

5-9 November
Athens Hilton
Athens, Greece

Session: B12

IDUG® 2007
Europe

Why z/OS for your Data Warehouse

Roger Miller
IBM Silicon Valley Lab

07 November 2007 • 04:15 p.m. – 05:15 p.m.

Platform: DB2 for z/OS



This presentation will discuss the wide range of data warehousing and System z. We'll include some recent announcements, architecture and products, then drill down to discuss the capabilities. DB2 began in 1983, with a design and architecture that were able to process both transactions and queries well. At that point DB2 was compared with the much more mature IMS for processing transactions, so the position was that DB2 was the query and reporting database. During the past 24 years, DB2 has improved in both transaction and query processing. Every release of DB2 added new query and optimization improvements, so that DB2 could handle larger amounts of data and more complex queries.

IBM LAUNCHES NEXT GENERATION OF BUSINESS INTELLIGENCE WITH DYNAMIC WAREHOUSING

ARMONK, NY, MARCH 13, 2007 – IBM (NYSE: IBM) today unveiled a comprehensive strategy to enable dynamic warehousing, a new generation of business intelligence capabilities that enable organizations to gain real-time insight and value from their business information. Today's announcement marks an important milestone in IBM's industry-leading pursuit of the global Information on Demand growth opportunity, which is helping customers transform their businesses by using information as a strategic asset.

GoFurther

See this recent announcement for the latest on warehousing from IBM. Note that System z was a key part of the announcement. DB2 for z/OS has a long history in warehousing. Indeed, DB2 V1 was introduced as a decision support option. The terminology has changed a lot, but every release of DB2 brought improvements for handling more complex queries on larger amounts of data.


The distinction between data warehousing and online transaction processing is blurring. Data warehousing and analytic applications are accessing operational or near-real-time data. The traditional strengths of System z and of DB2 for z/OS apply even more to the dynamic warehouse.

IDUG 2007 Europe

Dynamic Warehousing

A New Approach to Leveraging Information

Information On Demand
to Optimize Real-Time
Processes



Dynamic Warehousing

OLAP
to Un
Recom

Query &
to Un
What happened

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

Reporting

GoFurther

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. [click]

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). [click]

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*

GoFurther

DB2 9 provides more functionality in the Data Warehouse area. 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 ideally 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 V9 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>

IDUG® 2007 Europe

More Examples of Dynamic Warehousing in Action

Enabling Information On Demand for Business Advantage

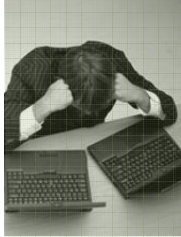
Traditional warehousing	Dynamic warehousing
Insurance fraud analysis and reporting	Identifying potentially fraudulent claims prior to approval and payment Transforms healthcare
Reporting on customer issues	Identifying possible related issues, churn risk and cross-sell opportunities while engaged with the customer Transforms customer service
Historical sales analysis and reporting	Discovering relevant customer information to identify cross sell opportunities and improve negotiating position at the point of sale Transforms sales effectiveness

GoFurther

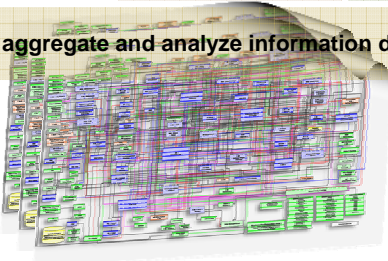
So, what are some other examples of how organizations are transforming their business with dynamic warehousing? [click for transition] Well, as an example, with traditional warehousing, insurance companies analyze the various claims that have been paid out over the past period and report on possible fraud. However, it is typically very difficult to recover any funds that have already been paid out. [click] Dynamic warehousing can transform health care by aggregating relevant information from across the organization and embedding analytics capabilities directly into the claims review process to identify potentially fraudulent claims prior to approval and payment. [click] Customer service can also be transformed by using the information you already have across your organization to identify related issues your customer may have, understanding the likelihood of the customer to leave or close their account, and even determining *appropriate* cross-sell opportunities *while* you are engaged with the customer, whether it be in person, on the phone or over the web. This will enable you to address customer problems more quickly, improve customer satisfaction and turn your customer support efforts into revenue generating opportunities. [click] Similarly, general sales efforts can be dramatically improved if you can better understand relevant information about the specific customer you are working with at the point of sale, instead of just using this information for historical analysis and reporting. This can directly impact profit margins by identifying more relevant cross sell opportunities and improving your negotiating position. [click]

And of course, we talked about how the New York City Police Department is using dynamic warehousing capabilities to more effectively fight crime. Where they used to focus primarily on reporting and analysis of crime statistics, they now aggregate and analyze relevant information as soon as they receive an emergency call reporting a crime. A four-to-six-page dossier, or report, is then generated and sent to the detectives being dispatched, allowing them to identify related incidents and potential suspects before they even arrive at the scene.

Why is it a challenge for organizations to leverage information effectively?

Information distributed in silos across the organization	▶	Not accurate Not complete Not trusted Not timely	
Volume and variety of information increasing			
Velocity of business driving real-time requirements			

Increased need to aggregate and analyze information dynamically



GoFurther

But just having information isn't enough. Unfortunately, the information that users typically have to work with is NOT accurate, NOT complete, NOT trusted and NOT timely. There are several factors driving this, which are preventing most organizations from being able to effectively leverage their information:

[click for transition]

-To begin with, information is scattered across the organization in various silos, as depicted by this real customer architecture diagram;

-In fact, [click for transition] this is only page 1 of 3! The amount of information an organization has to manage is rapidly increasing, and there is a growing variety of information types, such as RFID streams, XML, and unstructured content that must also be managed and analyzed – all of this is leading to extremely complex information infrastructures

-[click for transition] Finally, companies are being forced to innovate and respond to the market much more quickly, driving the need to deliver real time information and analysis

[click for transition] This is all driving the need to aggregate and analyze information more dynamically.

Creates challenges for traditional warehousing

Not just for traditional query and reporting purposes anymore

Warehouses must now:

- Address expanding needs for **analytics** and information on demand
- Leverage ALL types of information, including **unstructured**
- Serve **increasing numbers and types of applications** and users, with varying service level demands



Increasingly mixed workload environments
and the constantly changing needs of different business constituents
require more dynamic warehousing capabilities

GoFurther

Combined with the need to start leveraging information more effectively, this creates significant challenges for traditional warehousing approaches; warehouses are needed not just for traditional query and reporting purposes anymore.

There is a growing community of users that need the ability to do ad-hoc, exploratory analysis as organizations try to make business intelligence available to the masses. As we already discussed, there are now various operational applications that require real time business insights, necessitating analytics that can be embedded into business processes. And these applications must be able to leverage ALL types of information, including unstructured content, such call center logs, technician notes, contracts, call logs, etc. So warehouses must now address this expanding need organizations have for business insight and serve the increasing number of these different application types.

As a result, warehouses must also support the varying service level demands of these different applications. The combination of mission critical operational applications that require real-time responsiveness and traditional back-office reporting and analysis for strategic and tactical planning purposes is leading to increasingly mixed workload environments. And this is further complicated by rising data volumes, continuously expanding amounts of history and the growing number of users, which is causing requests for information to become more numerous and sophisticated. [click for transition]

These increasingly mixed workload environments, along with the challenge of addressing the constantly changing needs of the different business constituents that must be supported across the organization, requires more dynamic and balanced warehousing capabilities.


IDUG 2007 Europe


IBM provides more than just a warehouse

DB2 Warehouse provides extended capabilities and value

**Embeddable analytics
(Inline and as a Service)**

Multidimensional analysis
Data mining and visualization

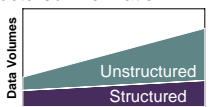




**Traditional
warehouse**

**Beyond traditional
structured data**

Generate and leverage
knowledge from
unstructured information



“As a direct effect of the mixed workload, with continuous loading and the increase in automated transactions from the functional analytics in OLTP, the transactional DBMSs have an edge that challenges the DW DBMSs (such as Teradata)”

Gartner Data Warehouse Magic Quadrant, 2006

**Benefits of a
data server**

Optimized for
High availability and reliability
Scalable, secure and auditable

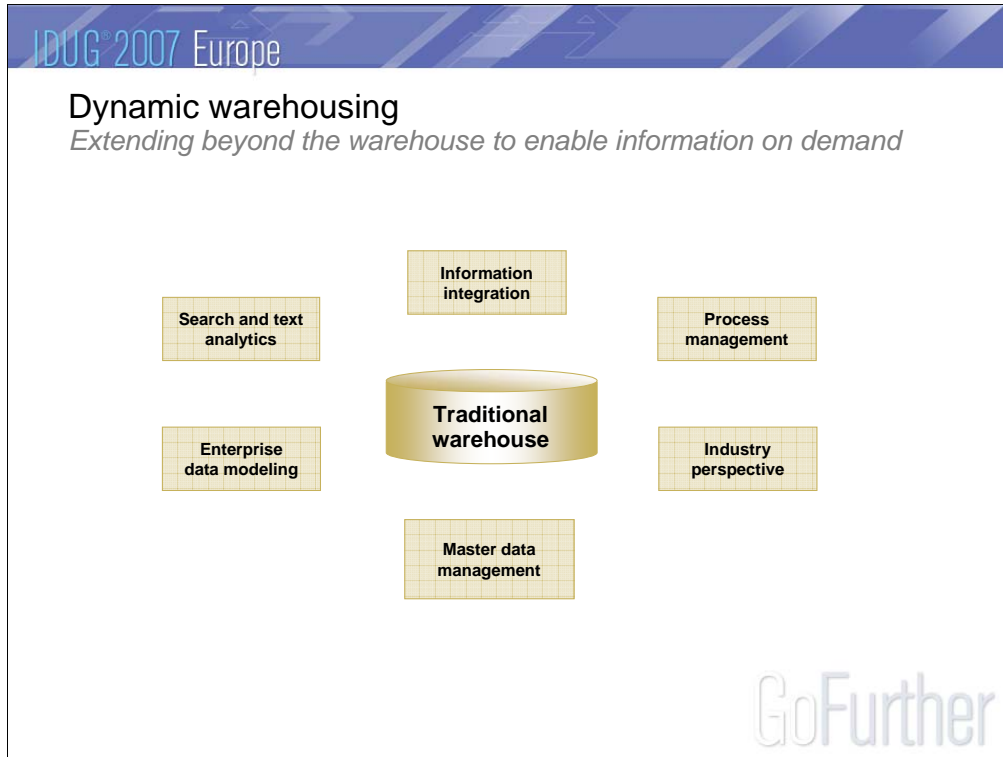
Reduced storage costs
Better disk utilization
Query speed improvement

**ed
sing**

Advanced parallelism
Advanced data partitioning
Workload management

GoFurther

Enabling dynamic warehousing requires a platform that goes beyond traditional warehousing capabilities. This is where IBM is differentiated — with our ability to provide more than just a warehouse. [click] As Gartner points out in its recent Data Warehouse Magic Quadrant report, as a result of these mixed workload environments, with continuous loading and increased “functional” analytics requirements in OLTP based applications, warehousing solutions based on transactional databases will have an edge over pure data warehousing databases such as Teradata’s. This is one of the areas where IBM’s DB2 family based warehouse solution is uniquely positioned. We have a “best of both worlds” architecture with the benefits of a transactional, OLTP oriented data server foundation, which has been optimized for real time, secure, auditable access, along with mission-critical reliability and high availability, while at the same time, providing dedicated warehousing capabilities, advanced data partitioning and workload management capabilities that when combined, provide a multi-use platform that can scale linearly to meet the demands of any organization. Scalable and dedicated warehousing capabilities that you can grow horizontally are available through either our DPF feature using a “Shared Nothing” architecture or Sysplex DataSharing using a “shared everything” architecture with hardware assist. On top of this, the DB2 Warehouse platform provides advanced compression capabilities that can dramatically reduce storage costs, while enabling better disk utilization and improved query speeds. [click] In addition, we actually provide analytics capabilities embedded within the warehouse. These analytics can be delivered inline within applications, or made available as a service. And this includes multidimensional analysis, data mining and visualization capabilities that are often required to optimize business processes and assist with decision making. We also provide support for better leveraging unstructured content, which is increasingly becoming the largest portion of enterprise information, yet remains extremely under-utilized. Not only do we enable unstructured content to be associated with business information in the warehouse, but we also now enable organizations to extract additional knowledge from that unstructured content, which can be used for deeper analysis and a better understanding of the data. [click] All of these things are part of the DB2 Warehouse, providing customers with extended value and a better platform upon which to build solutions that can deliver more dynamic business insights.



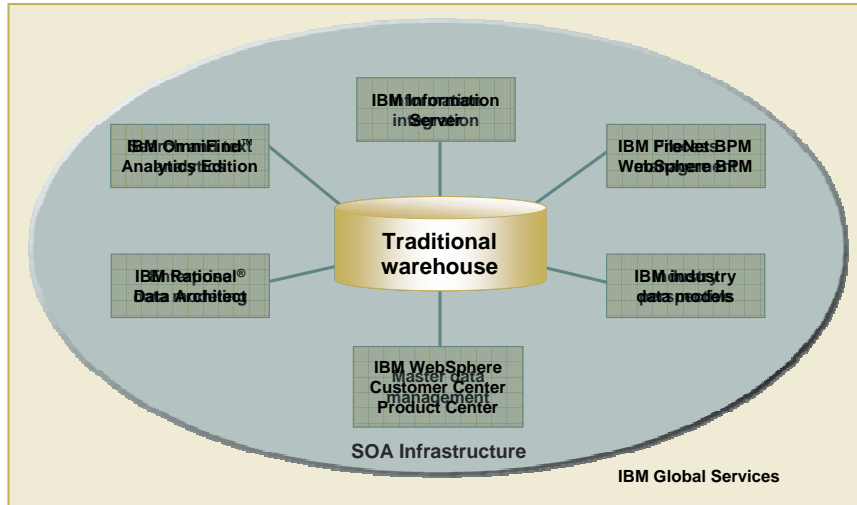
So, what is needed to achieve Dynamic Warehousing? Well, Dynamic Warehousing is about more than just the warehouse. It is about providing a set of services that extend beyond traditional data warehousing and reporting to support the increasing number of business processes and applications requiring analytic capabilities and address the demands for more dynamic business insight. It requires: [click] *Search and text analytics* capabilities to enable usable knowledge to be extracted from unstructured information *Information integration* capabilities to aggregate, cleanse and transform information from disparate sources across the enterprise and make that information available as a service *Process management* capabilities that can leverage analytics for improved decision making and process optimization by delivering business insights within the context of the activities being performed *Enterprise data modeling* to provide common metadata for working with all relevant information

- *Master data management* to ensure a common view of customers, partners and products across different applications, provide clean and authoritative dimensional data to the warehouse, and enable a single version of the truth *Industry perspective* to more effectively apply analytics to a particular domain.

Dynamic warehousing enables you to integrate, transform, analyze and harvest insights from structured and unstructured information that can be delivered in context for real-time operational business processes, in addition to the more traditional historical analysis required for strategic and tactical planning efforts. Making information available as a service enables all relevant information to be delivered to people and processes in the context of individual business activities, and allows you to ensure that only accurate, trusted information is being used to make business decisions. This approach can also be used to enhance the value of Data Warehousing investments. Many analytic systems today are limited in their reach. By extending the use of insightful information and analytics to a much broader community of users, in context of the activities they are performing, you can increase the value of investments being made to generate business intelligence. This is part of the road to dynamic warehousing.

How IBM Enables Dynamic Warehousing

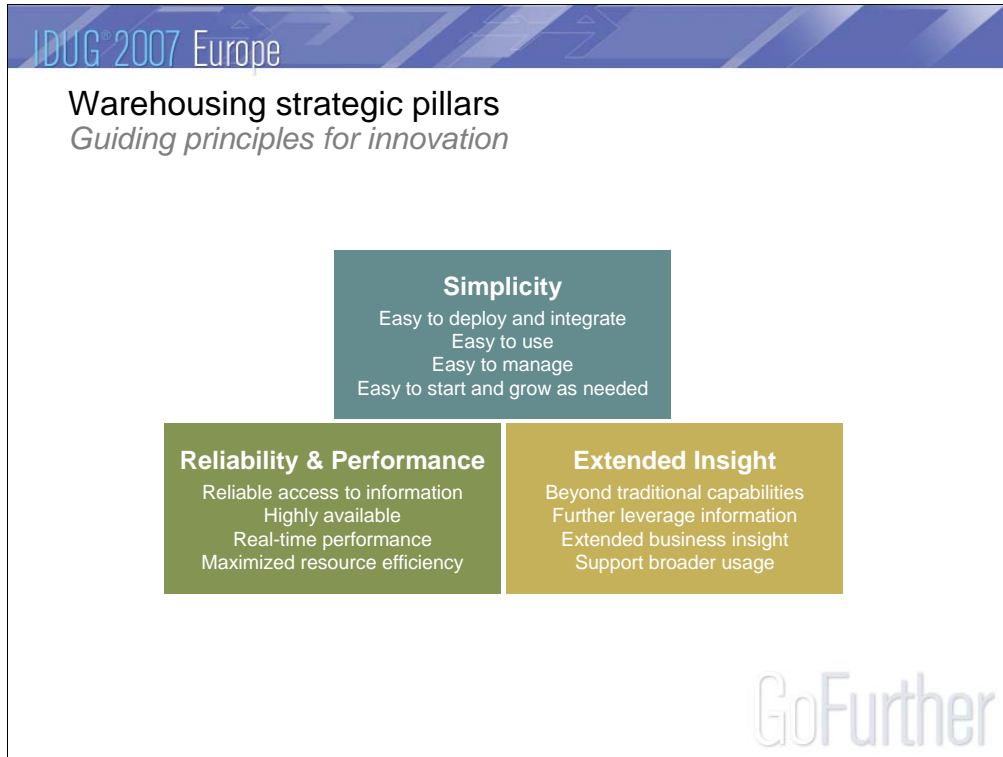
Integrated offerings to enable information on demand



GoFurther

So, IBM's DB2 Warehouse is at the heart of enabling Dynamic Warehousing. And as we reviewed on the prior slides, IBM's warehousing solution, DB2 Warehouse, includes a unique set of features and capabilities that make it better equipped to support the expanding needs for analytics and dynamic business insight, which must be incorporated into an increasing number of business applications with very diverse requirements. This includes: *Analytics capabilities embedded in the warehouse*, which can be delivered inline with business applications or made available as a service; *Integrated processing and transformation of unstructured information*, enabling a greater amount of knowledge to be extracted from what is the fastest growing set of information across the enterprise; and the combination of an OLTP-based *transactional data server foundation* and *dedicated warehousing*, providing a "best of both worlds architecture" that enables better handling of continuous loading and the increase in automated transactions from operational oriented analytics, while at the same time enabling better scaling and workload management capabilities. But this is just the start. IBM can deliver more effective Dynamic Warehousing capabilities than any other vendor because of the availability of offerings that address *all* of the critical services. This includes OmniFind Analytics Edition for search and text analysis, IBM Information Server for information integration, FileNet and WebSphere BPM offerings to address process management requirements, Rational Data Architect for enterprise data modeling, WebSphere Customer Center and Product Center for complete master data management capabilities, IBM Industry Data Models to provide out-of-the-box domain specificity, and industry leading SOA infrastructure that enables information to be delivered to people and processes more effectively.

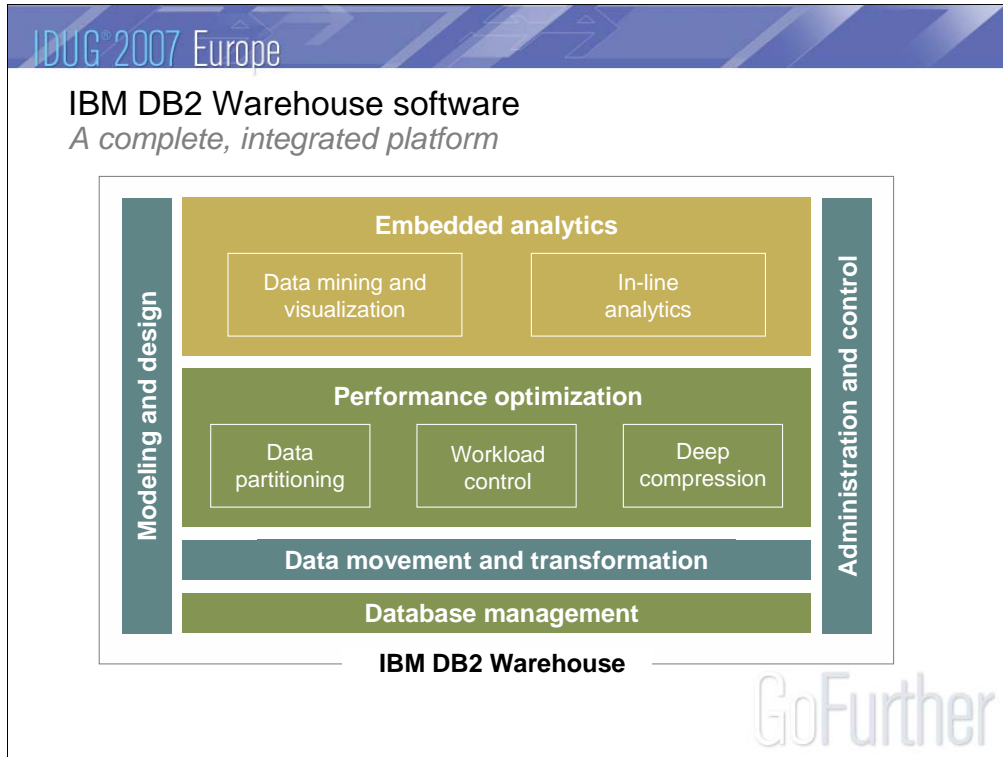
Making information available as a service enables all relevant information to be delivered to people and processes in the context of individual business activities, and allows you to ensure that only accurate, trusted information is being used to make business decisions. This approach can also be used to enhance the value of Data Warehousing investments. Many analytic systems today are limited in their reach. By extending the use of insightful information and analytics to a much broader community of users, in context of the activities they are performing, you can increase the value of investments being made to generate business intelligence. This is part of the road to dynamic warehousing. [click for transition] Not only does IBM have all of these relevant offerings, but we have begun tightly integrating them with the warehouse for more seamless delivery of dynamic warehousing. [click] And of course, IBM is also now offering a set of services to help organizations leverage industry best practices and achieve faster implementation times.



So, before we jump into the different components of IBM DB2 Warehouse, it's important to understand the guiding principles that drive what goes into our warehousing solutions. IBM has a set of core values, such as "Dedication to every client's success," which drive our actions as an organization. Likewise, we have identified a set of strategic pillars that we use as guiding principles when deciding what features to add to our warehousing offerings and how to innovate. These revolve around:

- *Simplicity*—to make our solutions as easy to use as possible;
- *Reliability and performance*—to provide efficient, reliable, highly available access to information for historical analysis OR operational purposes, with real time performance requirements, while maximizing resource utilization; and
- *Extended insight*—going beyond traditional warehousing capabilities to enable our customers to get more value out of their information by providing built in capabilities for generating greater business insights and supporting broader usage of the warehouse across the organization.

Every feature we add and each new offering we create is driven by one or more of these guiding principles.



Our goal with the DB2 Warehouse platform is to provide a complete set of tightly integrated services that meet the objectives of our guiding principles. Customers can leverage all of these services together or selectively implement desired options for ultimate flexibility.

It starts out with DB2 9 for LUW and z/OS, a highly scalable, enterprise class database that can address the performance characteristics of any type of application. We also provide out-of-the-box data movement and transformation capabilities to reduce the complexity and lower the costs typically associated with loading data into the warehouse and preparing that data so that it can be leveraged more effectively. We then add a set of performance optimization features that enable the warehouse to address broader enterprise requirements. This includes Database Partitioning or Sysplex DataSharing for dedicated warehousing that can scale linearly, Workload Control features for prioritizing queries to ensure that the most critical applications are serviced accordingly, and Deep Compression to increase efficiency of the warehouse and reduce storage costs.

We then extend the value of the warehouse by including analytics capabilities that can be delivered inline within applications or made available as a service, along with data mining and visualization to provide more dynamic business insight. These services are embedded in the warehouse to provide better performance, increased efficiencies and reduced costs. Finally, we provide a common set of integrated tools for data modeling and design, administration and control of the warehouse and all related services, making them easier to use and manage.

IDUG® 2007 Europe

Simplicity Extended Insight

Deliver inline visualization and analytics

Embedded analytics capabilities

Can be embedded directly into applications and Web pages

Out-of-the-box visualization tools

GoFurther

And DB2 Warehouse provides out of the box visualization tools for displaying the results of mining and analytic operations...

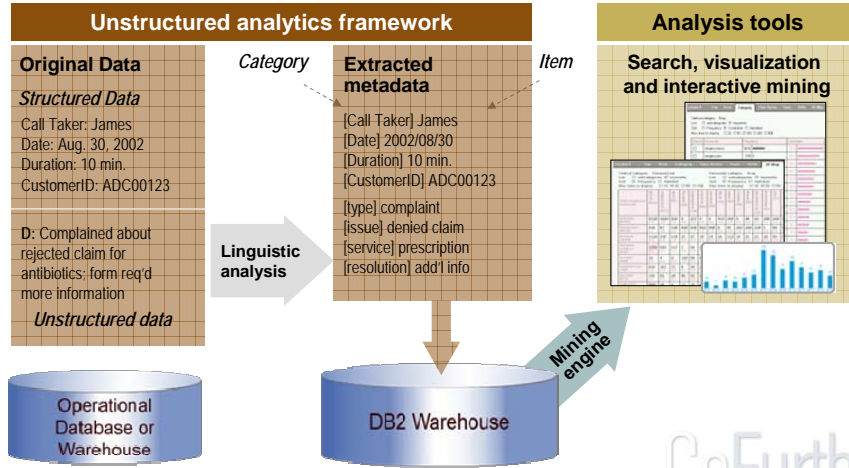
[click for transition]

...which can be embedded directly into your business applications or web pages. These can be simple charts or graphs to depict specific scorecard or dashboard type of information, such as revenue comparisons, portfolio mixes, or demographics, or can be highly interactive charts that allow users to drill through and further analyze the information.

These capabilities are not meant to replace your standard BI reporting tools, but to supplement them. The value of our embedded visualization and analytic capabilities is that they are made available as individual components that can be delivered inline directly within other applications, without requiring users to pull up a separate BI tool that takes them out of the context of their current activities and typically requires greater training for use.

Introducing IBM OmniFind Analytics Edition

- Rich analysis interface for combining structured and unstructured data
- Combines search, text analytics and data visualization



GoFurther

IBM is introducing a new solution to the market, OmniFind Analytics Edition, to provide a rich interface for extracting business insights from inter-related structured and unstructured information by combining a set of search, text analytics and visualization capabilities.

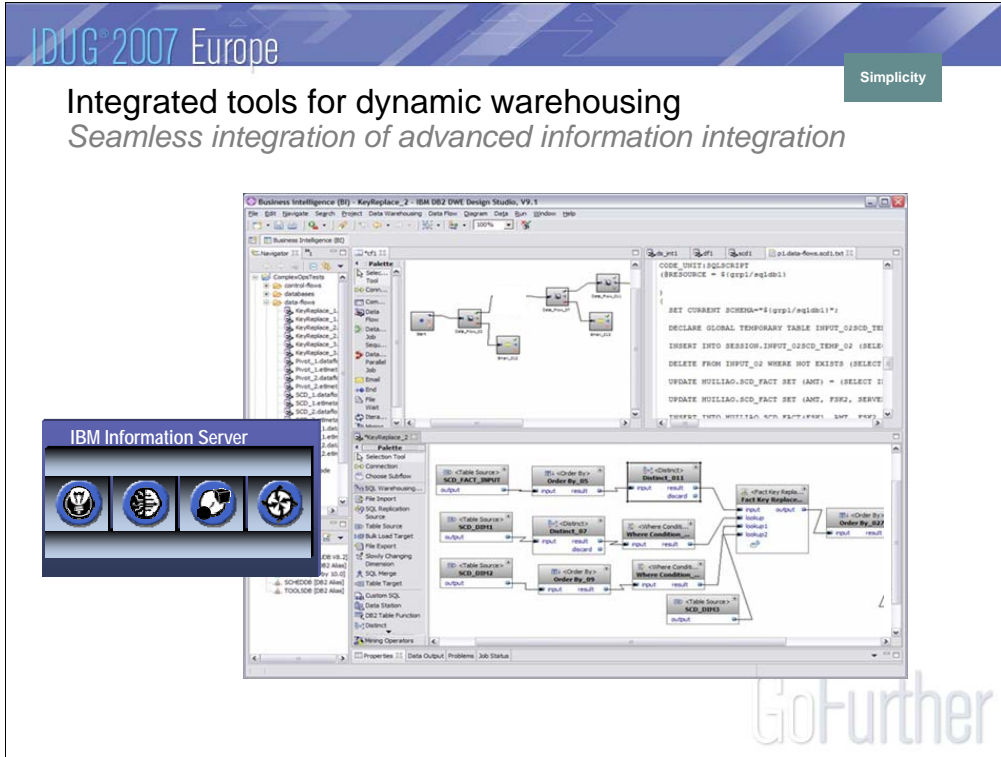
[click for transition]

An unstructured analytics framework, based on UIMA, provides linguistic analysis to interpret free form text and generate more useful metadata. In this example, we are showing how you could extract from call center notes the type of request, the type of service needed and the product components referenced. This information can then be sent to DB2 Warehouse to enable additional mining and reporting.

[click for transition]

OmniFind Analytics Edition also provides search, visualization and interactive mining tools to enable users to extract business insights through timeline analysis, topic extraction, correlation analysis, trend analysis, and distribution analysis.

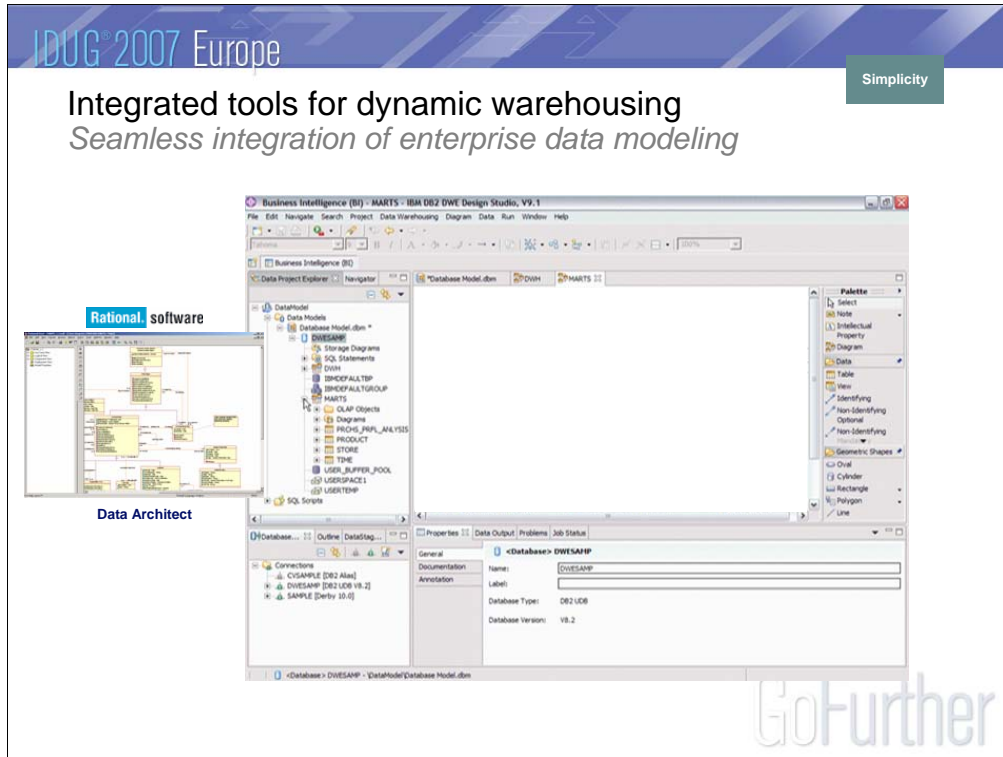
Integrated tools for dynamic warehousing
Seamless integration of advanced information integration



Not only does IBM have offerings to address all of the requirements for dynamic warehousing, but we have also tightly integrated those capabilities into our DB2 Warehouse platform. For example, we are providing tight integration with the IBM Information Server so that DB2 Warehouse customers can take better advantage of the advanced data quality, aggregation and transformation capabilities.

[click for transition]

This enables customers to seamlessly embed an Information Server process directly within a warehouse processing pipeline using the DB2 Warehouse design and administration tool.



Another key requirement we discussed earlier was the need for enterprise data modeling capabilities to establish common metadata for analysis and insight across business areas. IBM Rational Data Architect provides organizations with a solution that can address these needs. In our new release of DB2 Warehouse, not only do we enable customers to import data models created in Rational Data Architect...

[click for transition]

...but we are now embedding these data modeling capabilities directly into the DB2 Warehouse design and administration tool so that customers can design and implement their enterprise data models natively within the warehouse for simplified management and administration.

IBM Information Server for System z

Benefits of this architecture

Significant cost savings

- z/OS MIPs consumption dramatically reduced vs. USS or MVS approaches
Minimizes impact on other z/OS software costs
- All Job Processing is on zLinux (except the z/OS data access)
MIPs charged at IFL rate ... NOT z/OS rate
- DB2 workload on z/OS can qualify for offload to zIIP specialty engines

High performance z data connectivity

- Batch Pipes for DB2 load, DRDA to DB2 over hipersockets
- SQL to Classic over hipersockets
- Integration with MQ and therefore with the Data Event Publishers

Seamless integration with other IBM Information Server platforms

- Same operational architecture and metadata Repository
- Eliminates deployment issues
- Maintains value of DataStage for z/OS investments

GoFurther

IDUG® 2007 Europe

Industry data models
Leverage industry best practices for faster time to market

Over 400 Customers! Extended Insight

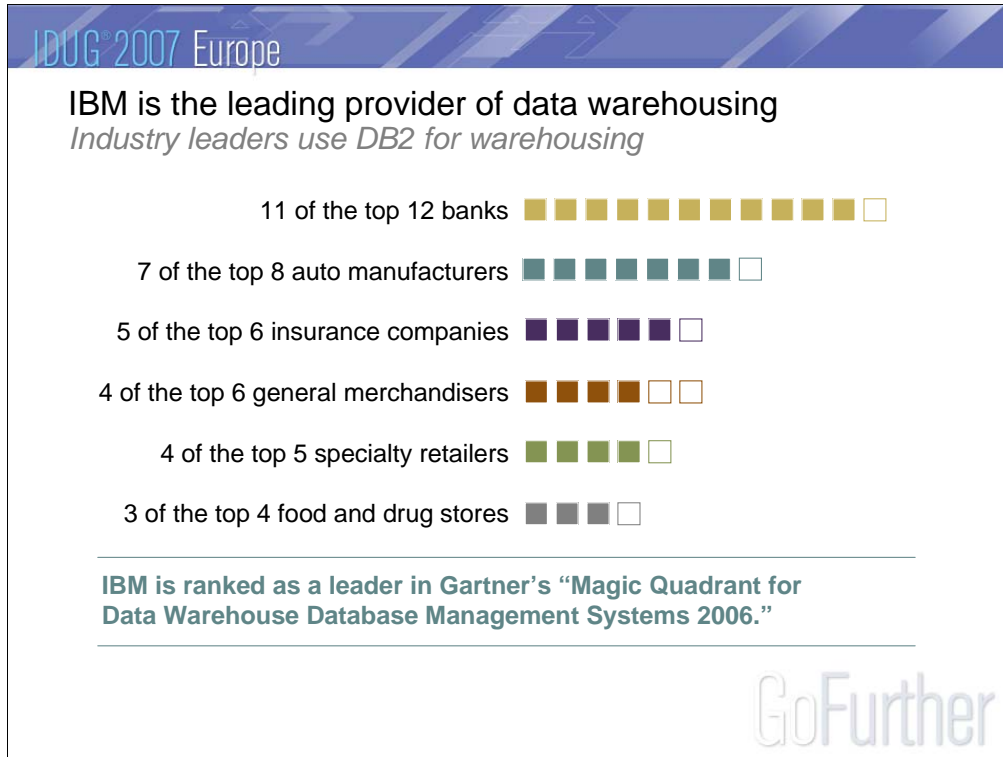
<p>Banking (Banking Data Warehouse)</p> <ul style="list-style-type: none"> ▪ Profitability ▪ Relationship marketing ▪ Risk management ▪ Asset and liability management ▪ Compliance 	<p>Financial Markets (Financial Markets Data Warehouse)</p> <ul style="list-style-type: none"> ▪ Risk management ▪ Asset and liability management ▪ Compliance 	<p>Health Plan (Health Plan Data Warehouse)</p> <ul style="list-style-type: none"> ▪ Claims ▪ Medical management ▪ Provider and network ▪ Sales, marketing and membership ▪ Financials <p>New Offering!</p>
<p>Insurance (Insurance Information Warehouse)</p> <ul style="list-style-type: none"> ▪ Customer centricity ▪ Claims ▪ Intermediary performance ▪ Compliance ▪ Risk management <p>Enhanced Capabilities!</p>	<p>Retail (Retail Data Warehouse)</p> <ul style="list-style-type: none"> ▪ Customer centricity ▪ Merchandising management ▪ Store operations and product management ▪ Supply chain management ▪ Compliance 	<p>Telco (Telecommunications Data Warehouse)</p> <ul style="list-style-type: none"> ▪ Churn management ▪ Relationship management and segmentation ▪ Sales and marketing ▪ Service quality and product lifecycle ▪ Usage profile

GoFurther

IBM also has several prebuilt industry data models to provide a starting point for addressing specific business areas. These include predefined vocabularies of business concepts and terminology specific to the industry, project views for particular business problems, like risk management, and business solution templates with key performance indicators for dashboards, scorecards or other reporting applications.

And we are actually bringing to market a new Health Plan data model that can help customers predict costs and analyze the impact of changes in the mix of disease incidence, medical services and demographics, while providing more accurate reporting of medical costs, and better understanding the claims handling cycle so that they can optimize their processes...

...and we are making available an enhanced Insurance data model that will help organizations improve efficiencies in claims and underwriting processes, reduce time to market for new products, and help address risk and compliance.



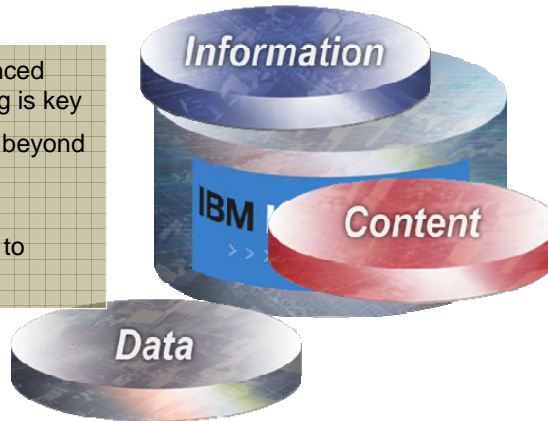
While IBM only recently introduced warehouse specific packaging with DB2 Data Warehouse Edition, the fact is that companies had already been using all of the different components, which were previously offered as piece parts on top of DB2, to address their warehousing needs.

In fact, IBM is actually one of the leading providers of warehousing solutions in the market. Over 3,000 companies are already using IBM DB2 for their warehousing needs, on distributed systems AND on System z – the mainframe. This includes the industry leaders in banking, auto manufacturing, insurance and retail. An impressive list of organizations.

IBM enables dynamic warehousing

Delivering greater value from information

- More dynamic and balanced approach to warehousing is key
- Broad set of capabilities beyond the warehouse required
- IBM provides the most comprehensive platform to address these needs



GoFurther

So as you can see, IBM is offering a broad set of products and capabilities to enable organizations to generate increased business insights and achieve more dynamic business optimization. A more dynamic and balanced approach to warehousing is key, but this also requires a broader set of capabilities beyond the warehouse.

[click for transition]

And IBM provides the most comprehensive portfolio of offerings to address these needs and help organizations deliver more business insight and greater value from their information.

How to benefit from a consolidation and/or modernization of your Warehouses and Data Marts to achieve near real-time requirements, single enterprise wide "version of the truth", best of industry security and regulatory compliance

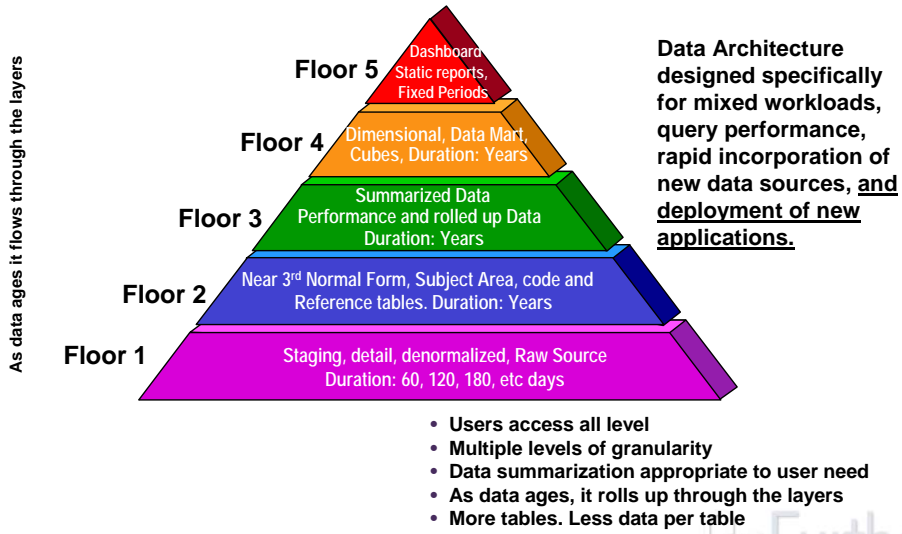
**DB2 for z/OS Enterprise Data Warehouse
using Layered Data Warehouse for
IBM Dynamic Warehousing**

GoFurther

The next section discusses how to benefit from a consolidation and/or modernization of your warehouses and Data Marts to achieve near real-time requirements, single enterprise wide "version of the truth", best of industry security and regulatory compliance.

The DB2 for z/OS Enterprise Data Warehouse using layered Data Warehouse for IBM Dynamic Warehousing is the basis for our discussion.

Layered “Data Architecture”



We have found that the layered data architecture enables the warehouse and marts, with a variety of levels. These levels are virtualized, so that we can have line of business warehouses, combined information, subject area data and enterprise level information. Here are suggested levels:

5: Alphablox, DataQuant, MQTs, WebSphere Portals...

4: Logical DM using Views for flexibility

Physical DM using MQTs for performance

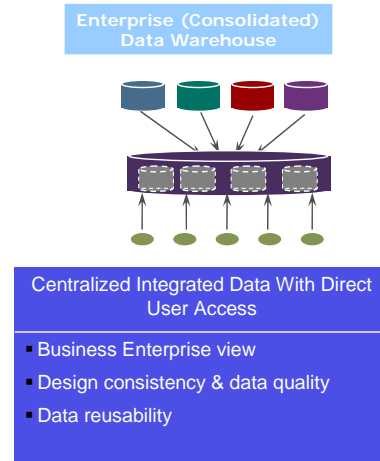
3: ELT, MQTs, DA, SQL Function...

2: Mining, MPP, scale, growth, Merge...

1: Near RealTime, MQ, LOAD, ||ism, multi-row insert & fetch...

Layered Data Architecture Offers a Proven Solution

- Enables Cross-enterprise View of the business – **single view of the truth**
- Reduces information latency
 - Enables business process management
 - Enables business performance management
- Reduce costs and complexity
 - Reduces data redundancy
 - Reduces ETL development costs and latency
 - Maximizes utilization of resources
- Enables iterative development of additional business areas



Two logical questions that might be asked at this point are: 1) Where does the layered data architecture fit relative to the other two approaches described earlier? 2) How does the layered architecture approach address the challenges of the other two? First, in both of the commonly used alternatives, the transaction history data and the business views are typically isolated into separate operating environments. In the layered data architecture design, both the historical business process (transaction) data and the information which has been organized to provide application specific business views are contained in the same database environment. This facilitates providing consistent cross-enterprise views of the business.

Second, because this layered approach enables the concurrent ingest of new transactional data while maintaining user access to both detailed data and business information views, it reduces the data latency associated with traditional batch oriented warehouse updates. As a result, it is possible to control business processes as they occur using up-to-date information.

Third, because this approach enables business information views to be provided with good performance through database views (providing a sort of logical data mart), business users have the chance to alter the views early on without incurring significant ETL development effort; later, when they have determined that the view meets their information needs, it can be converted to a physical data mart. And since the layered database design supports both the atomic historical data repository and the data marts in the same parallel database environment, resource use in the infrastructure is optimized as all the work makes use of all the resources under control of a workload manager to ensure SLAs are being met.

Finally, this layered approach provides the best of both of the commonly used approaches. Starting with your logical data model, you can choose which part of it to implement in both the atomic repository and the business information layer. This provides faster time to implementation than a “big bang” approach. When ready, the next subject area can be implemented. This provides an effective way to tie benefits to costs, while avoiding the issues we described with the independent marts approach.

What is available on System z to support and run your Data Warehouse



DB2 for z/OS Enterprise Data Warehouse in IBM Dynamic Warehousing

GoFurther

This next section discusses the primary options that run on System z to run and support the data warehouse.

What DW/BI Capabilities Exist Today for System z

- **DB2 z/OS V8 & 9 EDW Base:** many functional and performance enhancements, Easier online reporting and data management capabilities, Better DB2 family compatibility – including MQTs
- **Integration (EDW ETL):** WS DataStage EE, WS II Classic Federation, WS Classic Event Publishers, Distributed DBMS Event Publishers, BatchPipes for OS/390, DB2 Unload/Load Utilities, Replication & DataMirror
 - Information Server (WS Information Analyzer, WS QualityStage)
- **Analyze/Report:** Alphablox, QMF, DataQuant
 - Partners Cognos, Business Objects, SAS, IBI, ESRI, MicroStrategy, Hyperion, ...
- **Performance Management:** IBM Tivoli Omegamon XE for DB2 Performance Expert on z/OS, DB2 Query Monitor, Optimization Service Center & Optimization Expert (9)
- **Security and Regulatory Compliance:** Princeton Softech Optim, Consul Suite, DB2 Test Database Generator, DB2 Audit Manager Expert, IBM Encryption for DB2 and IMS Databases, DB2 Data Archive Expert, ...
- **Application Management:** DB2 Table Editor, DB2 Web Query Tool, DB2 Automation Tool, DB2 Query Monitor

GoFurther

Queries and data warehouses: DB2 V8


- Optimization Improvements
 - ❑ Improved techniques
 - ❑ Enhanced data
 - ❑ Visual Explain
- Enhanced index options
- New partitioning options
- Materialized Query Tables
- Schema evolution
- QMF improvements
- SQL enhancements



Announcing

DB2 UDB for z/OS

Version 8



GoFurther

Queries and data warehousing are improved a lot in V8. Optimization improvements provide a performance boost and make the job simpler. Improved optimization techniques like ability to use indexes more, star join and scale improvements allow reduced work for computers and for people. Enhanced data helps get the best access path. Visual Explain improves the ability to analyze and resolve any problems.

The many improvements for indexes, materialized query tables and partitioning can save space and add new options for improved performance and availability, even while simplifying the process. Not padded, clustering, longer and backward scans help indexes. Being able to add, rotate and rebalance partitions improve partitioning options. QMF enhancements build upon these strengths and add new function to reporting, dash boards, and a new platform in WebSphere. SQL enhancements on this page and the next improve portability of the SQL, improve the ability to express queries, and help with performance.

Specific DB2 V8 enhancements benefiting DW

- Usability, Availability & Scalability
 - Online Partitioning Changes, Up to 4096 Partitions, Rotate
 - Schema Evolution: ALTER instead of DROP & CREATE
 - VSCR with System z 64-bit Architecture
- Overall System & Query Performance
 - Locking Improvements
 - Multi-row INSERT & FETCH
 - Index Access for VARCHAR, data type mismatch, ...
 - DDF Performance Improvements
 - zIIP for improved cost of ownership, ability to use parallel
- Data Warehousing Support
 - Star Join Improvements
 - Materialized Query Tables

GoFurther

Queries and data warehousing are improved a lot in V8. Optimization improvements provide a performance boost and make the job simpler. Improved optimization techniques like ability to use indexes more, star join and scale improvements allow reduced work for computers and for people. Enhanced data helps get the best access path. Visual Explain improves the ability to analyze and resolve any problems.

The many improvements for indexes, materialized query tables and partitioning can save space and add new options for improved performance and availability, even while simplifying the process. Not padded, clustering, longer and backward scans help indexes. Being able to add, rotate and rebalance partitions improve partitioning options. QMF enhancements build upon these strengths and add new function to reporting, dash boards, and a new platform in WebSphere. SQL enhancements on this page and the next improve portability of the SQL, improve the ability to express queries, and help with performance.

DB2 SQL

z z/OS V8

common

luw Linux, Unix & Windows V8.2



z

Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning

c

o

m

n

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, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

l

u

w

Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE or DELETE, multi-site join, MDC

GoFurther

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.

IDUG® 2007 Europe

DB2 SQL

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


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, index compression**

c { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, 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, data compression**

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

u

w



GoFurther

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 V9 for z/OS, (V9 or DB2 9). V9 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 for LUW, code named Viper is already generally available. We are able to move more from the 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.

Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

IDUG® 2007 Europe

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**

C { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, 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, data compression, Session variables, DECIMAL FLOAT, optimistic locking, ROLE**

I { Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more functions**, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, **XQuery, arrays, global variables, 128 byte names**

U {

W {

GoFurther

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 V9 for z/OS, (V9 or DB2 9). V9 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 for LUW, code named Viper is already generally available. We are able to move more from the 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.

Sheryl Larsen provided the base for this information, but the mistakes are probably mine.

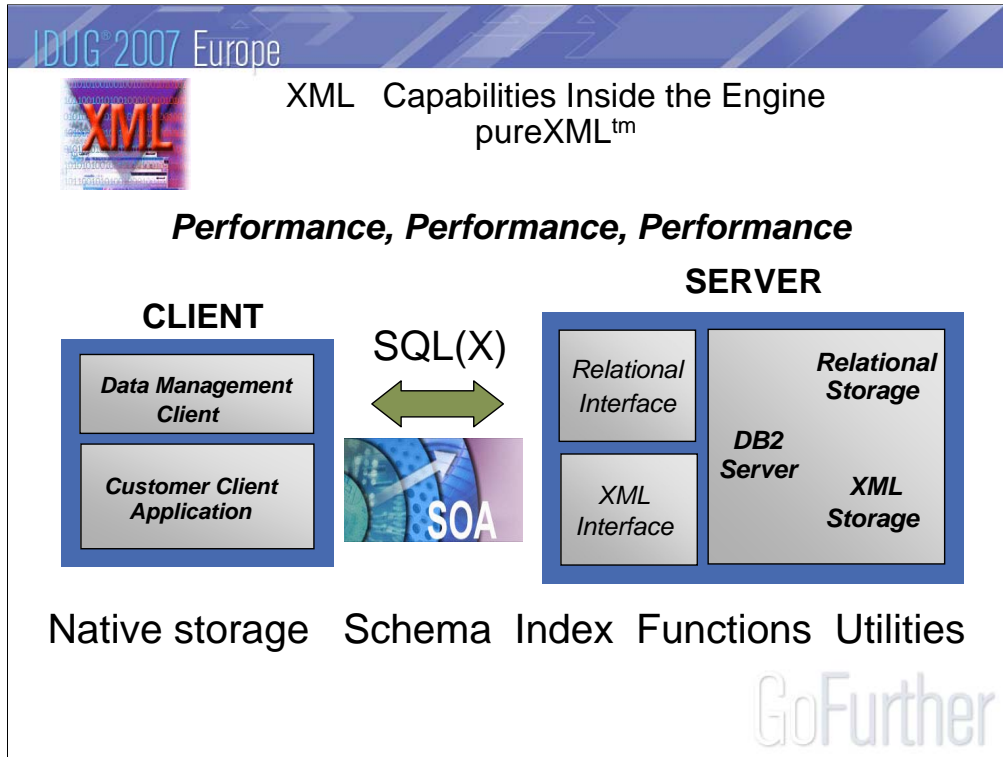
DB2 9 enhancements benefiting DW

- Usability, Availability & Scalability
 - XML, text, and spatial data
 - Universal & partition by growth table spaces
 - Clone Tables (Online Load Replace)
 - Insert and update scaling: logging, index changes
- Overall System & Query Performance
 - TRUNCATE for fast delete, MERGE for insert or update
 - Index compression, index on expression
 - Significant CPU reduction for DB2 utilities & many queries
- Data Warehousing Support: enhanced data, techniques, OSC & OE
 - Star join improvements (dynamic index ANDing)
 - SQL Enhancements (RANK, DENSE_RANK, INTERSECT, EXCEPT, ...)

GoFurther

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. V9 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. V9 is another big step in DB2 family consistency and in the ability to port applications to DB2 for z/OS.



The largest single item in DB2 9 for z/OS is the XML work, bringing pureXML™ for native storage and integrating XML with object-relational and DB2. Hybrid data server.

Powerful querying and transformation capabilities are included with XPath. Querying is not merely finding a document that matches a certain criteria. For e.g. example, 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.

JWJ: pure XML implementation on Z. Take XML doc, store in binary encoded format. New indexing technology for efficient search of XML. New query syntax w/ xpath expressions that can be used to perform the searches. Even have schema validation – as you insert the info, you can validate that the xml doc satisfies the schema. Allow the z DB server to participate in a lot of XML oriented apps and could be a key part of what people do with SOA.


JWJ: XQuery is not supported on Z. This is not a major limitation since SQL will likely see much higher use in most XML applications since XQuery does not support things like parameter markers.

Usage scenarios:

- 1) Document storage. XML column is provided. Decomposition SP is also provided to extract data items from an XML document and store into relational.
- 2) XML document retrieval. You can use SQL to retrieve entire documents from an XML column, just as you would for any other datatype. When you need to retrieve portions, you use XPath expressions via SQL/XML.
- 3) Admin support: XML schema repository for validation. Utilities support.

Text Search Server – delivering later in DB2 9

- Text search for CHAR, VARCHAR, CLOB & XML columns
- Provide a text index server
- Efficient communication interaction with DB2 for z/OS
- Text indexes are persisted into DB2 tables for backup & recovery purposes



GoFurther

A specialized text search engine on a separate server is expected to come in DB2 9, but will deliver after general availability. While the server is separate, the text indexes are saved into DB2 tables to provide improved backup and recovery.


IDUG 2007 Europe

DB2 9 Spatial Support

Enabling Open Geospatial Consortium (OGC) compliant geospatial applications

- **Spatial data types**
- **Spatial functions and predicates**
- **Spatial indexes**
- **Spatial search**
- **OGC-compliant spatial catalog**

Imminent delivery for Phase II (about 30 more functions)



GoFurther

GIS = Geographic Information System. Aka Geospatial. E.g. Google Maps.
 OGC = Open Geospatial Consortium

You can use IBM Spatial Support for DB2 for z/OS to generate and analyze spatial information about geographic features, and to store and manage the data on which this information is based. A geographic feature can be:

An object (that is, a concrete entity of any sort); for example, a river, forest, or range of mountains.

A space; for example, a safety zone around a hazardous site, or the marketing area serviced by a particular business.

An event that occurs at a definable location; for example, an auto accident that occurred at a particular intersection, or a sales transaction at a specific store.

IBM Spatial Support for DB2 for z/OS makes available facts and figures about the locations of geographic features. You can use functions to provide spatial data. IBM Spatial Support for DB2 for z/OS provides services to import spatial data in Shapefile and SDE Export formats.

Basically it allows you to add a spatial column (e.g. LOCATION) to a DB2 table (that has customer info including an address). You can then use a GEOCODER to translate the address into coordinates. The coordinates are stored in the spatial column LOCATION

You can then use ESRI's Arc Explorer to combine a map with your DB2 table to visualize where 33

Dynamic Index AND-ing for Star Join (Pair-Wise Join with Join Back)

- Multi-index access steps are considered independent
- Apply filtering to dimension tables before the fact table
 - Exploit single and/or multi-column fact table indexes
- Runtime assessment of filtering
 - Pre-fact dimensions with poor filter factors can be discarded at runtime and accessed post-fact
- Independent join of each dimension table to fact table via index
- The result of each pair-wise join is a set of fact table rids
- Perform Rid Sort and Rid Merge (ANDing) to generate final fact table rid list
- Final Rid list then used to retrieve data from Fact table
- Join back to dimension table(s) as necessary.

GoFurther

Better performance for star join queries.

Easier to design indexes – multi-col indices no longer needed.

More aggressive use of parallelism, more zIIP eligible.

Histogram Statistics - RUNSTATS

- V8 – DB2 has data skew awareness for single values
- Histogram statistics addresses skews across ranges of data values
- Summarizes data distribution on an interval scale
- DB2 uses equal-depth histograms
 - Each quantile has about the same number of rows
 - Example - 1, 3, 3, 4, 4, 6, 7, 8, 9, 10, 12, 15 (sequenced), cut into 3 quantiles

Seq No	Low Value	High Value	Cardinality	Frequency
1	1	4	3	5/12
2	6	9	4	4/12
3	10	15	3	3/12

OPT

GoFurther

RUNSTATS

Maximum 100 quantiles for a column

Same value columns WILL be in the same quantile

Quantiles will be similar size but:

Will try and avoid big gaps between quantiles

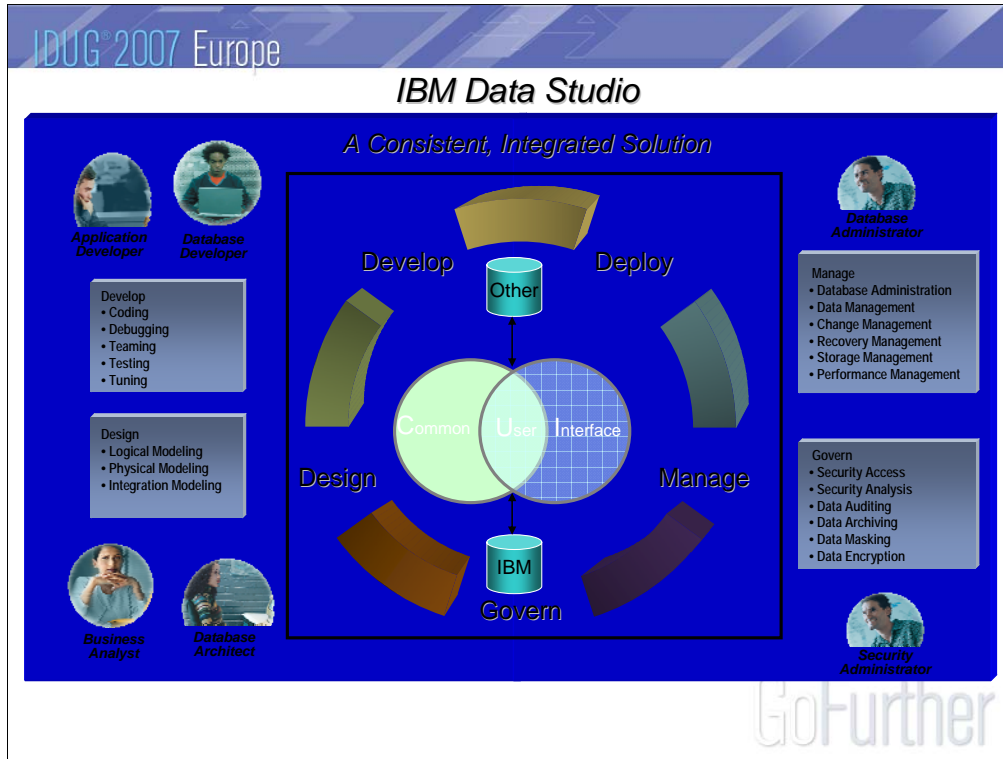
A column high value and low value may have separate quantiles

Null WILL have a separate quantile

If less than 100 column values, reverts to Distribution Stats

Not supported with LOAD and REORG

Supports column groups as well as single columns



Slide objective: set up for onion peel. First we give you a comprehensive look at the total Consul portfolio

Points:

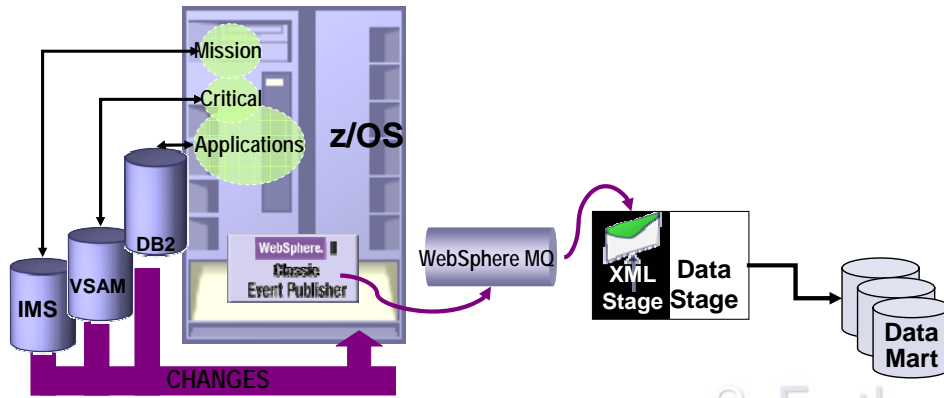
1. Consul provides the depth of security management on the mainframe, right side of the diagram and the breadth across the enterprise with comprehensive compliance management from the distributed environment on the left side of the diagram
2. Distributed portfolio is InSight
 1. Focus is on access monitoring, log management and compliance reporting
3. Mainframe portfolio is zSecure
 1. Focus is on integrated mainframe audit, monitoring, compliance and administration
 2. z/OS
 3. RACF
4. The intersection is the overlap and illustrates how Consul integrates mainframe data into InSight to provide the comprehensive compliance management view across the enterprise

*Data Integration and
Warehousing Tools*

GoFurther

Near Real Time data ingest using Classic Event Publisher and DataStage

- Dynamic, changed-data feed
 - Maximize data currency while minimizing and stabilizing bandwidth utilization
- Reliable and recoverable
 - Recoverability is built-in, and WebSphere MQ assures high performance delivery

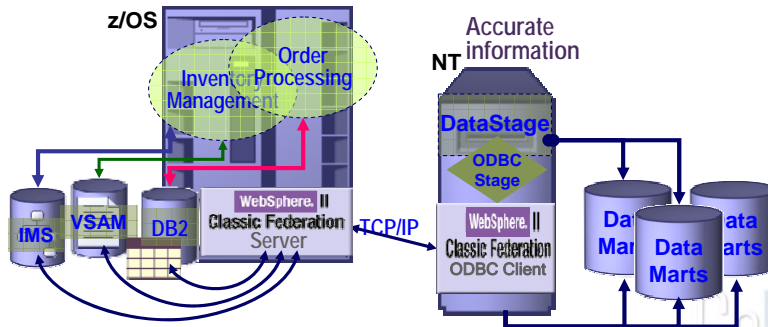


GoFurther

Operational data feed to ETL via SQL using Classic Federation

- **Dynamically connect data warehouse tool with mainframe data**

- No dependence on mainframe development
 - Dramatically simplifies management: one team owns it all
- One consistent process leverages “power” of ETL tooling
 - Development time “cut in half”
- Empowers additional uses
 - Dynamic query by business intelligence tools extends the warehouse



Available March 2007: IBM DataQuant

- Provides a comprehensive query, reporting and data visualization platform for both web and workstation-based environments
- Together with over 100 built-in analytical functions, IBM DataQuant allows organizations to derive maximum value from their data and rapidly build and distribute comprehensive data visualization solutions across the enterprise.
- Visual Dashboards, Enhanced Graphical Reporting, Security and Personalization, SOA Layer, Enhanced Analytics



GoFurther

DataQuant for z/OS, V1.1 (5697-N64) delivers a comprehensive query, reporting, and data visualization platform for both Web and workstation-based environments. While remaining compatible with the IBM Query Management Facility (QMF™) product line, DataQuant introduces a variety of powerful business intelligence capabilities, including:

- Powerful graphical reporting environment that allows rapid development and deployment of executive dashboards, information portals, and interactive data visualization solutions.
- Drag-and-drop development of OLAP analytics, SQL queries, tabular reports, graphical reports, pivot tables, and data analysis views.
- Dozens of charts, controls, and graphical primitives to visual reports, along with the ability to embed subqueries that provide supporting details.
- Full compatibility with QMF infrastructure and objects, including the ability to create, open, edit, delete, and save QMF queries, procedures, forms, and visual reports.
- Over 100 built-in mathematical and analytical functions accessible in both dashboards and visual solutions as well as in printable reports.

Rich security infrastructure, providing personalization of both the query and reporting environment as well as business intelligence content distributed within it. This infrastructure includes the following: ...

- Supports single sign-on and optionally interfaces with existing LDAP directories or internally defined DataQuant user directory.

What are the components of DataQuant?

- DataQuant for Workstation
 - Java-based, **rich client application** based upon the award-winning Eclipse platform
 - Provides a graphical environment for the development of query, report, dashboard and database visualization solutions
 - Supports multiple workstation environments, including Windows, Linux, HP-UX and Solaris
- DataQuant for WebSphere
 - Provides a comprehensive runtime environment, capable of displaying all DataQuant content, under a HTML, '**thin client**' model
 - Runs under IBM WebSphere Application Server (WAS) on all platforms supported by WAS, including Windows, Linux, Solaris, HP-UX, zLinux and z/OS

GoFurther

– Tailors look and feel, including available reports, visualizations, and data on a per user or group basis. For example, technical resources may see a traditional database-centric view whereas business users see a role-specific view of relevant reports and dashboards.

– Optionally logs on to data sources using accounts mapped from LDAP/internal security accounts.

– Tailors offline schemas to each user or group so that individuals see only those tables and columns relevant to their job function.

- Eclipse-based rich desktop environment that supports numerous workstation operating systems, including Windows™, Linux™, Solaris, and AIX®.

- High-performance, WebSphere®-based environment that extends key DataQuant functions to browser-based users across multiple platforms, including querying, content authoring, and drag-drop data analysis.

- Service-oriented architecture (SOA) infrastructure, enabling enterprises to share specific business intelligence content and data with partners over secure Internet or intranet connections.


- Support for DB2 and Informix™, as well as other popular database systems.

- Building and deploying information-rich business intelligence solutions that concurrently draw data from one or more heterogeneous relational databases.

With DataQuant you can derive maximum value from your data and rapidly build and distribute comprehensive business intelligence solutions across your enterprise.

*Hardware and Processor
Advancement*

GoFurther

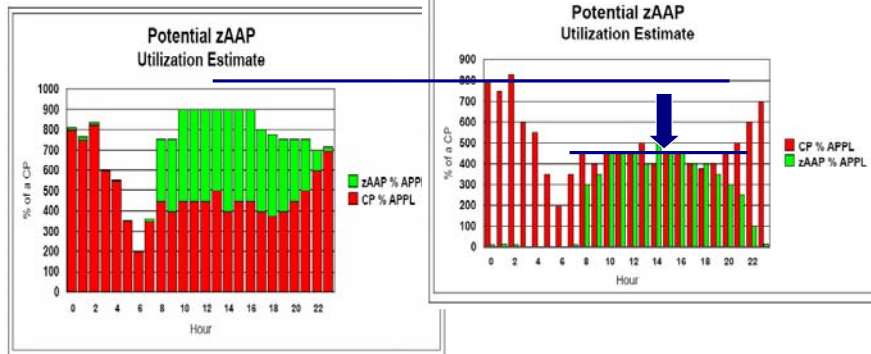


The IBM family of specialty engines have been delivered over a number of years for the diverse work loads, ranging from a Coupling Facility in 1997 and Linux in 2000 to the Java work loads in 2004 and some database work in 2006.

The latest change to deliver is the ability to use the zIIP for IPsec encryption in 2007. A statement of direction was provided for z/OS XML to be able to use zAAP or zIIP for XML parsing.

These specialty engines can be very effective at reducing overall costs in the data warehousing environment.

Why do zIIPs, zAAPs and IFLs Reduce Cost?



1. **Hardware costs:** By moving workload from general purpose processors to zIIP, zAAP and IFL processors (higher cost to lower cost processors).
 2. **Software Costs:** license/maintenance costs based on number of and usage of general purpose central processors. Specialty engines can reduce use of CP's.
- No z/OS software charges based on zIIP, zAAP and IFL processors or usage.

GoFurther

zIIPs and zAAPs do not add functional capabilities, but they do address hardware and software costs. In addition to being lower cost processors for the specialty purposes, they allow you to reduce the license and maintenance cost for software on z/OS, as there is no z/OS software charge for processing running on zIIP, zAAP or IFL processors.

I've talked this over with several customers who estimated that the zIIP is 5 times to 20 times less costly. If 50% of the query work can be run within a zIIP or zAAP, then the net would be a 40% reduction in costs for running the work load. There are other improvements in getting to the latest technology, but this is a significant step.

DB2 & IBM zIIP value

F06 tuning to use zIIP

Portions of DB2 for z/OS V8 and **DB2 9 (in blue)** workloads benefit from zIIP*:

- 1 - ERP, CRM, Business Intelligence or other enterprise applications
 - Via DRDA over a TCP/IP connection (enclave SRBs) (enclave SRBs, not stored procedures or UDFs)
 - **DB2 9 for z/OS Remote native SQL procedures**
 - **DB2 9 XML parsing via DRDA to fully utilize zIIP (statement of direction)**



- 2 - Data warehousing applications*: SQL Requests using parallel queries, including star schema
 - **DB2 9 higher percentage of parallel queries eligible for zIIP**
- 3 - **DB2 Utilities LOAD, REORG & REBUILD*** (DB2 utility functions used to maintain index structures)


* 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.

DB2 9 uses zIIP in two new ways, remote native SQL procedures and increased use of parallelism. See presentation F06 by Terry Purcell. 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/>

Typical zIIP redirect for a DW workload



GoFurther

There is a wide range of query work that can be redirected to a zIIP. The primary work for a data warehouse is the parallel processing. For large queries that are processed in parallel, as much as 80% of the work can be redirected to a zIIP. Some cases will be smaller, since part of the work will not be run in parallel. Small queries and the initial part of large queries (under 100 ms of cpu time) do not use the zIIP – unless the work is from remote SQL. If the work comes in remotely over TCP/IP and DRDA (not stored procedures), then the DRDA redirect can be used as well for the initial work.

IDUG 2007 Europe

A vision for System z advanced data serving System z Enterprise Hub for Mission Critical Data

- With a strong foundation for transaction processing, built on 40+ years of technology innovation, System z servers with z/OS and DB2 can provide a premier platform for data serving, today and into the future*
- IBM plans to continue to invest in new solutions to address customers' strategic information on demand goals*

Today's Capabilities

- Industry-leading data integrity and security
- Data sharing solution for centralized view of data
- Scalability and availability for enterprise class workloads
- Comprehensive systems and data management environment

Extension of capabilities*

- New specialty engine (zIIP) with DB2 exploitation - for mission critical distributed, ERP, CRM, and Data Warehousing workloads *
- DB2 9 and tools improve regulatory compliance and autonomies
- Encryption capability (TS1120 tape subsystem) with z/OS centralized key mgmt
- Data protection to achieve highest levels of security certifications

Future direction*

- Additional zIIP exploitation: DB2 9 adds Native SQL procedures
- DB2 enhancements to help improve usability and reduce complexity and management costs.
- Encryption for IPsec zIIP
- XML parsing capability zIIP or zAAP
- DB2 table scan acceleration via DS8000
- Encryption capability (disk subsystem) with z/OS centralized key mgmt
- Handle larger volumes of data, with improved scalability

*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

We have been working on specialty processors for a long time, and have just delivered new capabilities with the zIIP. The improved cost of ownership can help a lot for some customers, but not at all for others. There have been several big recent changes:

The z9 Business Class and Enterprise Class extend zIIP capabilities to many more customers. Only the largest customers needed the z9-109 processors, and the upgrade steps were very large ones. The new z9 Business Class and Enterprise Class processors have a much greater range of processing power with more granular upgrade options. The entry level z9 processor now starts at under \$100,000.

Query work is broadened beyond just the star joins to all large parallel queries. If you have a warehouse that uses parallel processing and significant CPU time, then the zIIP may provide a benefit.

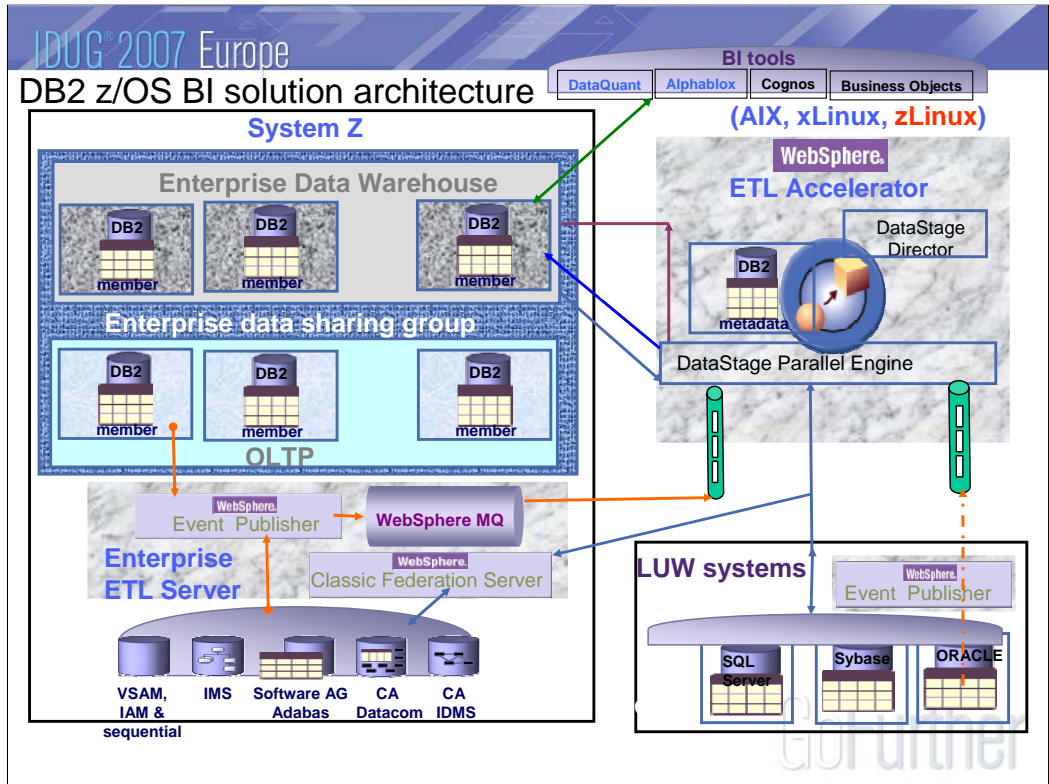
The TS1120 tape subsystem has added encryption capability with several options for centralized key management.

The Data Serving Roadmap provides more information about this slide:
<http://www.ibm.com/systems/z/ziiip/data.html>

BI Solution Architecture



GoFurther



EDW

Highest Qualities of Server: highly available, best security,

Single DB system view with OLTP

Flexible yet high performance layered DWH

Flexible logical data marts (views, partitions, online DDL) with performance objects and the flexibility to convert to physical data marts for even more performance once flexibility is not as important

Performance objects (MQT, user aggregates, parallelism)

ETL Accelerator

Single integration metadata store

Single event manager

Flexible transformations

Parallel engine

Near real time data streaming and/or high performance bulk transfers

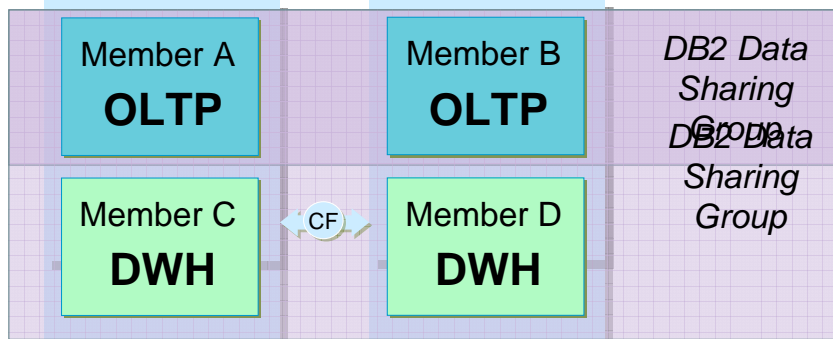
Extensive data sources on z/OS and distributed systems

Modern BI reporting and Analytics

GenevaERS for static report consolidation (satisfy many reports with one access to the DWH)

DataQuant and/or Alphablox for reporting and analytics

Solution Architecture



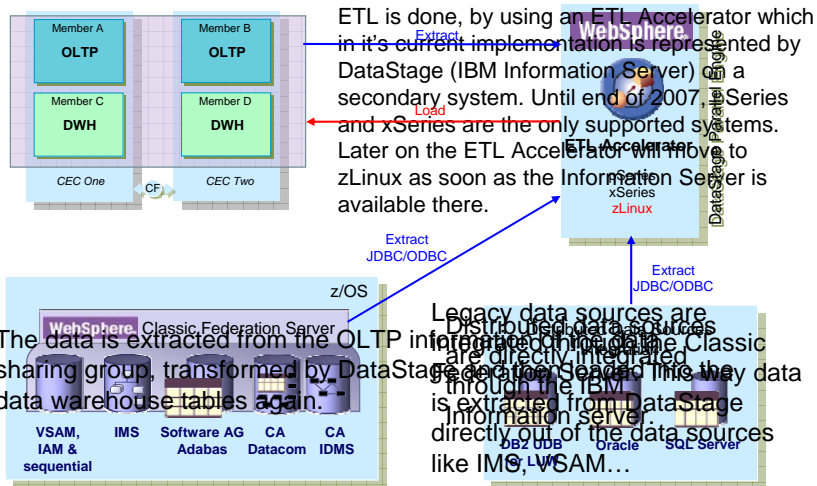
Within a Data Sharing environment, the Data Warehouse resides in the same group as the transactional data.

GoFurther

If the customer already uses a parallel sysplex and DB2 data sharing, the data warehouse should be implemented as additional members within the same sharing group. This enables the single view on the complete system (OLTP and DWH) but allows specialized optimizations for the workloads. Data of both environments can still be joined and accessed as one database.

Having the members on different CECs allows implementation of HA requirements.

Initial load of the Data Warehouse

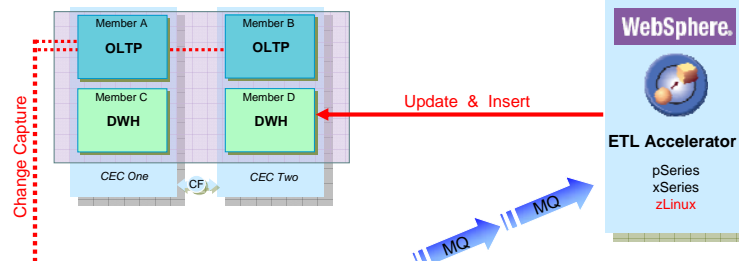


For the initial load, the existing System Z data, stored outside DB2, can be accessed through WebSphere Classic Federation server. It allows a relational access on non relational data and is responsible for mappings between hierarchies (IMS) and flat tables or also allows the DBMS like access to VSAM or flat files. All data is access through a single ODBC driver, no matter if it is stored within IMS or VSAM. This way, the ETL Accelerator can directly read all necessary information w/o having to care about the different access methods of the host data sources.

The distributed data sources can directly be integrated through the Federation Server of the IBM Information Server, which builds the foundation for DataStage on the ETL Accelerator. The accelerator can be a pSeries, Linux or even Windows System. In future, zLinux will be recommended as platform for this component.

The existing OLTP data is accessed directly by DataStage through DRDA. All transformed, cleansed and joined results are stored back to the DB2 warehouse within the same data sharing group.

At runtime, the Data Warehouse is updated incrementally




Once the data warehouse is loaded, only incremental updates are performed. The changes on the original data sources (z/OS and distributed) are captured by WebSphere Data Event Publisher and sent through MQ to the ETL Accelerator. Special DataStage stages take the change information from MQ and updates or inserts are performed on the data warehouse data. This might of course be delayed to batch windows by just receiving the changes and only updating the warehouse at a later point in time.

GoFurther

Once the data warehouse is loaded, it needs to be updated in incremental mode. On the host side, WebSphere Classic Event Publisher is used to capture any changes in the data sources like IMS, VSAM or even the OLTP part of DB2. These changes are then pushed through MQ series pipes directly into specialized DataStage input stages. The changes are then fed into the data warehouse immediately or within defined windows.

Distributed data sources like DB2 LUW or in future Oracle can also be used for change capture and feed their updates into the accelerator.

In Database ELT is triggered by DataStage



Simple example: Wherever possible, "In Database" transformations (ELT) are used to spare the transport of the data to the accelerator. But the used SQL is still sent from the ETL Accelerator to the database to have one place of documentation for all ETL steps.

```

-- Aggregate by salary by department into aggregates
INSERT INTO AGGSALARY ( DEPTCODE, AVGBAND, AVGSALARY )
SELECT DEPTCODE, AVG( BAND ) AS AVGBAND, AVG( SALARY ) AS AVGSALARY
FROM EMP
GROUP BY DEPTCODE
    
```

This can also be used to shift the data up the hierarchy within the Layered Data Architecture.

GoFurther

Most efficient transformations happen within the database without the data ever leaving the DB2 system. This can be done whenever SQL delivers all necessary functionality to perform the transformation. Aggregations from detailed data into day, week or month summaries can be one example for this kind of in database transformations, here described as ELT (Extract Load Transform) vs. ETL.

5-9 November

Athens Hilton

Athens, Greece

IDUG® 2007
Europe

Dynamic Warehousing in System z - Summary



GoFurther

When is System z the Preferred Platform?

- **True Real-time Operational Data Store (ODS)**

- Data must be insync with the operational data
- Availability and Resiliency
- Operational data is on System z

- **Existing marts or warehouses are on System z**

- **System z and SAP: operational and warehouse,**

- E.g. Using SAP BW when SAP R/3 is on System z

- **Existing skills and investments on System z**



GoFurther

Data warehousing is becoming more real time, with more connections to operational data and more security concerns. These changes make the key advantages of System z even more applicable. Today's complex applications include both transactions and reporting, so performing both well is imperative. Operational data stores are often kept on z/OS, since that is the source of most of the information, and synchronization across platforms is more challenging. When availability and resiliency are very important, many of the robust, unique characteristics of System z and DB2 for z/OS provide an advantage.

Another common situation is having the operational and warehouse on the same platform. For example, if the SAP operational systems are on z/OS, then having the warehouse there too makes the integration simpler.

If customers have existing skills and investments on System z, they know the advantages. Further improvement like the zIIP and improvements in the latest versions of DB2 for z/OS can make System z the best total cost of ownership choice.

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.

Value proposition of DW on System z

- **Qualities of Service**
 - Superior Quality
 - Continuous Availability
 - Security and Regulatory Compliance
 - Scalability
 - Backup and recovery
- **Positioned for the future**
 - Web-based applications
 - XML support
 - Service Oriented Architecture (SOA)
- **Operational data and the ODS together means**
 - Reduced complexity
 - Reduced cost
 - Shared processes, tools, procedures
 - Streamlined compliance and security
- **zIIP and other specialty engines improve Total Cost of Ownership**
- **Better leverage System z skills and investment**

GoFurther

IDUG 2007 Europe

DB2 9 for z/OS RedBooks & RedPapers

- Powering SOA with IBM Data Servers SG24-7259
- LOBs with DB2 for z/OS: SG24-7270
- Securing DB2 & MLS z/OS SG24-6480-01
- DB2 9 Technical Overview SG24-7330
- Enhancing SAP - DB2 9 SG24-7239
- Best practices SAP BI - DB2 9 SG24-6489-01
- DB2 9 Performance Topics SG24-7473
- DB2 9 Optimization Service Center SG24-7421
- Index Compression with DB2 9 for z/OS paper
- DB2 Stored Procedures SG24-7083 update

See the DB2 library for detailed information.

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

Eight redbooks and one paper with substantial DB2 9 content are on the web, with one more being written now. Check for updates.

- Powering SOA with IBM Data Servers, SG24-7259

<http://www.redbooks.ibm.com/abstracts/SG247259.html>

- LOBs with DB2 for z/OS: Stronger & Faster SG24-7270

<http://www.redbooks.ibm.com/abstracts/SG247270.html>

- Securing DB2 & MLS z/OS, SG24-6480-01

<http://www.redbooks.ibm.com/abstracts/sg246480.html>

- DB2 9 Technical Overview, SG24-7330, <http://www.redbooks.ibm.com/abstracts/SG247330.html>

- 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>

- DB2 9 Performance Topics, SG24-7473,

<http://www.redbooks.ibm.com/abstracts/SG247473.html>

- Index Compression with DB2 9 for z/OS, redpaper REDP4345

<http://www.redbooks.ibm.com/abstracts/redp4345.html?Open>

- DB2 9 Optimization Service Center, SG24-7421,

<http://www.redbooks.ibm.com/abstracts/SG247421.html>

- DB2 for z/OS Stored Procedures: CALL & Beyond SG24-7083-01 update being written

IDUG® 2007 Europe

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

Name	Size	Type	Modified
IOD1427johndeereESP2007Oct.pdf	1.03 MB	Adobe Acrobat 7.0 ...	11/5/2007 12:02 PM
IOD1094newSQLdb29zos2007Oct.pdf	730 KB	Adobe Acrobat 7.0 ...	11/5/2007 12:02 PM
IOD2316_db2zos_trendsdirections_2007Oct.pdf	3.10 MB	Adobe Acrobat 7.0 ...	10/23/2007 5:02 PM
IOD1022Adynamicwarehousesystemz2007Oct.pdf	1.31 MB	Adobe Acrobat 7.0 ...	10/22/2007 12:02 PM
IOD1649_justenOct2007.pdf	1.01 MB	Adobe Acrobat 7.0 ...	10/18/2007 6:02 PM
IOD1298_ADadminTrendsDirections_CotnerOct2007.pdf	1.59 MB	Adobe Acrobat 7.0 ...	10/18/2007 6:02 PM
IOD1297A_DBAccessJava_CotnerOct2007.pdf	297 KB	Adobe Acrobat 7.0 ...	10/18/2007 6:02 PM
IOD2082_userexp_Campbell2007Oct.pdf	395 KB	Adobe Acrobat 7.0 ...	10/18/2007 1:02 PM
IOD1019_bestpracticesecurity2007Oct17.pdf	1.52 MB	Adobe Acrobat 7.0 ...	10/18/2007 1:02 PM
IOD1018_migration2007Oct11.pdf	1.71 MB	Adobe Acrobat 7.0 ...	10/18/2007 1:02 PM
IOD1015_DB2storageIczkovits2007Oct.pdf	567 KB	Adobe Acrobat 7.0 ...	10/18/2007 1:02 PM
IOD_db29zos_sql_2007Oct16.pdf	881 KB	Adobe Acrobat 7.0 ...	10/18/2007 1:02 PM
IOD2007_OSC_OE_overview.pdf	6.90 MB	Adobe Acrobat 7.0 ...	10/17/2007 1:02 PM
IOD2007_OSC_OE_demo.exe	7.90 MB	Application	10/17/2007 12:02 PM
DB2zOS_v8_upgrade082407.pdf	367 KB	Adobe Acrobat 7.0 ...	10/8/2007 4:04 PM
S5_Access_DB2zOS_data_SOA_2007Aug08.pdf	1.26 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S4_XML_DB2zOS_2007Aug08.pdf	1.52 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S3_DB2zOS_evolution_SOA_2007Aug08.pdf	2.18 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S2_zIIP_DB2_workloads_2007Aug08.pdf	979 KB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S1340_best_practice_security_2007Aug09.pdf	1.34 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S1304_migration_2007July29.pdf	1.47 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S1301_New_DB29_Beyondr22n2.pdf	864 KB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S1300_cmc DB2zOS_Intro_CatherineCox.pdf	940 KB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM
S1 DB2 SOA todav 2007Aug08.pdf	3.60 MB	Adobe Acrobat 7.0 ...	8/14/2007 1:04 PM

GoFurther

This is the primary DB2 for z/OS ftp server. It has hundreds of presentations, with all of the ones shown (and more) updated in the last two months.

To see the most current files, you can use the ftp server directly with Internet Explorer, then sort by the date modified.

System z road show

- 1 DB2 and SOA today
- 2 Leveraging zIIP for DB2 workloads
- 3 DB2 for z/OS evolution in an SOA world
- 4 XML in DB2 for z/OS
- 5 A brave new world:

Accessing DB2 for z/OS data in SOA

- 6 DB2 tools update: The key to optimized DB2 environments and reduced TCO
- Handouts for all road show sessions:

<http://www.ibm.com/software/zseries/db2seminars/handouts>

<http://www.ibm.com/software/os/zseries/roadshows/handouts/> old

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



GoFurther

1 DB2 and SOA today: DB2 for z/OS Version 8 supports a number of SOA capabilities, and this session will offer an update on how IBM customers are benefiting as they upgrade to Version 8 and fully exploit these rich functions. This session starts with an introduction to SOA to provide a level set, then looks at some current capabilities.

2 Leveraging zIIP for DB2 workloads: IBM announced its new zIIP specialty engine offering in 2006. The zIIP is designed to help process certain DB2 related tasks at a reduced software license and processor cost. The workloads which benefit are distributed SQL access, parallel queries, and some utility processing. SAP customers are a good example. In this session, we will explain how you can start making use of the zIIP and which workloads qualify for the zIIP. We will also provide some measurements and best practices on how to maximize the usage of the zIIP.


3 DB2 for z/OS evolution in an SOA world: This presentation will provide an overview of trends and directions of DB2 on z/OS. Special focus will be on SOA enablement and on DB2 9.

4 XML in DB2 for z/OS: DB2 for z/OS provided an XML extender in DB2 Version 7, then added XML publishing functions in DB2 Version 8. DB2 9 provides extensive changes for pureXML. In this session we explain how to use XML with DB2 for z/OS and provide different usage scenarios including how to leverage XML data in an SOA.

5 A brave new world: Accessing DB2 for z/OS data in an SOA: DB2 data and stored procedures can be accessed via Web services, or via Java programs using JDBC or SQLJ. Come learn about the benefits of each of these technologies, along with tips on how to select the right method for your requirements, and best practices in the

Data Warehousing for DB2 on System z – more info

- **Whitepaper** : Why Data Warehousing on System z
<http://www.ibm.com/software/data/db2bi/systemz.html>
- DW on system z – **Demo** available in the Technical Marketing Competence Center, Böblingen, Germany, TMCC@de.ibm.com
- DW on system z – **Customer Briefings** possible in the Technical Marketing Competence Center, Böblingen, Germany or at Silicon Valley Lab TMCC@de.ibm.com
- DW on system z – **Best Practices** paper planned
- Redbook: Best Practices for SAP BI using DB2 9 for z/OS, SG24-6489-01
- ...



GoFurther

Here are some resources on the web about System z and data warehousing.

IDUG® 2007 Europe

United States [change]

IBM

Home Business solutions · IT services · Products · Support & downloads · My IBM ·

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

Software > Information Management

Dynamic Warehousing from IBM
A new approach to delivering enhanced business insight.

MOVE FROM WAREHOUSING TO DYNAMIC WAREHOUSING

Leverage embedded analytics in your business processes.
Incorporate knowledge from unstructured information.
Support real-time access to aggregated, cleansed information.
Use integrated capabilities for Information On Demand.

Why IBM

IBM is the first to fully understand and define this next generation of data warehousing. Only IBM can deliver all of the Dynamic Warehousing capabilities with a product portfolio that is best in class.

New white paper

From conflicting unintegrated historical data to actionable insight.

Related links

- DB2 Data Warehouse Edition
- DB2 Alphablox
- Business Intelligence on System z
- OmniFind Analytics
- OmniFind Discovery for Business Intelligence
- DB2 Data Server Family
- Master Data Management
- Warranties, licenses and

See these web pages for more on Dynamic Warehouse.
<http://www.ibm.com/software/data/db2bi/data-warehousing/>

Software > Information Management >

Business Intelligence on System z

Description	Product offerings
<p>The distinction between data warehousing applications and online transaction processing (OLTP) applications is blurring. Data warehousing and analytic applications are accessing operational data or near-real-time data.</p> <p>IBM provides a sophisticated data warehousing solution including - integrating/building warehouses, analyzing and reporting on that data, streamlining security and regulatory compliance, performance tuning, and application management capabilities. It combines the traditional strengths and new capabilities of DB2 Universal Database and the</p>	<ul style="list-style-type: none"> • WebSphere DataStage for z/OS takes ETL performance to a new level with powerful data integration and transformation. • WebSphere Replication Server for z/OS replicates data between DB2 for z/OS and other databases for high availability • DB2 Alphablox imbeds analytics directly into web based applications. • Rational Data Architect helps data architects understand information assets, design federated databases, and streamline database projects.

Information Management software

IBM Information On Demand 2007

Oct 14-19, 2007

Register Now!

Early, Early Bird Rate \$1595

The Premier Information Management global conference

Highlights

→ Announcement Letter 206-181

White papers

[Why Data Serving on a Mainframe](#)

Events

→ Replay: 'Why BI on z9/System z makes "cents"'

→ Sept 7 Webcast on 'Data Warehousing on ...'

Go further

Business intelligence on System z:

<http://www.ibm.com/software/data/db2bi/systemz.html>

Thank
YOU

GoFurther

Disclaimer

The information contained in this presentation has not been submitted to any formal IBM review and is distributed on an "As Is" basis without any warranty either expressed or implied. The use of this information is a customer responsibility.

The materials in this presentation are also subject to

- enhancements at some future date,
- a new release of DB2, or
- a Programming Temporary Fix (PTF)

IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, zSeries, System z.

THE FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC and WINDOWS 95.

For additional information visit the URL

<http://www.ibm.com/legal/copytrade.phtml> for "Copyright and trademark information"

GoFurther

Roger Miller

IBM Silicon Valley Lab

millerrl@us.ibm.com



GoFurther