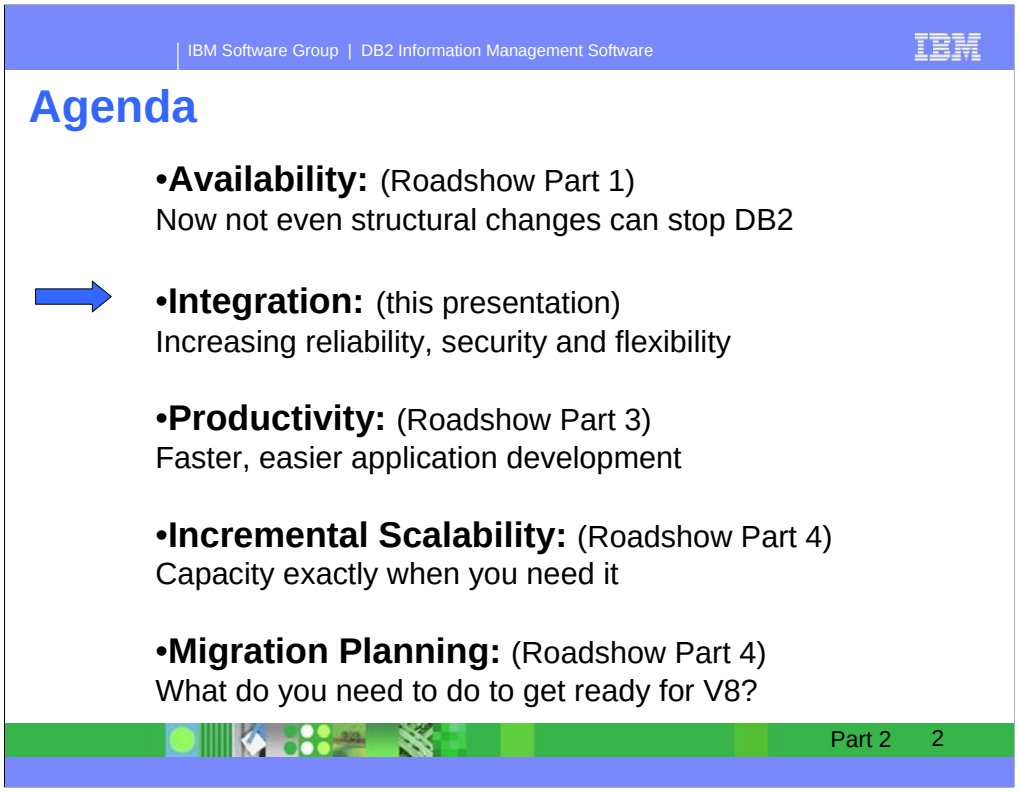

Many of the items I'm talking about today improve performance, availability, productivity and scalability. We'll try to keep the items that are closely related together, but categories don't always work when you need to choose and the real answer is both.

The agenda for today is roughly
8:15 am Registration and Continental Breakfast
8:45 am Welcome and Introduction
9:00 am Availability: Now not even structural changes can stop DB2
10:30am Break
11:00 am Integration: Increasing reliability, security and flexibility
12:30 pm Lunch
1:30 pm Productivity: Faster, easier application development
3:00pm Break
3:30 pm Incremental Scalability: Capacity exactly when you need it
Migration Planning: How to get there
5:00 pm Close of Program

Customer Value Highlights: Integration Theme

- **Reliability & security:** zSeries platform leverages unique synergy between the hardware, operating system, and database to provide a highly reliable and secure DBMS
- **Access to business data anywhere in the enterprise:** QMF provides enhanced data visualization via multiple interfaces for robust secured access to key business data
- **Easier application deployment:** WebSphere & Java with DB2 application connectivity feature of DB2 makes it easier than ever to deploy new Java applications on zSeries, no matter where you develop them
- **Increased effectiveness & productivity:** Data Management Tools enhance the value of DB2 by automating & simplifying key tasks, leveraging new V8 functions from day one

We will discuss integration with the zSeries and z/OS platform and with the middleware here. Integration with the applications is noted in the other sections, especially part 3.

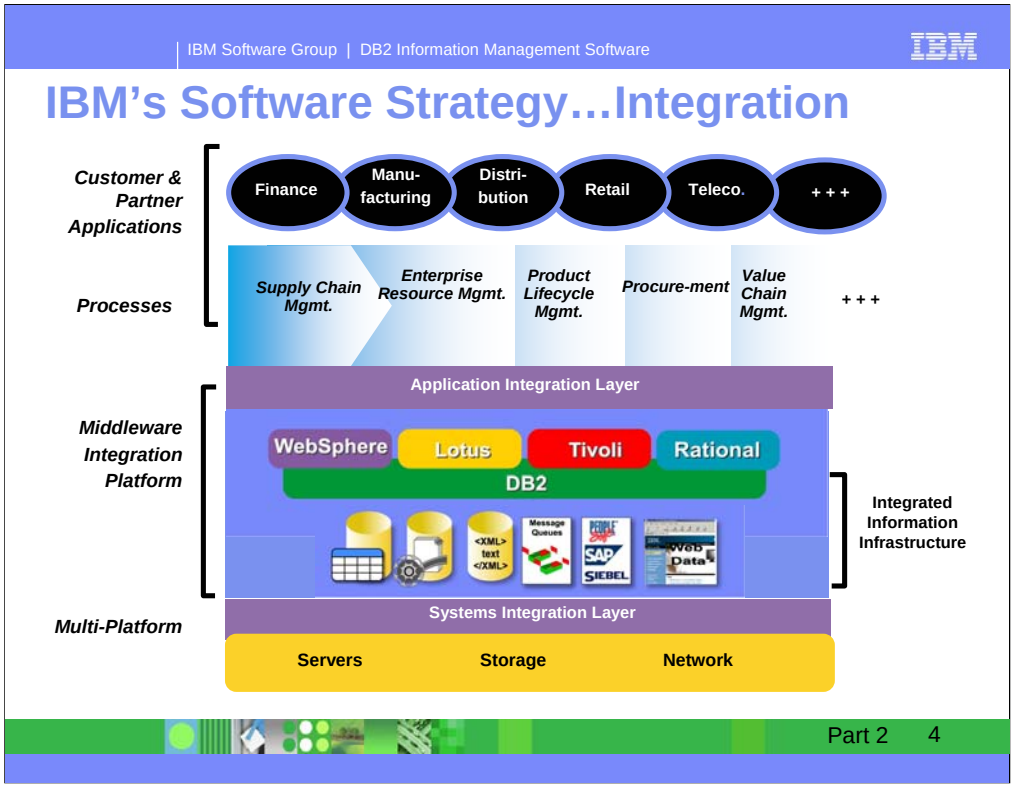

This is our map for infrastructure integration, working to integrate well with applications, across the middleware layer, and deep systems integration with our platforms: operating systems and hardware..

This section will discuss some examples of the platform integration and the middleware. Integration with applications from SAP, PeopleSoft and Siebel is included across every section.

IBM middleware has an unmatched breadth and depth of offerings. Today, all of our products and market-leading brands - including our leading transaction management, data management, collaboration, and systems management solutions - are evolving to provide customers with a scalable and modular computing environment -- one which, by virtue of being open, integrated, virtualized and autonomic, meets the rigorous computing requirements of the on demand era.

First, we will look briefly across the options and improvements in the zSeries and z/OS platform, and discuss how DB2 uses those facilities for improved:

- Availability
- Integrity and security
- Scalability and performance

DB2 UDB for z/OS Version 8 & zSeries z990 z890

1.6X Faster Processors
Up to 32 Processors
More memory, better value; 64 bit virtual storage
✓New backup and restore
✓Multilevel Security

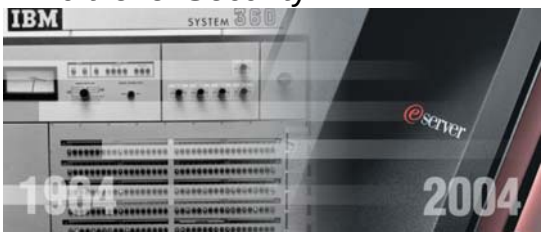

✓Parallel Sysplex enhancements
✓Unicode conversion
✓Compression
✓Cryptography
✓zSeries Application Assist Processor
✓z/Architecture Long-displacement Facility
✓WLM …

These are the key functions of the latest zSeries z990 and z890 that DB2 UDB for z/OS uses, almost everything to deliver **zSeries® and z/OS™** synergy**.** DB2 has used the function of the zSeries and z/OS platform extensively for many years. DB2 benefits from zSeries large real memory support, faster processors, and better hardware compression. DB2 uses Parallel Access Volume and Multiple Allegiance features of the IBM Enterprise Storage Server™ (ESS). ESS FlashCopy® is used for DB2 backup in combination with log suspend / resume. DB2 makes unique use of the z/Architecture™ instruction set, and a number of instructions provide improvement in reliability, performance and availability. DB2 continues to deliver synergy with hardware data compression, FICON™ (fiber connector) channels, disk storage, advanced networking function, and Workload Manager (WLM).

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

zAAP Architecture and Workflow: Executing Java under IBM JVM control

- Java work unit "eligible" to be dispatched on zAAP
- z/OS dispatcher attempts to dispatch eligible work on zAAP
- zAAP ineligible work only dispatched on standard processors
- May dispatch zAAP eligible work on standard processor
  - installation control

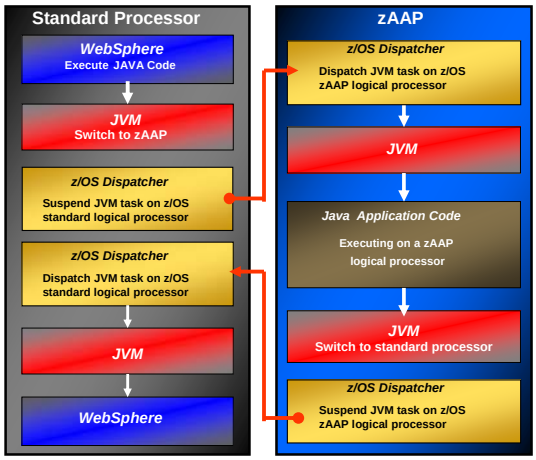

IBM JVM, parts of LE runtime, and z/OS Supervisor needed to support JVM execution can operate on zAAPs

IBM JVM communicates to z/OS dispatcher when Java code is to be executed

When Java is to be executed, the work unit is "eligible" to be dispatched on a zAAP

z/OS dispatcher attempts to dispatch zAAP eligible work on a zAAP (when present)

zAAP ineligible work only dispatched on standard processors

If there is insufficient zAAP capacity available, or standard processors are idle, the dispatcher may dispatch zAAP eligible work on a standard processor

There is an installation control to limit the use of standard processors to execute zAAP eligible work (see Java code execution options)

z/OS Support

		z990 z890	z900 z800	G5/G6 Multiprise 3000	G3-G4	End of Service	Coexistence Migration Policy	Ship Date
OS/390	2.10*	xc	x	x	x	9/04	1.4	
z/OS	1.1*		x	x		3/04	1.4	
	1.2*	xc	x	x		10/04	1.5	
	1.3	xc	x	x		3/05	1.6	
	1.4	x	x	x		3/07	1.7	9/02
	1.5	x	x	x		3/07*	1.8	3/04
Current Releases	1.6	x	x			9/07*	1.8	9/04
	1.7*	x	x			9/08*	1.9*	9/05*
	1.8*	x	x			9/09*	1.10*	9/06*

* Planned dates and releases

x^{c} - Compatibility support – does not exploit new z990 features: 30 LPARs and multiple Logical Channel SubSystems

Bimodal Accommodation Offering is available for z/OS 1.2, 1.3, and 1.4. It will not be provided for z/OS 1.5

z/OS 1.6 has been available since September 2004. It has the same hardware prerequisites as DB2 UDB for z/OS Version 8. If you plan to be current, you need to be on this z/OS release by early 2007.

Note the z/OS end of service dates as well. Note that end of service for z/OS 1.3 is March 2005. Check the latest information by going to the z/OS web page, then to Support and then to the z/OS, z/OS.e, and OS/390 marketing and service announce, availability, and withdrawal dates
http://www.ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html

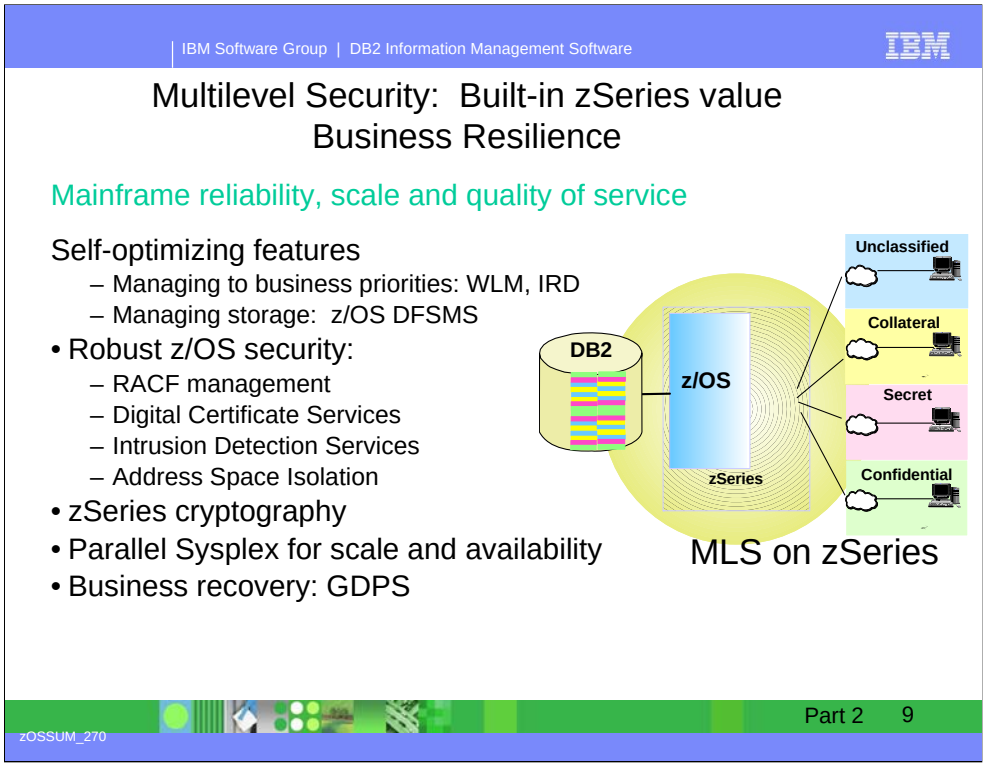

z/OS 1.5 and RACF 1.5 or Security Server add another type of security, called multilevel security, labeled security or mandatory access control (MAC) to our capabilities. The only option in the past with a high degree of separation has been physical separation. In the database world that might mean another machine or LPAR or perhaps another subsystem, another database or another table. With multilevel security, we still have a high degree of security even with data in the same table.

Access control is consistent across many types of resources using RACF, so that multilevel controls apply for data sets, for communications, for print and for database access – both objects and now with row level granularity. The DB2 controls are for both SQL access and for utility access.

For an more on multilevel security, see **Planning for Multilevel Security and Common Criteria (GA22-7509)**

http://publibz.boulder.ibm.com/epubs/pdf/e0z2e111.pdf

Multilevel Security and DB2 Row-Level Security Revealed, SG24-6480

DB2 QMF Version 8

DB2 V8 support

QMF for Windows
- Visionary Studio

QMF for WebSphere
- Browser access
- Java class & Web services

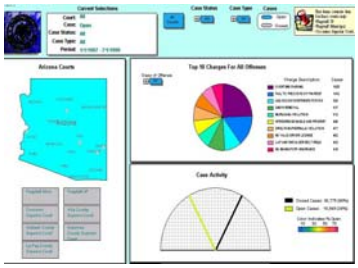

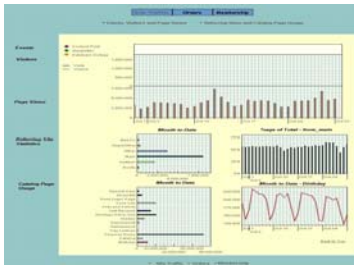

What's in QMF Version 8? Easier, faster, and more global on-demand access to enterprise data and analysis through support for DB2 Version 8 plus:

QMF for Windows: - new drag-and-drop data visualization with Visionary Studio (in addition to existing summary reports, charts, and spatial data maps): across / pivot / top formatting, conditional formatting, rich HTML reports, multi-dimensional analysis (OLAP), visual query building interface, new visual database explorer, support for DB2 V8.1 features
QMF for WebSphere: new, greatly enhanced user interface for Web-based data access through an ordinary browser, visual display of customized report libraries, rapid, robust query development: Expression Builder, Java class API & Web services API for custom Web-based applications.
http://www.ibm.com/software/data/qmf/

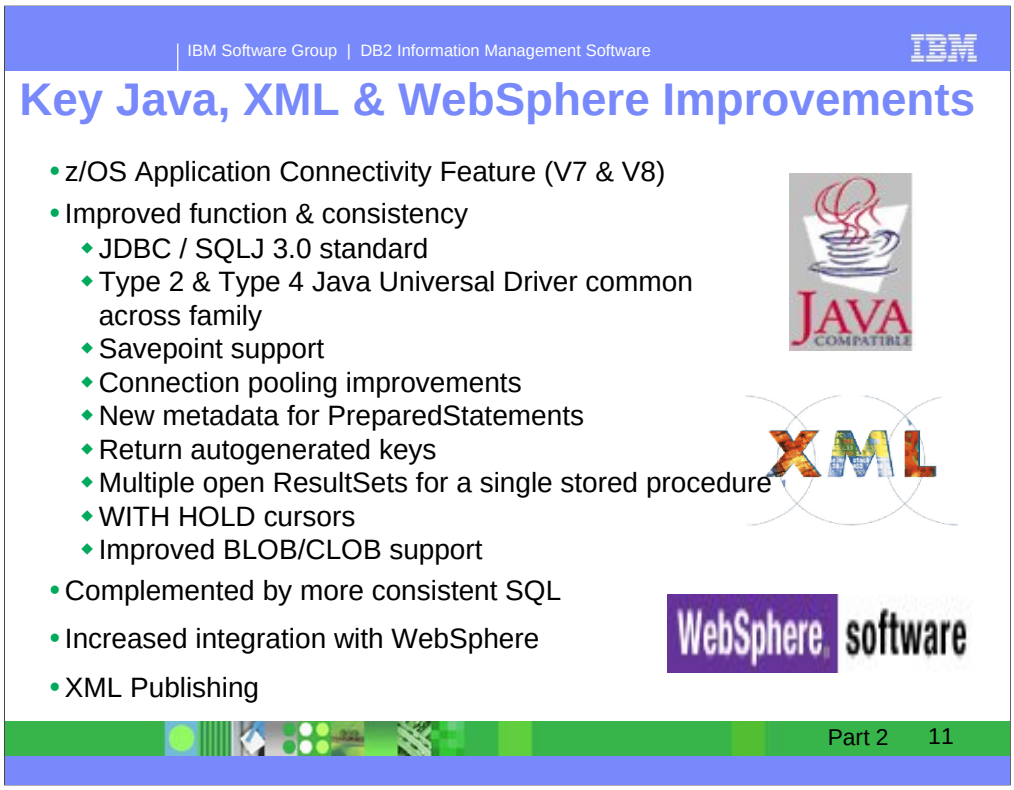

Java support will be more consistent across platforms as we use a single Java code base across the DB2 family. The improved consistency also adds new function to DB2 and improves integration with WebSphere and Java.

The Java Universal Driver is updated to support the JDBC/SQLJ 3.0 standard, including improvements like savepoints, connection pooling improvements, the ability to reuse PreparedStatements, multiple open ResultSets for a single stored procedure, WITH HOLD cursors, and improved BLOB and CLOB support. Very substantial improvements in Unicode, allowing join of Unicode tables with EBCDIC and ASCII and converting DB2 catalog to Unicode.

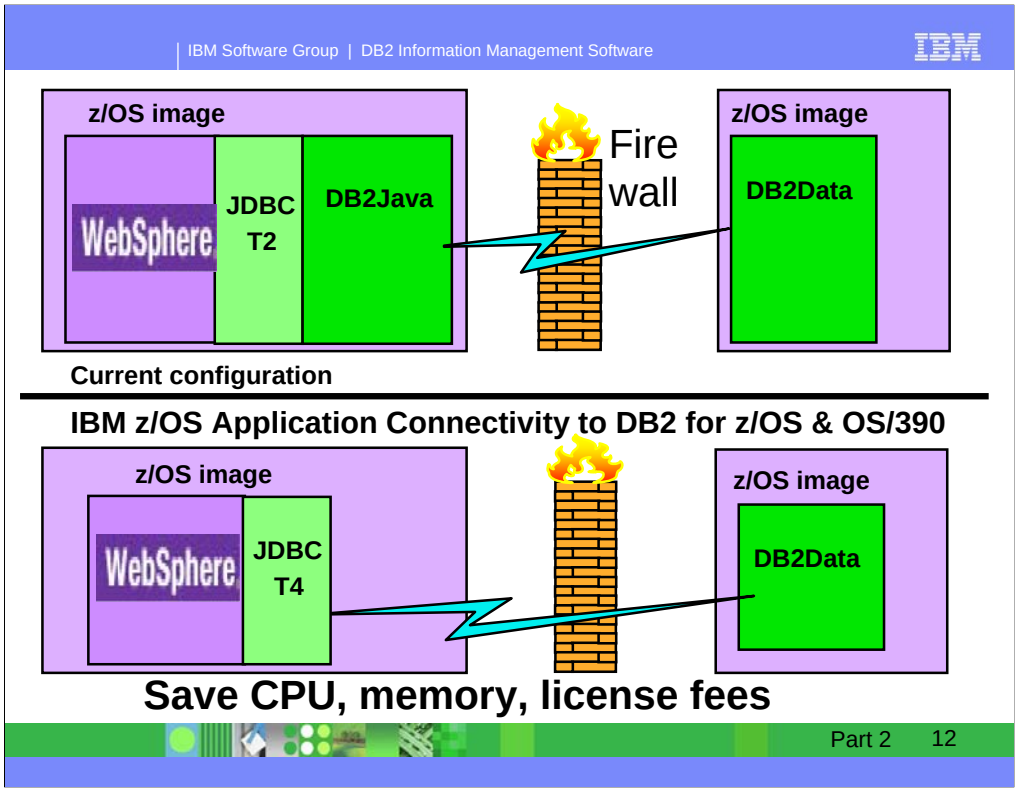

z/OS ® Application Connectivity to DB2 ® for z/OS and OS/390 ® is a no-charge, optional feature of DB2 Universal Database ® Server for z/OS V7 and V8. This feature consists of a component known as the DB2 Universal Database Driver for z/OS, Java ™ Edition, a pure Java, type 4 JDBC driver designed to deliver high performance and scalable remote connectivity for Java-based enterprise applications on z/OS to a remote DB2 for z/OS database server. The driver:

· Supports JDBC 2.0 and 3.0 specification and JDK V1.4 to deliver the maximum flexibility and performance required for enterprise applications

· Delivers robust connectivity to the latest DB2 for z/OS and WebSphere ® Application Server for z/OS

· Provides support for distributed transaction support

· Allows custom Java applications that don't require an application server to run in a remote partition and connect to DB2 z/OS

See the December 16, 2003 announcement or the web for more:

http://publib-.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/tips0356.html?Open

Utility Improvements

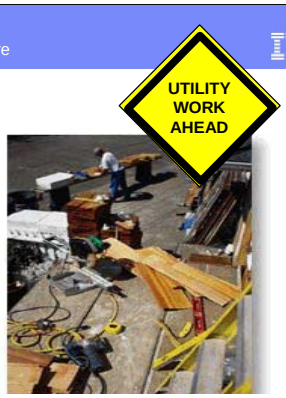

- **Easier restart** (V7 PQ72337)
- **Schema Evolution**
  - **REBALANCE partitions**
- **On-line REORG Enhancements**
  - **DISCARD**
  - **Avoid BUILD2 with DPSI**
  - **REORG DB2 catalog SHRLEVEL REFERENCE**
- **LOAD & UNLOAD delimited input & output**
- **SCOPE PENDING**
- **RUNSTATS non-uniform statistics on non-index columns**
- **System-level log point backup and recovery**
- **Improved defaults for performance**
- **LOBs in cross load, LOAD, UNLOAD, online check index**

Many utility enhancements are part of the base changes in this version, supporting long names, Unicode, 64 bit addressing, DPSIs, system backup and recovery, multilevel security and schema evolution. These utility enhancements improve our value for the money.

Schema evolution uses utility support to rotate the first partition to the last partition. The new REBALANCE function can balance the sizes of a partition range or of all partitions.

The REORG DISCARD can be performed with SHRLEVEL CHANGE. DPSIs can be reorganized without a BUILD2 phase. The DB2 catalog tables can all be reorganized in SHRLEVEL REFERENCE or read only mode.

Delimited files can be used as input to LOAD or output from UNLOAD.

SCOPE PENDING provides improved usability. SCOPE PENDING indicates that only partitions in a REORP or AREO* state for a specified table space or partition range are to be reorganized.

ftp://ftp.software.ibm.com/software/db2storedprocedure/db2zos390/techdocs/Z06m.pdf

Enhanced Instrumentation

- Package level accounting
- Accounting rollup for DDF & RRSAF
- New and improved traces, larger monitor area
  - Long-running reader
  - Lock escalation
  - Full SQL statement
  - PREPARE attributes
  - High water marks
  - Secondary authorization ids
  - Storage 64-bit, READS
  - Dynamic statement cache
  - Temporary space usage
  - Auditing for multilevel security
  - Option for EBCDIC or Unicode data

Many enhancements are made in instrumentation, helping to monitor and account for the larger and more varied workloads. Additional information is provided at a package level, if those traces are on. Accounting can roll up multiple trace records into one for DDF and RRSAF. A new IFCID is provided for lock escalation. The full SQL statement (not just 5000 bytes) can be traced with a new IFCID. The PREPARE statement attributes can be traced. The statement id is added to dynamic statement cache traces. Secondary ids can be retrieved with a synchronous read in an APAR that was added to V6 and V7. Additional fields were added to storage IFCIDs 225 and 217 for 64 bit addressing. Dynamic statement cache traces were improved to be more usable. A new IFCID 0342 was added for temporary space use by agents. Auditing was added for multilevel security. See the Release Guide, Appendix F for new and changed instrumentation.

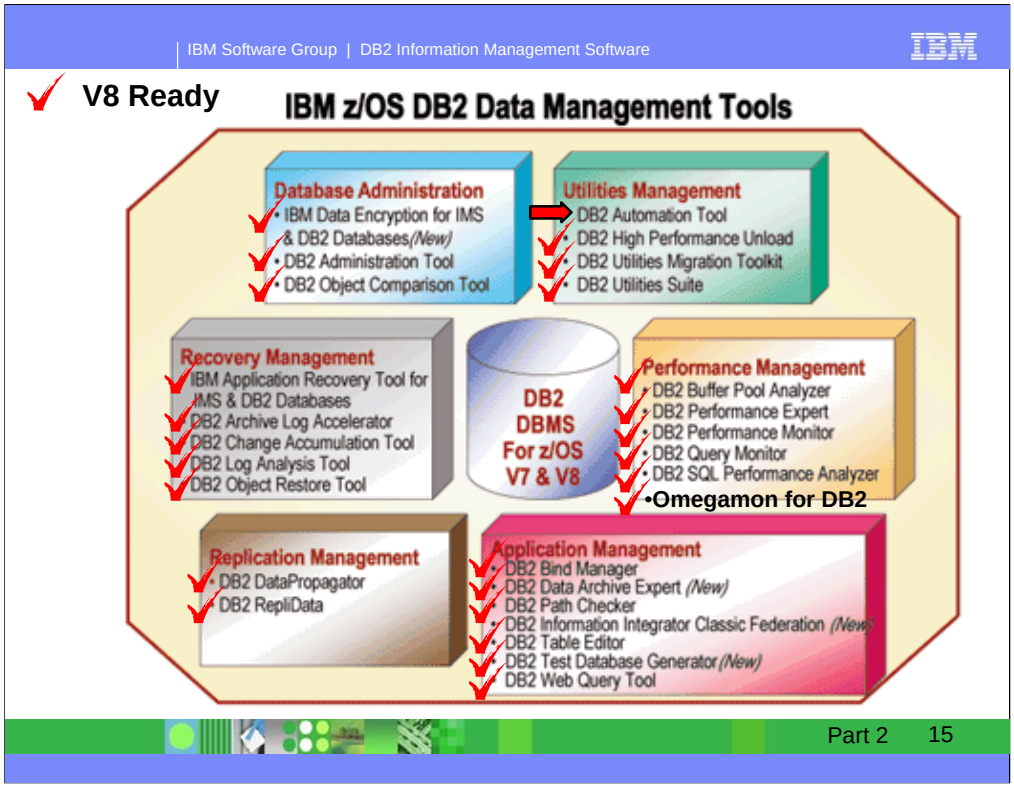

Check marks indicate tools which were ready for DB2 UDB for z/OS V8 by June 2004. All are ready now. With the exception of DB2 RepliData, ALL of these tools now exploit, not just tolerate, new DB2 V8 functions. See the tools web site for details and planning information. Click on Support to see exactly which levels are needed for V8.
http://www.ibm.com/software/data/db2imstools/

Fundamental to our tools strategy is to be able to extend and exploit the data base. To that end all of our tools exploit the features of DB2 V8 immediately at GA time.Our long term goal is to create tools that provide expert advice and automatic management features for DB2 to enable DB2 environments on all platforms to be easier to manage, require less administrative effort and less expertise to get outstanding performance and results.

Tools are a long term and strategic initiative for IBM. We have increased our investment every year since we started in 2000. This year is no exception and we will be releasing new tools, releases, versions and features every quarter, even as the portfolio is broadened. Candle tools have been added to our portfolio, and they are also ready for V8.

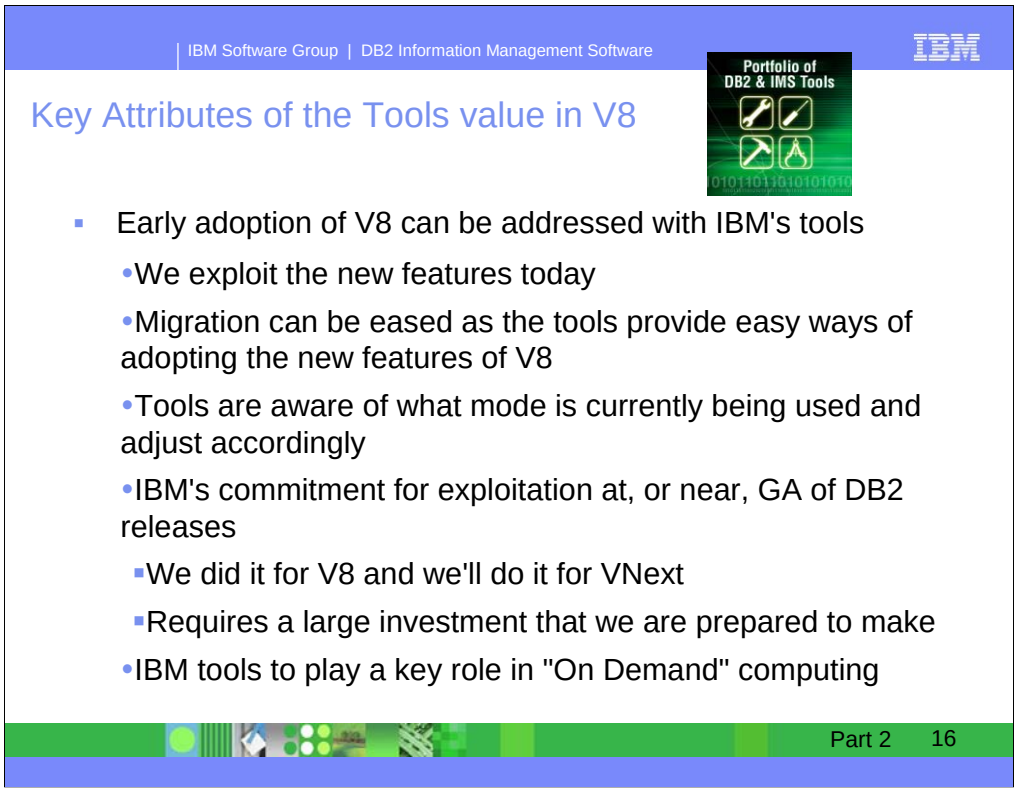

One goal of our support is to enable early adoption of V8. That is why we made the investment for support at, or near, GA. Depending on the tool, up to 70% of the available resource was required to provide this support. Why did we think this important, DB2 V8 has many new, key features. By having the tool support them you are able to make use of them in a much more timely fashion. For example, with our support of MQTs in DB2 Admin you can quickly define, create and view these new objects in a simple straight forward way.

As V8 has 3 modes during the migration, as the Catalog can be a different structure in each one, we've made sure the tools know what phase DB2 is in and adjusts accordingly. Assuring you a smooth transition.

IBM's tools play a key role in on demand computing. We'll talk about this later. Finally, we made this commitment for V8 but, more importantly, we make this commitment for each new release. You can rest assured that the IBM tools will be ready for Vnext.

DB2 V8 Function Supported

- Longer names
  - Tools using new ISPF features
  - Allows for scrollable "columns"
  - Prereqs for tools -- z/OS 1.3 + APAR OW57368, PTF UA02839 (ISPF)
- Sequence objects
- Identity column enhancements
- Unicode
- New parameters
- Materialized Query Tables
- 4096 Partitions

The next series of charts are a subset of all the functions in DB2 V8 that the tools support. I'm not going to go over each of these but, rather, will highlight a few of the more important ones.
Perhaps the most costly and time consuming item we had to support was the longer names. This was due to volume of changes (nearly every screen) All of our tools have ISPF interfaces, displaying a 128 character ID on an 80 character screen presented our first challange. Rather than have each tool "do their own thing" we worked with ISPF to allow a panel to be defined with scrollable "columns". These work just like the scrollable "rows" did. When you have one of these fields you'll see a to indicate it is scrollable. Use PF10/PF11 for scrolling with the cursor on the field you are interested in. Want to see the entire name, no problem, just use the "expand" command and a new window opens with the full 128 name. If you are on z/OS V1.3 or V1.4 you need the ptf listed. V1.5 contains the support in the base.

Other key items on this slide are MQTs, which I talked about earlier, and the >254 Partitions. As with MQTs, you can use our tools to alter your existing tablespaces to take advantage of this new capability.

DB2 V8 Functionality Supported

- Support for VOLATILE tables
- Stored Procedure and User Defined Function Enhancements
- Default for ROWID GENERATED clause
- Support for EXCLUDING clause
- Online Schema Evolution
- Data-Partitioned Secondary Indexes

The only item on this slide I'll talk about is Online Schema Evolution. Change is inevitable. All the tool vendors provide a complex alter capability and a way to manage change. Most of these changes were what I call "distructive alters" because the tool had to unload the data, drop and then (re)create all the objects. And this meant ALL the objects -- tables, views, indexes, authorizations, PLANS, etc. While the DROP was relatively straight forward, recreation of all the dependent objects was time consuming and lead to outages where your data wasn't available. With V8, many of those "distructive alters" can be handled with Online Schema Evolution(OSE). The tools detect if the change you are making can be accomplished by OSE and, if so, use it rather than the destructive drop.

DB2 V8 Functionality Supported

- Support for Plan Table changes
- Support for RUNSTATS REPORT NO UPDATE NONE
- Aliases for Plan Tables
- Distribution Statistics for Non-Indexed Columns
- System Level Point in Time Recovery
- Plus more ....

As with the previous chart, I'll concentrate on only one item -- BACKUP/RESTORE SYSTEM, shown on the slide as System Level Point in Time Recovery. As you know, these new utilities allow you to backup (or restore) a DB2 subsystem to a prior point in time. While we support the generation of these, you can also use one of our tools, Automation Tool, to set up a schedule that you wish to backup the system. This schedule can be based on either some statistical exception you want monitored, by date/time or a combination of the two. Thus you define your criteria for backup and then you're done.

DB2 Tools Product Portfolio 2005

Database Administration
- Administration Tool
- Object Compare
- IMS and DB2 Encryption

Utilities
- Utilities Suite
- High Performance Unload
- Automation Tool

Application Management
- DB2 Bind Manager
- DB2 Data Archive Expert
- DB2 Path Checker
- DB2 Table Editor
- DB2 Test Database Generator
- DB2 Web Query Tool

Performance Management
- Performance Expert
- Omegamon for DB2
- Bufferpool Analyzer
- Performance Manager
- Query Monitor
- SQL Performance Analyzer

Recovery
- IBM Application Recovery Tool for IMS and DB2 Databases
- DB2 Archive Log Accelerator
- DB2 Change Accumulation Tool
- DB2 Log Analysis Tool
- DB2 Object Restore Tool
- Batch Thread Cancel

Replication
- Websphere Information Integrator (II) Replication for z/OS
- Websphere II Event Publisher for DB2 UDB for z/OS
- Websphere II Classic Event Publisher for IMS
- Websphere II Classic Event Publisher for VSAM
- Websphere II Classic Federation for z/OS

Other tools:
IMS, CICS, Applications
Tivoli, Candle, Ascential …

This slide lists all the current DB2 tools available from IBM. I'm not going over each tool today but you can find information about them and other tools from IBMon the web pages.

http://www.ibm.com/software/data/db2imstools/
http://www.ibm.com/software/data/db2imstools/infointeg.html
http://www.ibm.com/software/awdtools/
http://www.ibm.com/software/tivoli/

One point to draw from this chart is the breadth of the tools support. There is a tool to cover almost every need of a DBA, application programmer or system programmer.

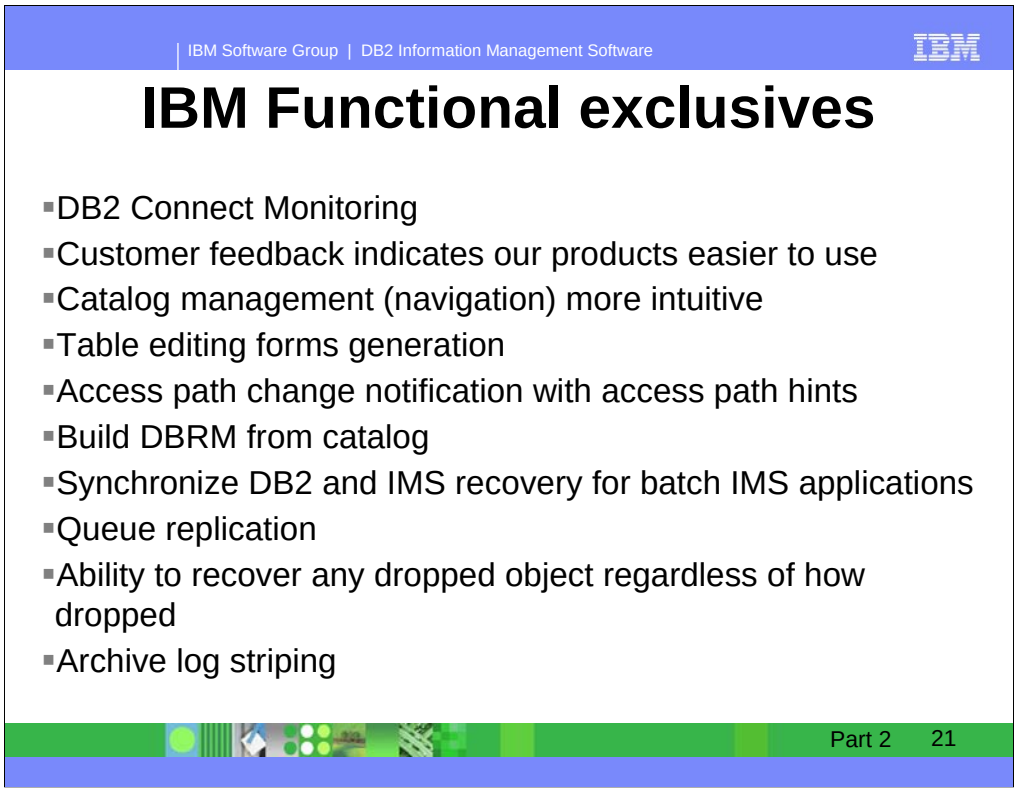

Now let's talk about things that make the IBM tool set unique. First and foremost is DB2 Connect Monitoring. While there are other tools available that do Connect Monitoring none, to my knowledge, are integrated into a system monitor as IBM has done. By integrating it with DB2 PE a user can get a "picture" of performance throughout their environment. That is, the performance is presented to them from the Connect server up to DB2 and back. The next two items are subjective but what our customers are telling us -- our tools are easy and intuitive to use.

The ability to customize an Edit form allows the tool to be useable by more than just the DBAs. In fact, we have one customer that built a custom form and has their call center using it for data entry.

As customers start migrating to V8 they have to be concerned over access path changes due to Optimizer changes. We have a tool that will tell you whether a BIND will cause a change and, if it will, gives an option to have a "hint" built. Along those same lines, many folks have lost the DBRMs created years ago, our tool allows for these DBRMs to be (re)built from the catalog.

The remaining 4 items are all self-explanatory and illustrate additional exclusives in the IBM tools.

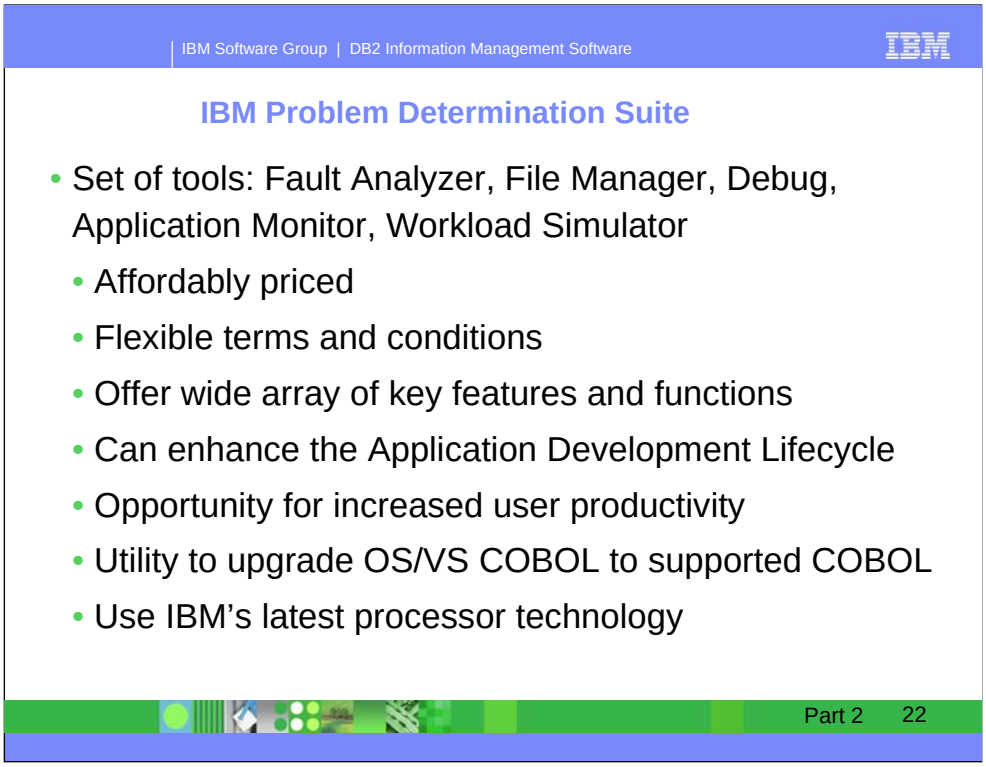

Fault Analyzer for z/OS (FA) Helps you rapidly pinpoint cause of failed application (abends)

File Manager for z/OS (FM) Data management tool supporting key file structures like VSAM, DB2, and IMS

Debug Tool for z/OS (DT)

Debug Tool Utilities & Advanced Functions for z/OS (DTU)

Source code debugging to improve development productivity

Utility to upgrade old OS/VS COBOL to supported levels of COBOL

Application Monitor for z/OS (AM) Helps IT (application programmers) isolate the cause of online and batch application performance bottlenecks with ability to drill down to source

Workload Simulator for OS/390 and z/OS (WS) Application stress and regression testing

For more information on **z/OS Problem Determination and Deployment Tools**

www.ibm.com/software/awdtools/deployment

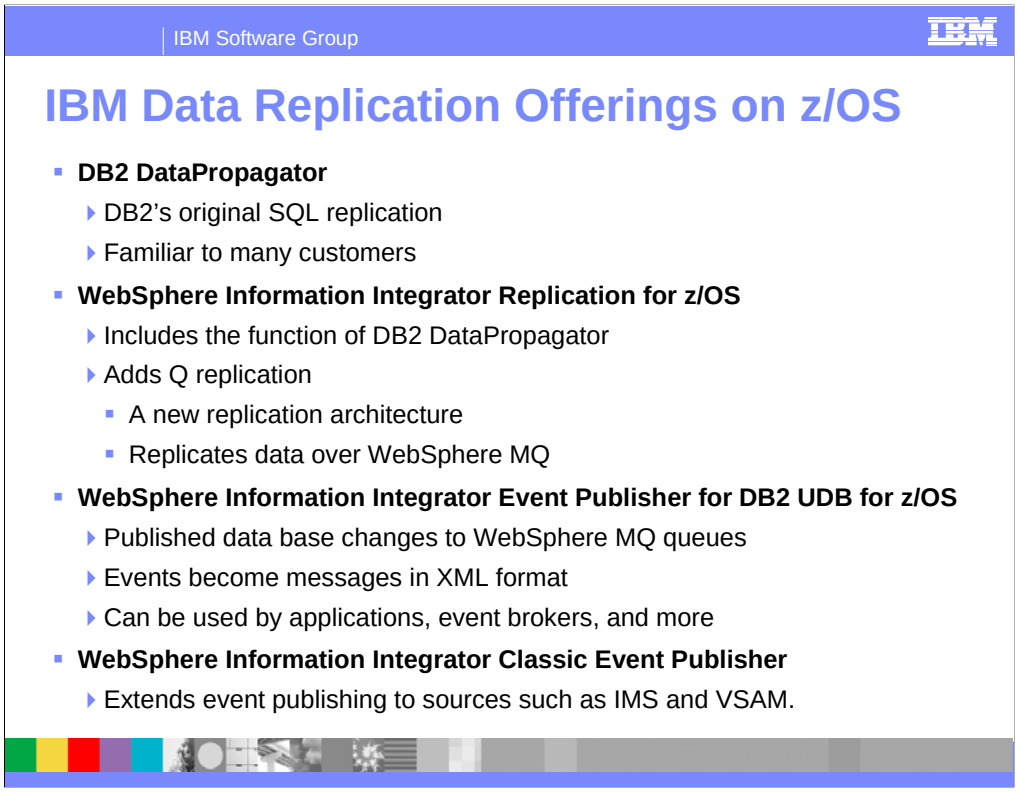

These are IBM's data replication offerings on z/OS. The DB2 replication and event publishing offerings all support DB2 UDB for z/OS V8. The following charts describe each in more detail, along with the types of scenarios where they are used.

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

http://www.ibm.com/software/data/integration/replication.html

http://www.ibm.com/software/data/integration/replication/

http://www.ibm.com/software/data/integration/eventpub.html

Why Replicate?

- **Availability**
  - Scheduled outage, failover, disaster recovery
    - Can use Hardware, Software, or a combination of methods
  - Move query or reporting work to a separate system
    - Other methods such as flash copy also possible
  - Peer to peer - split workload
    - This is only possible through replication
- **Distribution / Consolidation**
  - Move data between central to branches, branches to central, or both
  - Federate or Replicate?
    - where does the application need the data to be? - what db?, what platform?
    - does the data need to be real time or not?
    - what is the change volume?
- **Warehouse / Business Intelligence / Application Integration**
  - Move data to new platform/database, transform data
  - ETL or Replicate?
    - latency needs
    - change volume versus total volume
    - complexity of transformation and/or cleansing

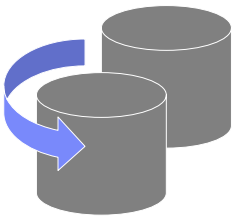

Customers are using replication products to satisfy a wide variety of application needs.

Warehouses and ad-hoc query databases can be built using changed data, real time rather than through less frequent full extract/load processing. So in this case the customer must weigh the relative advantages and disadvantages of traditional ETL processing versus change capture replication.

Data can be accessed in place using federation, but when availability and/or performance of the application is critical, frequently the choice is made to replicate data to a local copy or cache.

The most frequently cited application of replication technology in recent years is availability. Copies of data may be used for scheduled outages, unscheduled outages, disaster recovery, or combinations of these. There are many choices to consider in a high availability scenario- hardware/software, logical/physical, synchronous/asynchronous. Some customers opt for a combination of methods for best coverage.

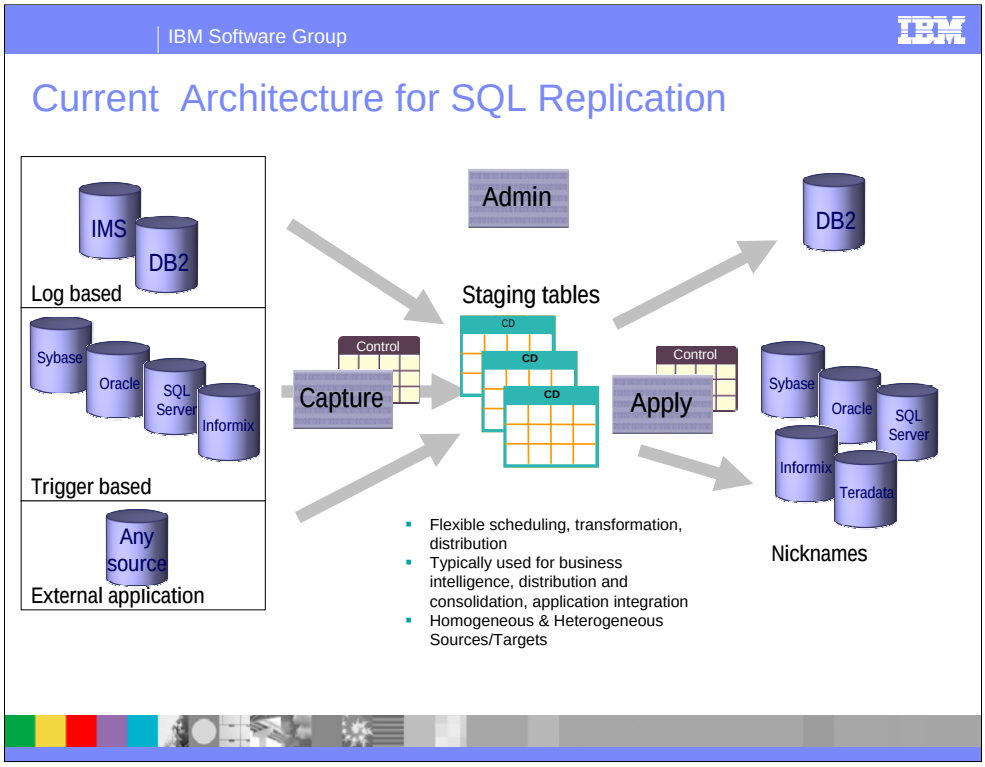

•This is the architecture that has been available for the last 10 years. A Capture program or trigger captures changes and moves them into a staging table, called a changed data (CD) table. A single staging table can serve as source for multiple subscriptions or multiple staging tables can be created for a single source depending on the application requirements. The staging table typically resides on the same system as the source table. Staging table format is published to enable applications or ISV to provide capture function

•The Apply program fetches data from the staging tables using client/server db2 communications and applies it to the target tables using standard SQL statements

- •One or more apply programs can subscribe to a CD table
- •One apply program can replicate data to one or more target tables
- •Target tables can be user copies, history tables, or staging tables
- •Apply program handles column and row subsetting, performs SQL transformations, manages commit scope based on subscription sets and table vs transaction consistent delivery ➔ note that RI cannot be guaranteed for foreign sources as ordering across tables is unknown from trigger capture mechanisms
- •Apply program references foreign source and target tables and control tables via nicknames

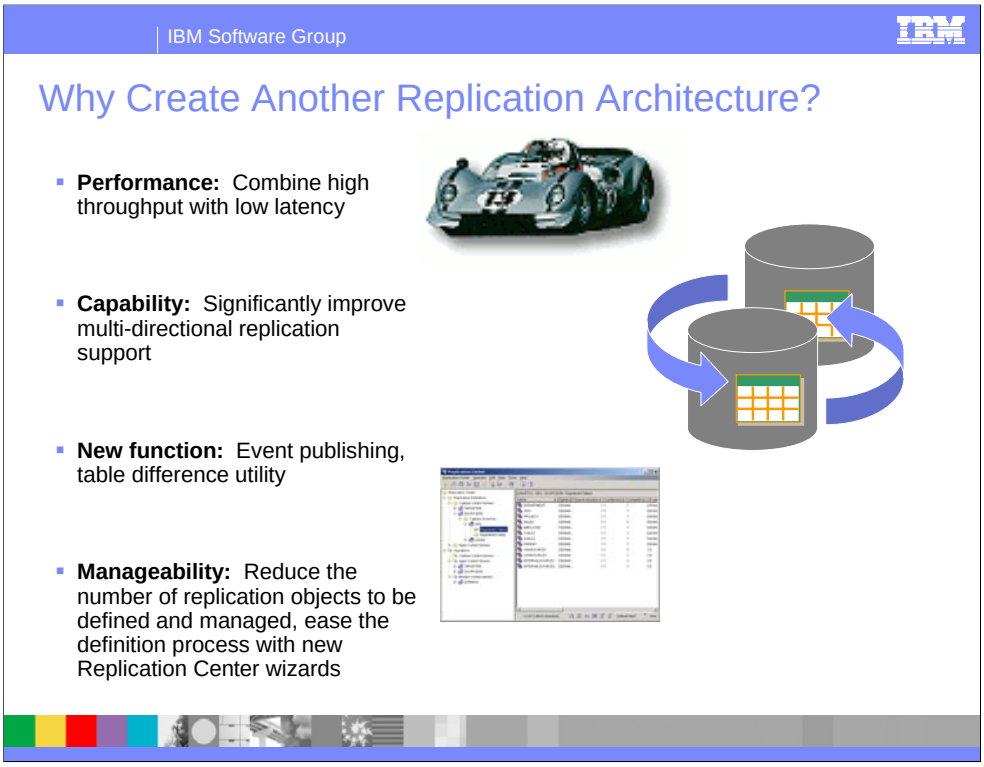

There is a growing demand for high speed low latency replication, primarily for the purposes of meeting high availability requirements. The implementations are varied, and include geographically distributed peer to peer applications, workload balancing, and primary/secondary failover configurations. In addition to speed, these implementations require robust methods for conflict detection and resolution.

We also see the need for a solution that is easy to manage - requiring reduction in the numbers of objects to create and manage, and with easier methods to create and manage those objects.

In creating this new architecture, we also see an opportunity to create an infrastructure that can serve application messaging/publishing in addition to replication.

http://www.ibm.com/developerworks/db2/roadmaps/sqlrepl-roadmap-v8.2.html

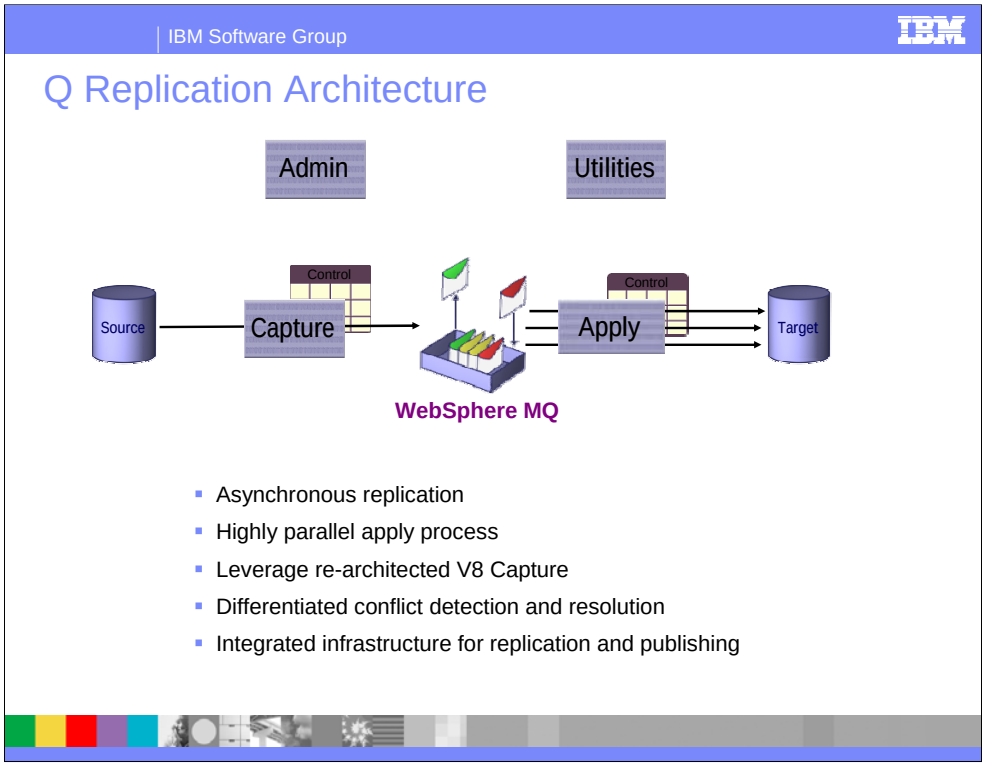

- Capture program stages data in queues
  - Each message represents a transaction
  - One or more data transport queues per source/target database pair
- Apply is significantly re-architected
  - Highly parallel in how it applies transactions to tables
  - Data is always applied per source transaction units
  - Data is applied such that source commit order is observed where necessary for data integrity
- Conflict detection very robust, including ability to handle deletes and key changes
- Data can also be published in XML format for external applications, using the same capture infrastructure
- http://www.ibm.com/developerworks/db2/roadmaps/qrepl-roadmap-v8.2.html
- http://www.redbooks.ibm.com/redpieces/abstracts/sg246487.html

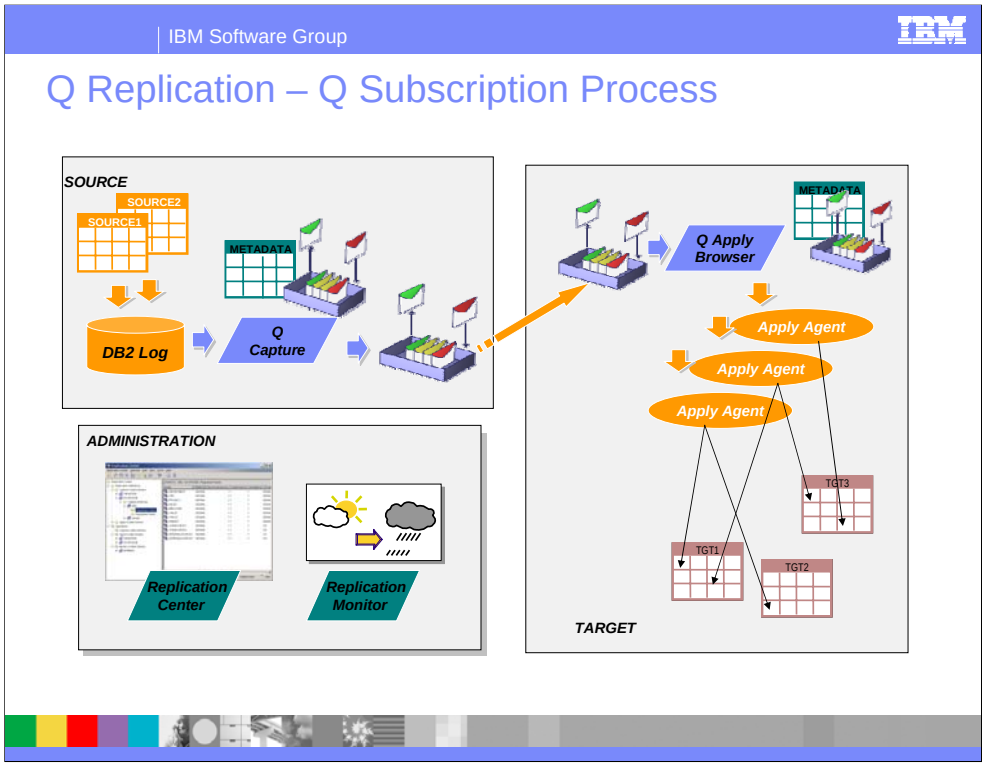

This animation takes you through the implementation details of Q replication.

(1) First you install the programs and set up infrastructure for queues and control table metadata

(2) Subscribe to those tables of interest

(3) Changes to those tables will appear on the DB2 recovery log

(4) The changes will be read by Q Capture and stored in memory

(5) Committed transactional data will be put to the data transport queue

(6) At a commit interval the data will be sent by MQ to the target receive queue

(7) Q Apply browser reads transactions from the queue, examines and tracks dependencies between transactions, and feeds transactions to Apply agents

(8) The alert monitor program keeps an eye on the both SQL and Q Replication server metadata and statistics.

Event Publishing: Why Publish Data?

- **Application to Application Messaging**
  - Drive downstream applications or APIs based on the transactional changed data of database events
  - Reduce application development and maintenance, performance impact to source applications, and availability impact to source applications
- **Meet Auditing Requirements**
  - Capture and store information regarding what changes were made to critical business data and by whom
- **Event Notification**
  - Stream changed data information to Web interfaces
  - Stream only particular events of interest (filter data)
- **Warehouse / Business Intelligence**
  - Integrate captured changed data with an ETL tool
  - Perform very complex transformations
  - Use a specific transaction format to update target

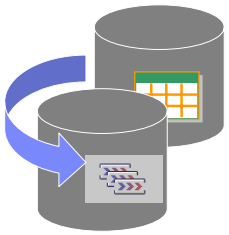

More and more customers are using message queuing to provide application to application communication. When the need exists to combine database activity with application messaging, then a strong advantage can be gained by using an asynchronous log based infrastructure to post messages that coordinate with database events. This eliminates the cost of 2 phase commit, or works in databases that cannot support a 2 phase commit. This also avoids availability concerns posed by an application that would otherwise require both the message queue server and the database to be available. In addition to the performance and availability gains, there is the simplicity of a central publishing mechanism that can be used without any special coding changes or modifications to old or new applications.

Database changes can be posted to a queue and then sent to downstream applications for further processing. Examples: streaming stock prices or wholesale item prices, moving data from an order database to shipping and/or billing databases, notifying all systems of customer information changes….

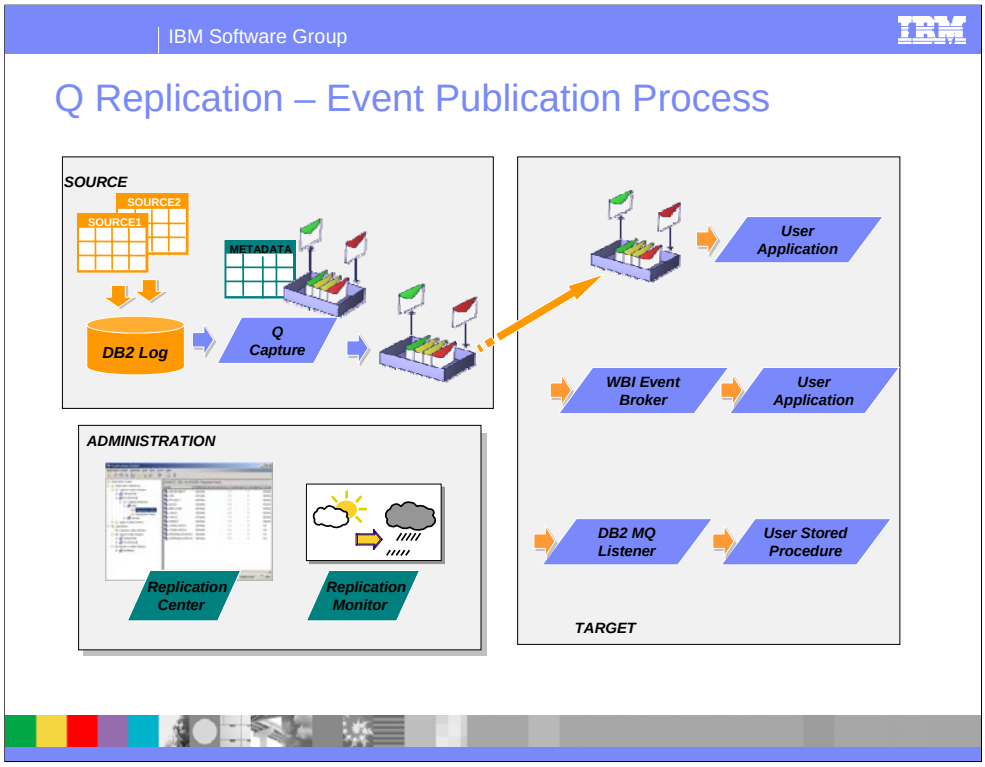

This animation is showing various configuration suggestions for event publishing. In addition to receiving published data directly from a user application, the data could first be brokered by the Websphere Business Integration Event Broker (formerly known as MQSI – MQ Series Integrator), and then passed on to other applications, or the data could be brokered by the MQ Listener function of DB2 on LUW or z/OS, and then passed on to your user written stored procedure.

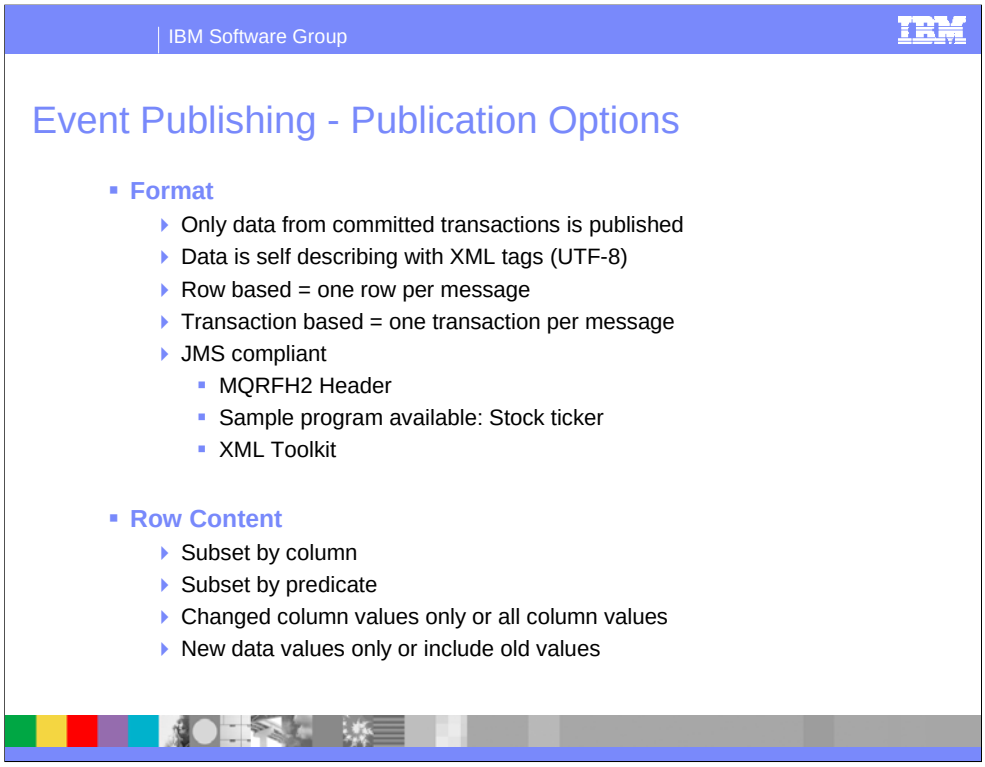

Data is captured in the same way that it is captured in Q Replication – transactional data is stored in memory until a commit record has been seen on the log. Then the committed data is translated into UTF-8, tagged with descriptive XML tags, and is put to a queue.

You can choose for the messages to be made up of individual row changes, or of all associated row changes that were in a transaction.

You can choose whether to have the message JMS compliant or not.

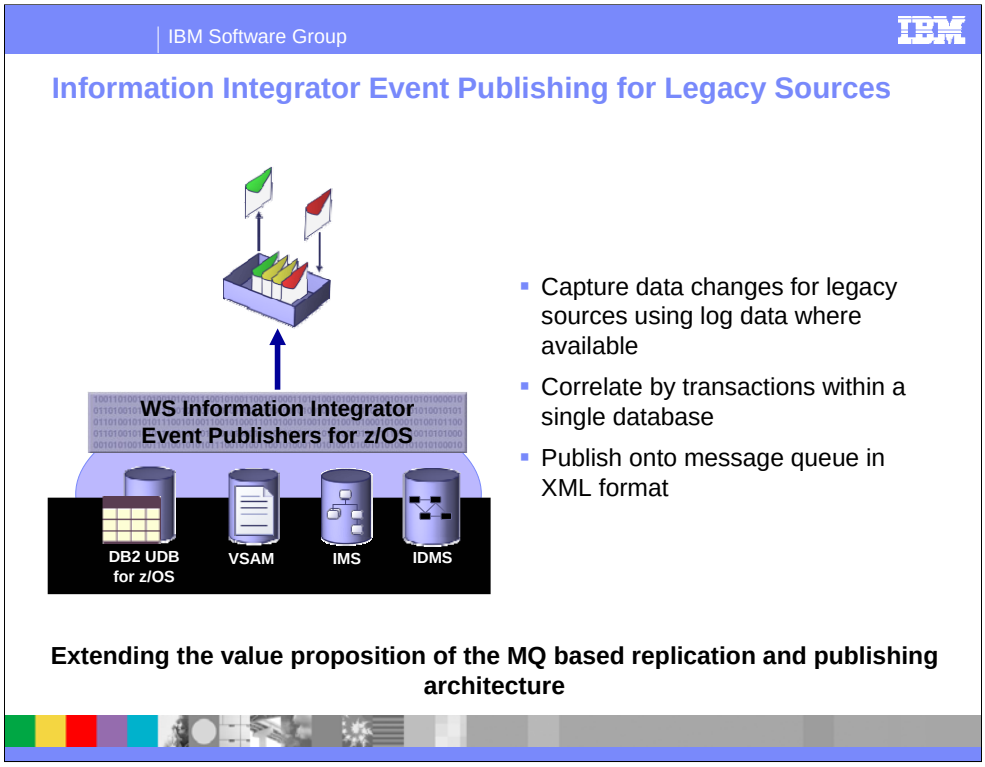

In addition to the event publisher for DB2, we are extending this technology by also offering event publishing from other legacy data sources. The first to be offered will be event publishers for IMS and CICS VSAM. All publishers provide MQ messages in the same UTF 8 XML format. Most options that are offered for DB2 event publishing are also offered for the legacy sources. The major differences are (a) the non relational data must first be mapped to a relational format through a mapping tool, and (b) the legacy capture does not offer filtering at the logical “row” level – this must be performed at the application level.

For legacy data sources, check on DB2 Information Integrator Classic Federation for z/OS.

http://www.ibm.com/software/data/integration/iicf/

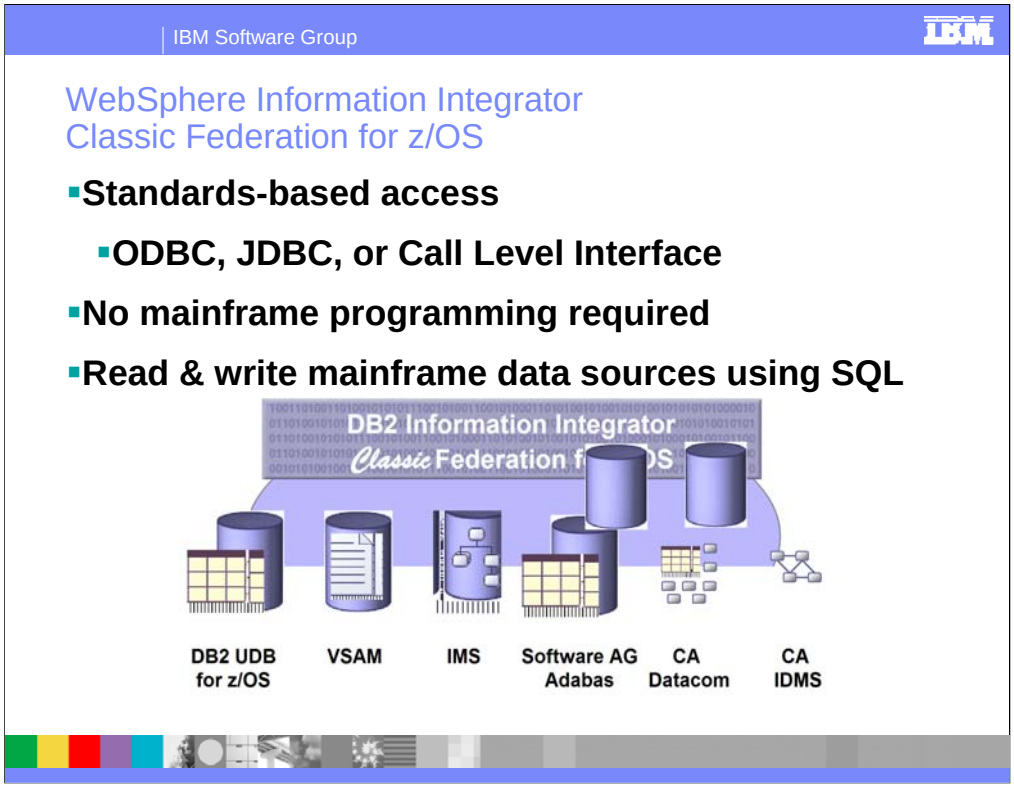

IICF provides standards-based access via ODBC, JDBC, or Call Level Interface. The code is multi-threaded with native drivers for scalable performance. Being Metadata-driven means No mainframe programming required; Fast installation & configuration; and Ease of maintenance.

Classic Federation works with existing and new:

Mainframe infrastructure

Application infrastructure

Toolsets

You can read from and write to mainframe data sources using SQL. This product can help in situations needing to federate the data or in conversions to DB2. We also have a product for VSAM transparency.

http://www.ibm.com/software/data/integration/iicf/

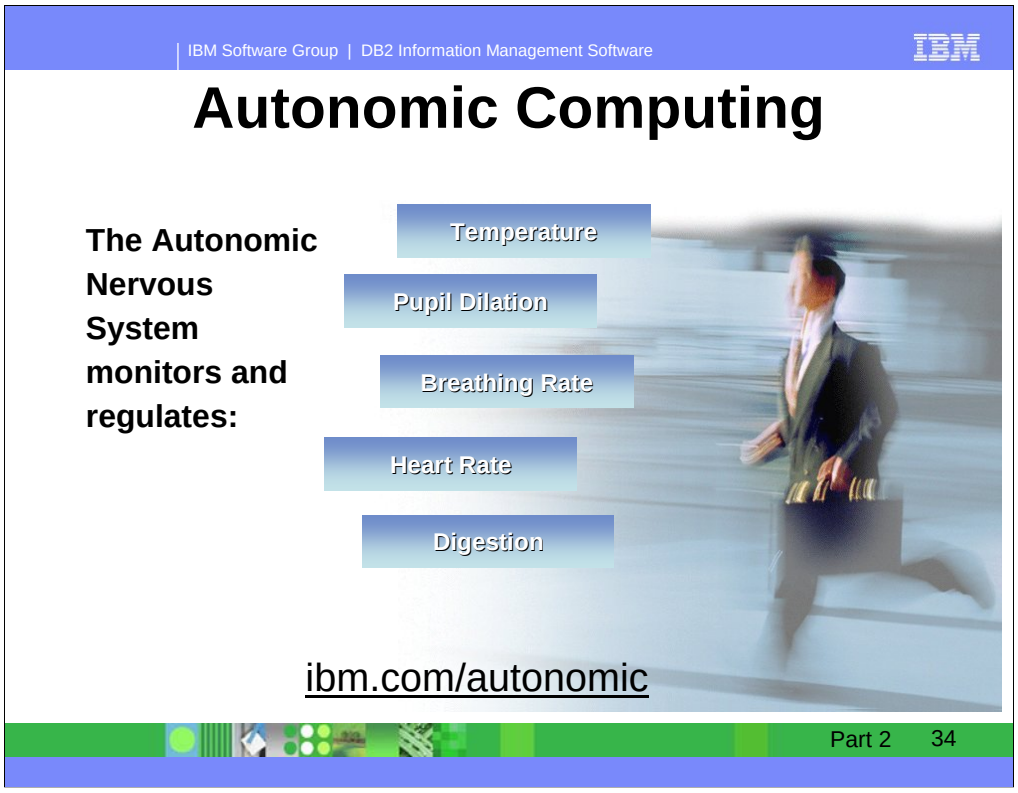

Now let's talk about the future and our role in On Demand processing. There are several different parts to IBM's On Demand strategy. The area with the most payback for the tools is Autonomic Computing.

The source of the term autonomic computing comes from the autonomic nervous system

A system that

Self-manages

Self-heals

Self-protects

Self-optimizes

In other words, a system that takes action based on some predefined exception.

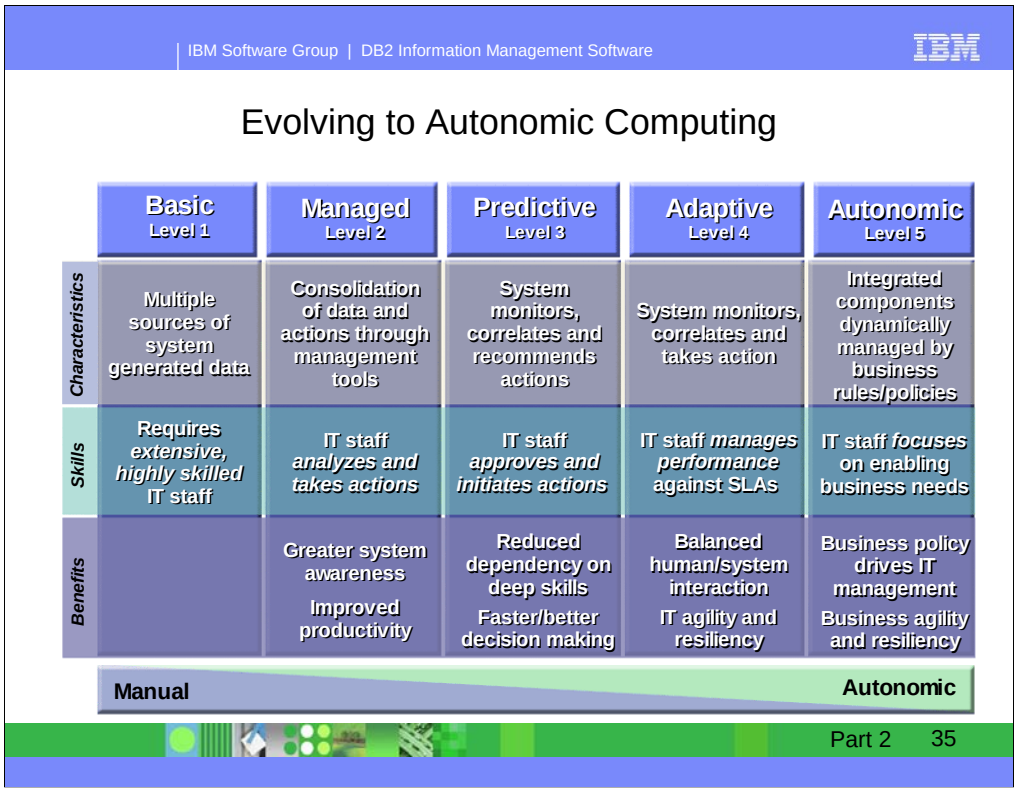

The shift to autonomic computing is an evolution that will take place over a number of years, providing customers with value every step along the way.

After listening to customers talk about how they would like to build on their existing infrastructure, IBM has developed a 5 level deployment model for the evolution of autonomic capabilities.

Different parts of the IT environment can exhibit behaviors at different levels.

Managed capabilities provide system administrative teams with productivity gains through consolidation of data and controls.

Predictive functions provide customers with the ability to make more accurate and faster decisions, having a positive impact on both efficiency and resiliency.

The IT staff becomes responsible for translating business requirements into inputs that control the adaptive decision making process.

Finally, at level 5, The system now has responsibility for decision making and initiation of management actions.

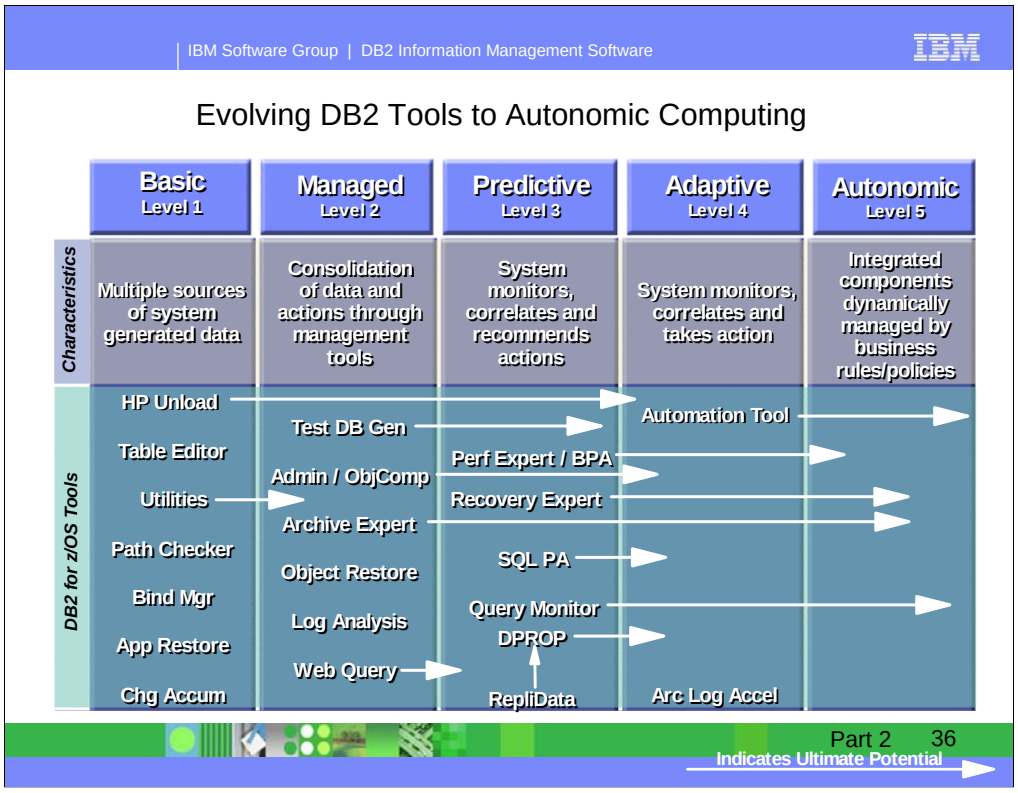

This plot contains an indication of where we feel the tools for DB2 for z/OS fit in the Autonomic levels. Where the name appears is where we believe it currently is.

Autonomic computing is not an all or nothing concept as we see from the levels. The arrow to the right of a tool indicates its utlimate potential in the autonomic computing evolution and represents our plans to move these tools further towards full autonomic computing

However, some tools will not by themselves have the highest ultimate potential. Utilities are a good example of this. They are absolutely vital building blocks to being able to deliver Level 5 tools. For example, the DB2 Automation Tool relies upon the utilities to deliver its ability to perform routine system maintenance in a set-it-and-forget-it style.

It is important to understand the difference between Level 4 and Level 5, so let's use the DB2 Automation Tool as an example. Today, this tool provides the ability to detect a situation needing attention and taking action. That sounds pretty autonomic all by itself, and it is.

However, the next level requires a higher level of self-management. It requires a policy to be specified that will enable you to manage any of your DB2 objects, not by specifying the situation in which you want the action to take lace, but by specifying your business goals. So it's a higher level of management specification.

Summary

- IBM Data Management Tools strategy gains increased focus on autonomic computing
- Our database tools have autonomic characteristics today
- Our database tools will follow the autonomic evolution in an effort to make IT less complex
- Value statements apply:
  - Tools from the database company
  - Comprehensive cross-platform integrated tools with common interfaces
  - Improving customer time to value and total cost of ownership
  - Long term commitment to DB2 and IMS customers

IBM Data Management Tools for DB2 on the web

ibm.com/software/data/db2imstools/

The On Demand strategy is not a new strategy -- it is a new phase in a strategy that began many years ago.

The autonomic computing characterictic of our strategy is not a new concept, but we have an increased focus on it now within Data Management.

Many of our tools exhibit autonomic computing characterictics today, and you will see more improvements in upcoming announcements.

The web site shown contains links to more data on all of our tools -- documentation, analysis reports, customer testimonials, etc.. It will always contain the most current information so I suggest you visit it often.

www.ibm.com/software/data/db2imstools/

Click Support For Compati-bility Matrix & Other Support

→ → →

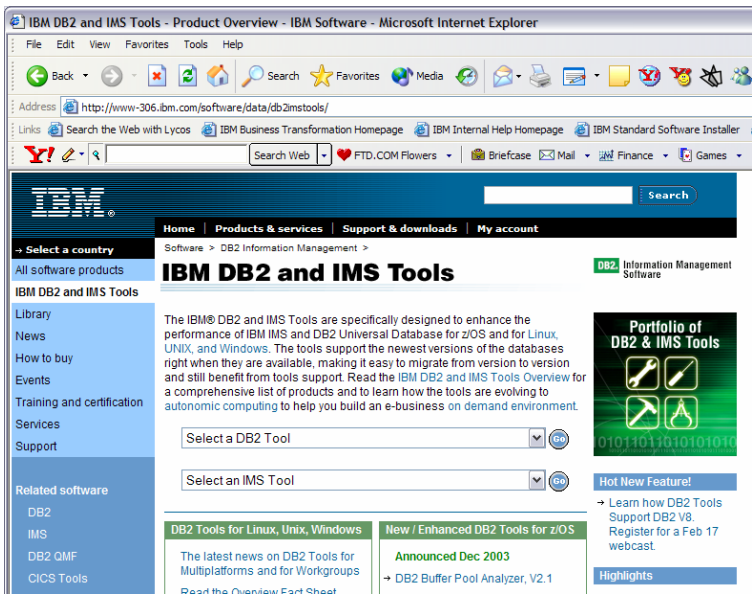

This is the primary web page for DB2 and IMS tools. If you want to know exactly which levels of each tool work with DB2 UDB for z/OS Version 8, then go to the Support page for the tools. Then click Technotes (FAQs) and search for items that include V8 and PTF. Select the item DB2 Data Management Tools and DB2 for z/OS V8.1 Compatibility. This table provides the minimum maintenance required for DB2 Tools to support DB2 for z/OS V8. The Support page has a wide range of other detail about these products.