

Agenda



Announcing: DB2 Information Integrator real time data integration

- Information Integration in on demand computing
 - ▶ A fundamental underpinning
- DB2 Information Integrator 8.1
 - ▶ Federated query server and replication server combined
 - ▶ June 13, 2003 GA
- Pricing and packaging details
 - ▶ Five licensing options
- World-Class Partner Support
 - ▶ Broad range of tool interoperability
- Customer Experiences and Value
 - ▶ Cross-industry and geography




On Demand Business

Requires an On Demand Operating Environment

Business Transformation

- Align IT processes with business priorities
- Enable business flexibility and responsiveness

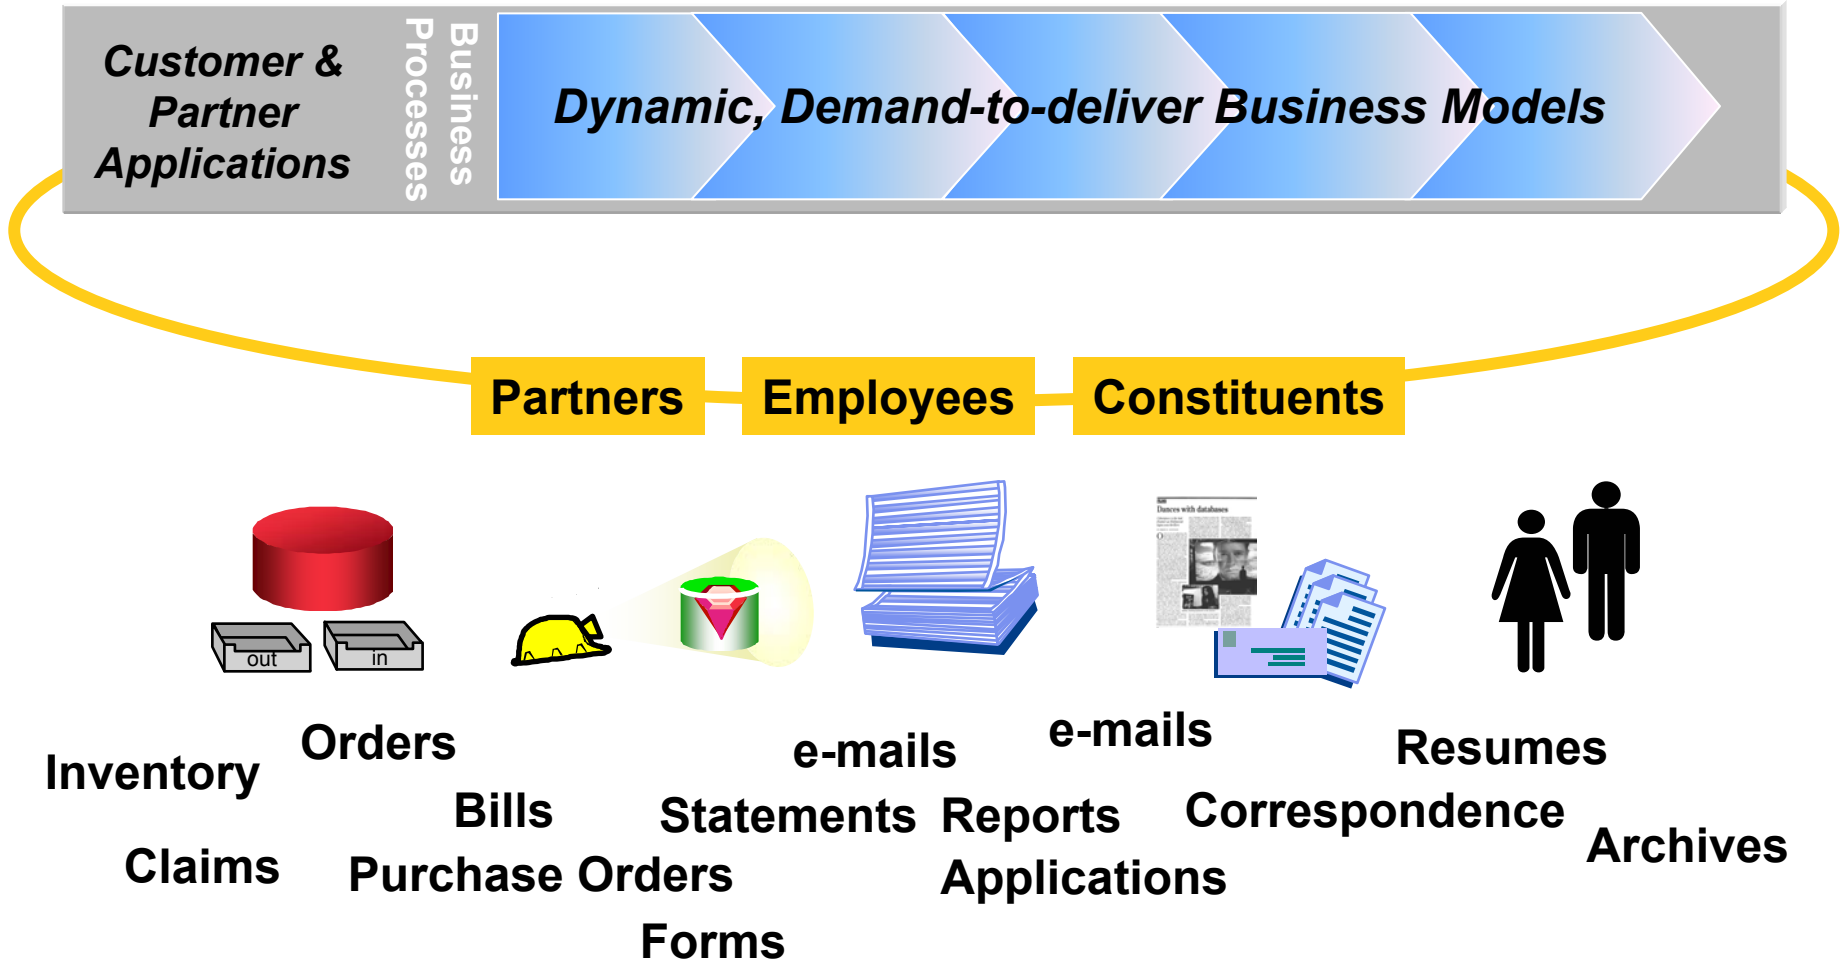


- Reduce cost
- Improve asset utilization
- Address new business opportunities

On Demand Operating Environment



On Demand Requires a Seamless Flow of Information



Information Management Technology Evolution



Capabilities

...From single subject databases to an information infrastructure...



DB2 Information Management

Providing the Best Value to Our Customers

- Information fuels e-business on demand
- Information Management is More Than a Database



Data Challenges

Variety, Velocity, and Volume

- New composite applications need data from multiple sources
 - ▶ Consumers expect holistic, personalized, and value-added content
 - ▶ Relational, XML, packaged applications, content repositories, file systems all contain critical business information

- Increasing emphasis on current data
 - ▶ Real-time analytics
 - ▶ Business activity monitoring

- Petabytes will be the measure of available online data
 - ▶ All client interactions are important (e.g. instant messages, audio records, Web traffic,...)
 - ▶ Internet and intranet content



Complementary Information Integration Approaches

- Consolidate data for local access
 - ▶ Access performance or availability requirements demand centralized data.
 - ▶ Currency requirements demand point-in-time consistency, e.g. close of business
 - ▶ Complex transformation is required to achieve semantically consistent data
 - ▶ Production applications, data warehouses, operational data stores
 - ▶ Typically managed by ETL (Extract, Transform, and Load) or replication technologies

- Integrated access to distributed sources
 - ▶ Access performance and load on source systems can be traded for overall lower cost implementation
 - ▶ Currency requirements demand a fresh copy of the data
 - ▶ Data security, licensing restrictions, or industry regulations restrict data movement
 - ▶ Combining mixed format data, e.g. customer ODS with related contract documents or images
 - ▶ Query requires real-time data, e.g. stock quote, on-hand inventory
 - ▶ Represents an emerging category of technology – Enterprise Information Integration (EII)



What is Enterprise Information Integration?

Enterprise Information Integration refers to a category of middleware which lets applications access diverse and distributed data as if it were a single source, whether or not it is.

It enables the integration of structured and unstructured data


- ▶ to provide real-time read and write access,
- ▶ to transform data for business analysis and data interchange, and
- ▶ to manage data placement for performance, currency, and availability.



Agenda



Announcing: DB2 Information Integrator real time data integration

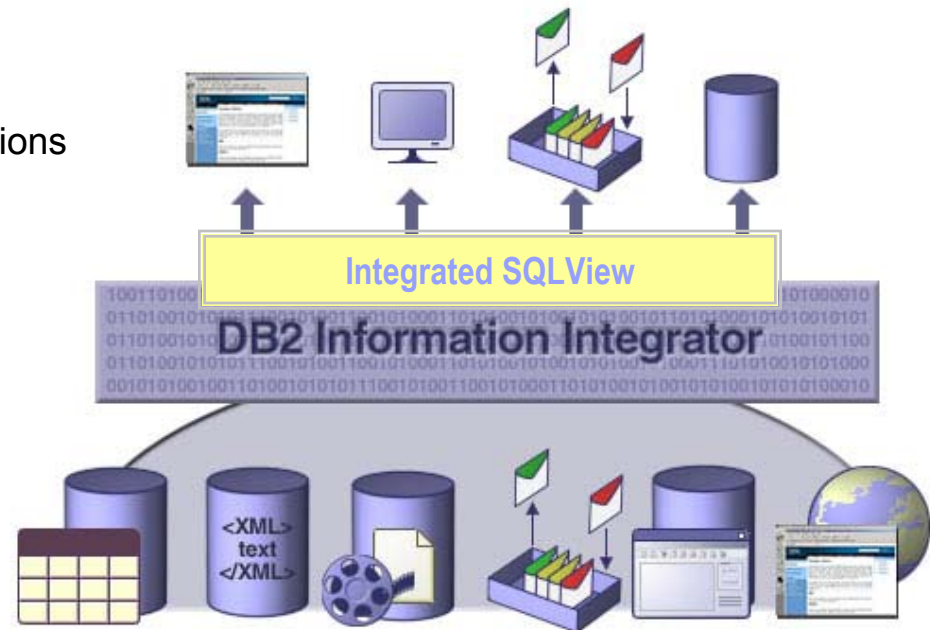
- Information Integration in on demand computing
 - ▶ A fundamental underpinning
- ▶  DB2 Information Integrator 8.1
 - ▶ Federated query server and replication server combined
 - ▶ General availability June 13, 2003
- Pricing and packaging details
 - ▶ Five licensing options
- World-Class Partner Support
 - ▶ Broad range of tool interoperability
- Customer Experiences and Value
 - ▶ Cross-industry and geography



DB2 Information Integrator 8.1

A Federated Data Server – Query distributed data as if it were a single source

- Define integrated view across diverse and distributed data
 - ▶ Wide range of data and content sources
 - ▶ Extensible to virtually any data source
- Query as if a single source
 - ▶ Use standard SQL query and SQL expressions
 - ▶ Include text semantics in the search
 - ▶ Surface specialized functions into SQL
 - ▶ Leverage query optimization and caching
- Compose XML documents
 - ▶ Combine diverse sources
 - ▶ Validate against DTDs or schema
- Publish results to a message queue
 - ▶ Familiar DB programming model
- Single source, relational updates

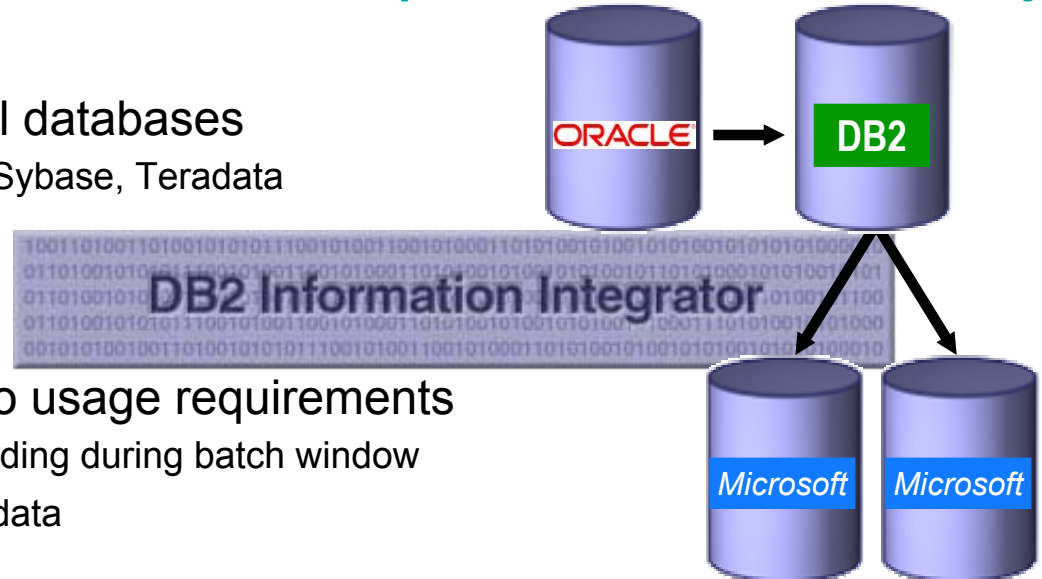


DB2, Oracle, SQL Server, Sybase, Teradata, OLE DB, ODBC, Excel, XML, message queues, Web services, flat files, document repositories, content repositories, LDAP directories, WWW, email databases, and more.

DB2 Information Integrator 8.1

A Replication Server – Manage consolidation for performance and availability

- Distribute data among relational databases
 - ▶ DB2, Informix, Microsoft, Oracle, Sybase, Teradata
- Support flexible topologies
 - ▶ Distribution: One to many
 - ▶ Consolidation: Many to one
- Match data movement modes to usage requirements
 - ▶ Table-at-a-time for warehouse loading during batch window
 - ▶ Transaction-consistent for online data
- Choose latency characteristics
 - ▶ Scheduled, interval-based, continuous
- Apply transformations in-line
 - ▶ Standard SQL expressions or stored procedure execution.



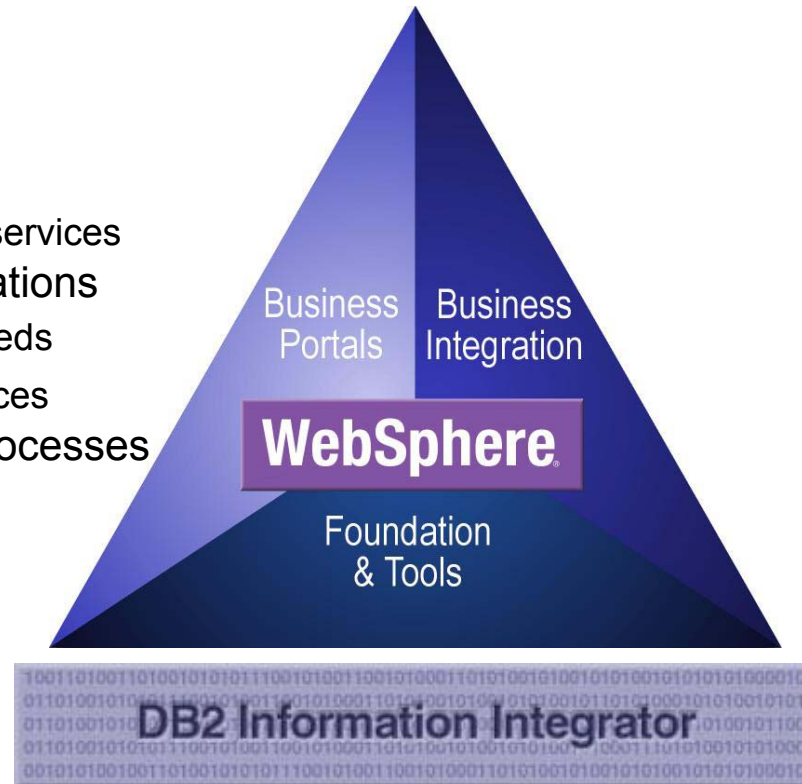
Value to Clients

- Extend current investments
 - ▶ Work within your existing infrastructure
 - ▶ Consolidate data or access distributed data as if it were a single data source
 - ▶ Combine existing data and content assets in new ways
 - ▶ Use familiar SQL programming model and existing tools
 - ▶ Build on a standards-based, strategic integration platform
- Speed time to value for composite applications
 - ▶ Reduce hand-coding 40%-65%
 - ▶ Reduce skill requirements
 - ▶ Reduce development time by half
- Control costs
 - ▶ Reduce payroll costs
 - ▶ Reduce need to rip and replace
 - ▶ Reduce need to manage redundant data

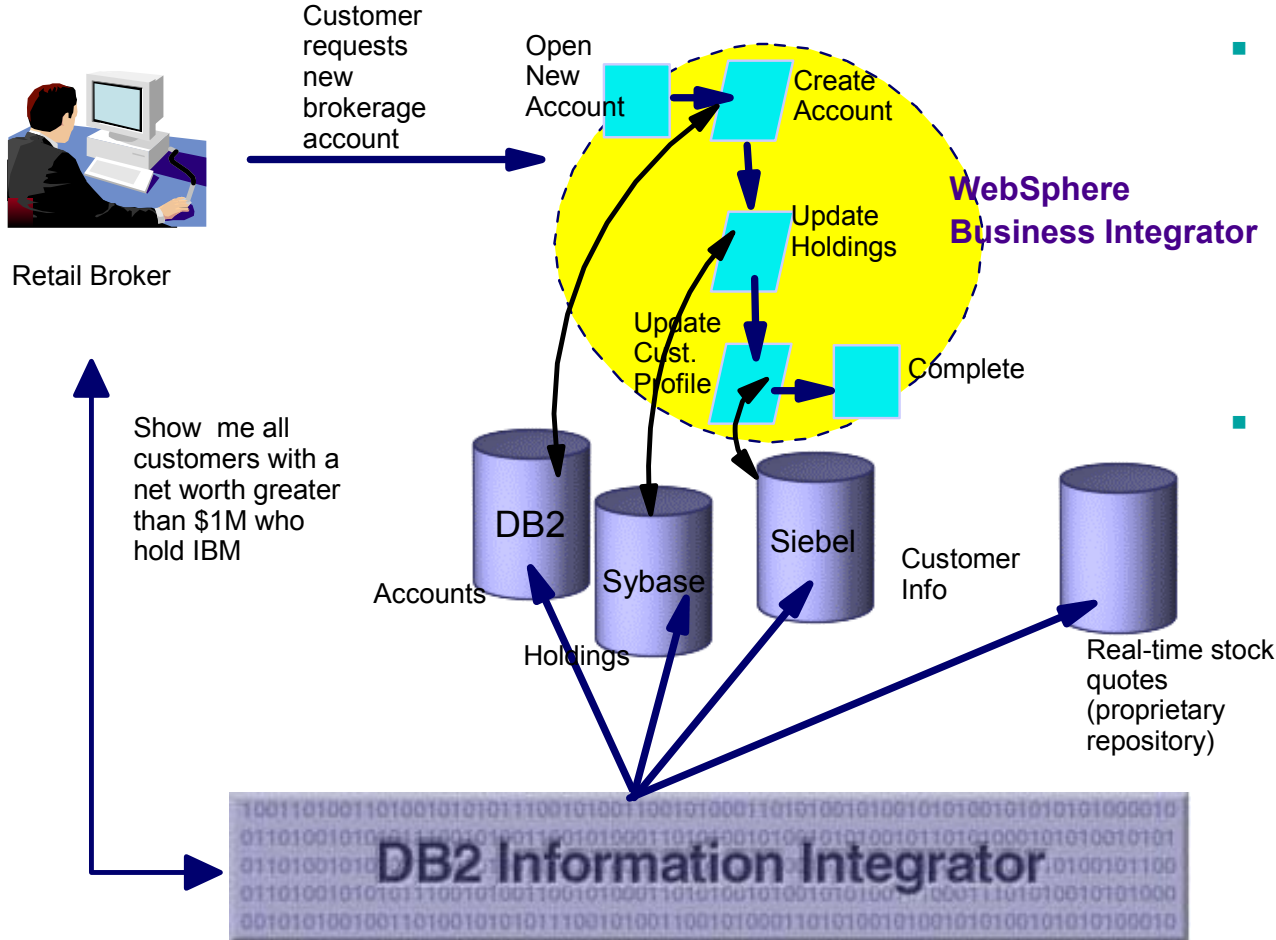


WebSphere Synergy

- Plug into a services oriented architecture
 - ▶ All the power of EII through Web service
 - ▶ Extend reach and transformation richness with Web services
- Improve developer productivity for composite applications
 - ▶ Reduce lines or code, development time, and skill needs
 - ▶ Give applications access to all the relevant data sources
- Reduce cost and complexity of complex business processes
 - ▶ Simplify development for composite business objects
 - ▶ Leverage set-base relational processing
- Enable business activity management
 - ▶ Correlate real-time and historical data
- Increase return on messaging system deployment
 - ▶ Analyze queue traffic with standard analytical tools
 - ▶ Make messaging infrastructure more easily accessible by DB community



Retail Banking Scenario



- **Business Process Integration**
 - ▶ Connects applications, data sources - one at a time
 - ▶ Workflow automates sequence of events
 - ▶ example: opening a new brokerage account

- **Information Integration**
 - ▶ Can look across many information sources simultaneously
 - ▶ example: show me all brokerage customers with a net worth of \$1M, who hold IBM



Agenda



Announcing: DB2 Information Integrator real time data integration

- Information Integration in on demand computing
 - ▶ A fundamental underpinning
- DB2 Information Integrator 8.1
 - ▶ Federated query server and replication server combined
 - ▶ General availability June 13, 2003
- ▶ Pricing and packaging details
 - ▶ Five licensing options
- World-Class Partner Support
 - ▶ Broad range of tool interoperability
- Customer Experiences and Value
 - ▶ Cross-industry and geography



DB2 Information Integrator Packaging

- DB2 Information Integrator, Replication Edition, priced at \$5K/processor
 - ▶ Includes a limited-use data store for use as a replication control server
 - ▶ Plus connectors priced on a per connection basis, at \$15K/connection
- DB2 Information Integrator, Standard Edition, priced at \$20K/processor
 - ▶ Includes a limited-use data store with Net Search Extender for use as a local cache
 - ▶ Plus connectors priced on a per connection basis, at \$15K/connection
- DB2 Information Integrator, Advanced Edition, priced at \$40K/processor
 - ▶ Includes a full function data store with Net Search Extender
 - ▶ Plus connectors priced on a per connection basis, at \$15K/connection
- DB2 Information Integrator, Advanced Edition Unlimited, priced at \$125K/processor
 - ▶ Includes a full function data store with Net Search Extender & unlimited use of available connectors
 - ▶ Break-even on 2-way at 11 connectors, on a 4-way at 23 connectors
- DB2 Information Integrator, Developer Edition, priced at \$999/user
 - ▶ Same content as Advanced Edition, restricted for use in development environments

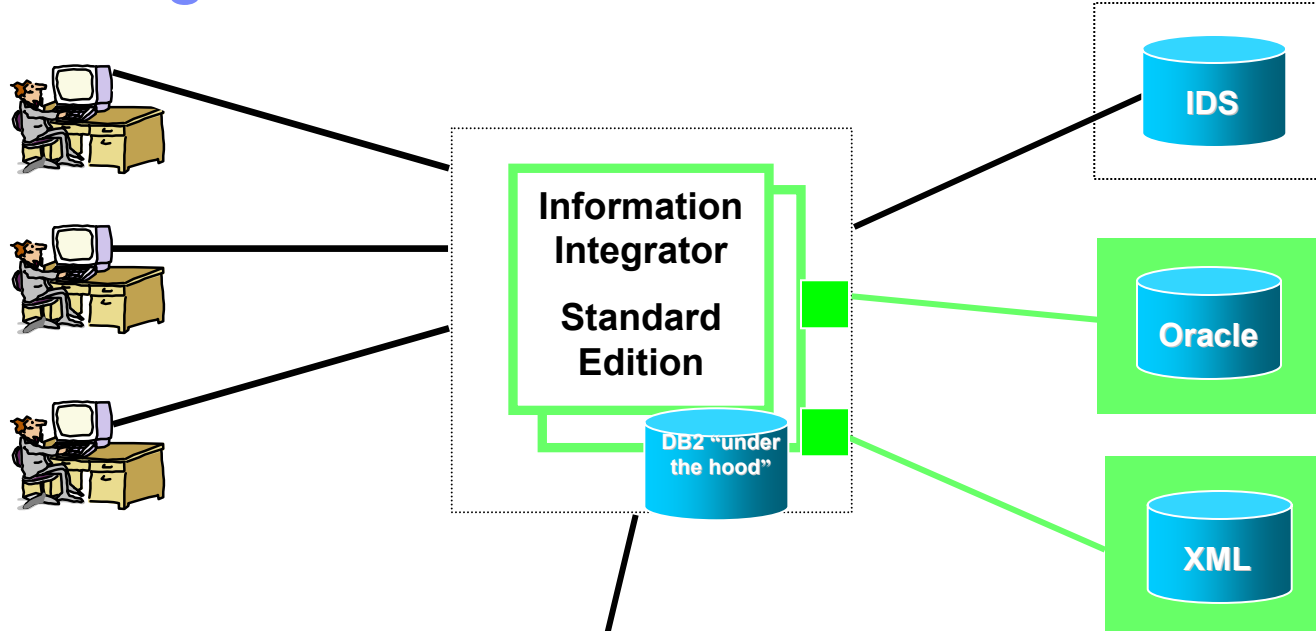


Connector Licensing

- Included with base server
 - ▶ DB2 UDB, Informix IDS, Informix XPS, WebSphere MQ, Web services, and OLE DB
- Unlimited by type
 - ▶ Microsoft Excel - access to one or more Excel documents that reside on one or more servers.
 - ▶ Flat files - access to all flat files that reside on one or more servers.
 - ▶ XML - access to all XML documents that reside on one or more servers.
 - ▶ Extended Search - access to all data sources accessible through Extended Search
- Per instance
 - ▶ Documentum - access to one repository of Documentum documents.
 - ▶ Relational databases - access only one database instance of Oracle, Microsoft SQL Server, Sybase or Teradata databases, or any one data source instance accessed using the ODBC connector.
- Life Sciences
 - ▶ BLAST - access to all data sources accessible through one BLAST daemon.
 - ▶ Entrez - access to all supported data sources accessible through Entrez. Access to PubMed and Nucleotide databases is supported by DB2 Information Integrator.
 - ▶ HMMER - access to all data sources accessible through one HMMER daemon.
 - ▶ BioRS - access to all the databanks supported by a single installation of BioRS.



Sample Configuration: Standard Edition



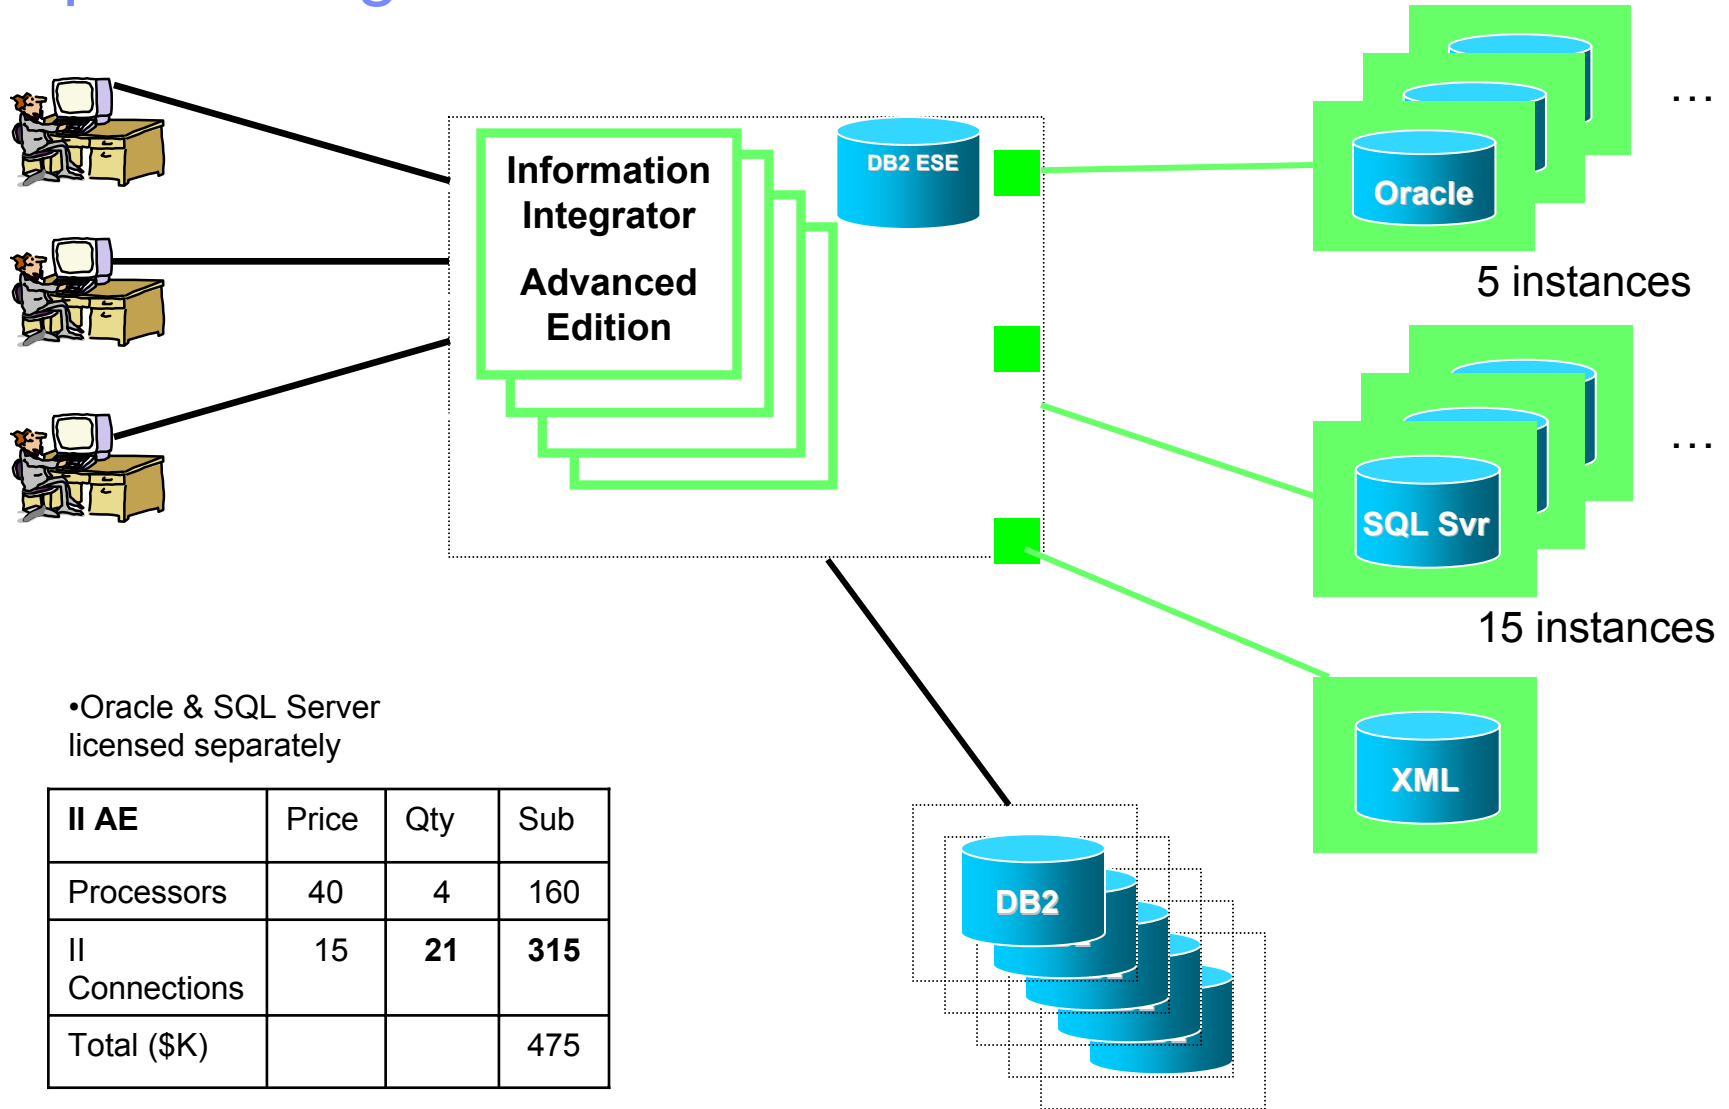
• DB2 UDB “under the hood” limited use

•IDS, DB2 z/OS, Oracle & XML licensed separately

II SE	Price	Qty	Sub
Processors	20	2	40
II Connections	15	2	30
Total (\$K)			70



Sample Configuration: Advanced Edition



•Oracle & SQL Server licensed separately

II AE	Price	Qty	Sub
Processors	40	4	160
II Connections	15	21	315
Total (\$K)			475

Extend Existing Investments


- Migrate existing customers
 - ▶ Entitlement under software maintenance for DB2 DataJoiner, DB2 Relational Connect, DB2 Life Sciences Data Connect, IBM Discovery Link
 - ▶ Migration scenarios exercised in beta
 - ▶ Migration services available
 - ▶ Trade-up from DB2 UDB available
- Significant new function for migrating customers
 - ▶ Performance improvement
 - 80% of tests were faster than DataJoiner
 - Half of those twice as fast!
 - Plus federated caching
 - ▶ Usability improvement
 - Discovery and auto-mapping features
 - Replication center enabled for mixed environment
 - ▶ New sources
 - Content rich access
 - Extensibility to virtually any source
 - ▶ New SQL
 - ▶ Write access (for Relational Connect base)



Agenda



Announcing: DB2 Information Integrator real time data integration

- Information Integration in on demand computing
 - ▶ A fundamental underpinning
- DB2 Information Integrator 8.1
 - ▶ Federated query server and replication server combined
 - ▶ General availability June 13, 2003
- Pricing and packaging details
 - ▶ Five licensing options
- ▶  World-Class Partner Support
 - ▶ Broad range of tool interoperability
- Customer Experiences and Value
 - ▶ Cross-industry and geography



World-Class Partner Ecosystem

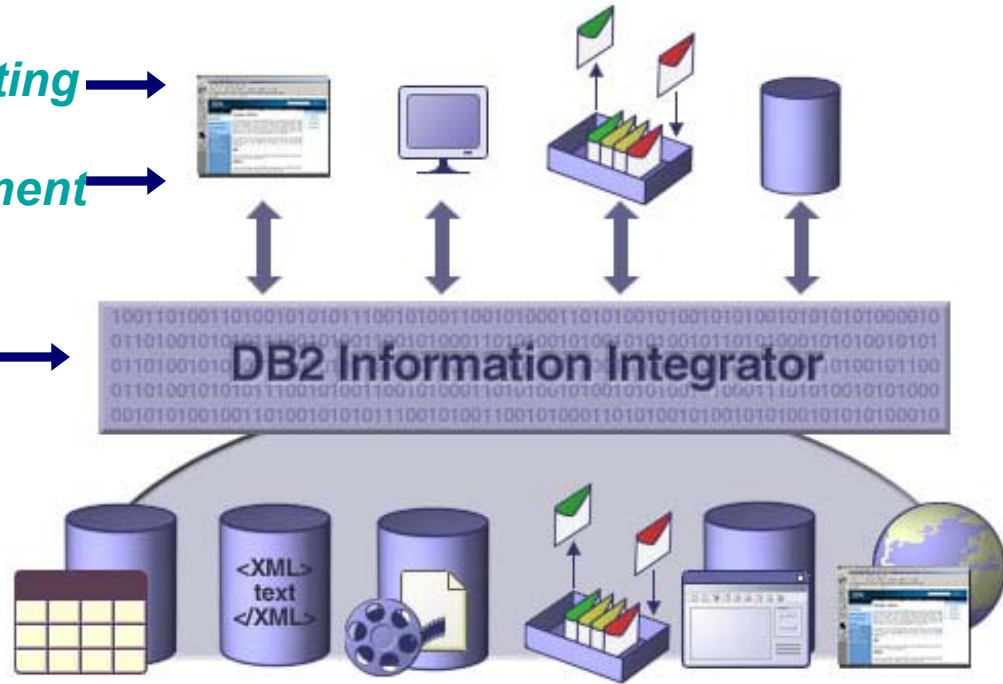
Enhance analysis and reporting →

Enable application development →

Enrich function →

Extend reach →

Expedite deployment ← →



Enhance Analysis and Reporting



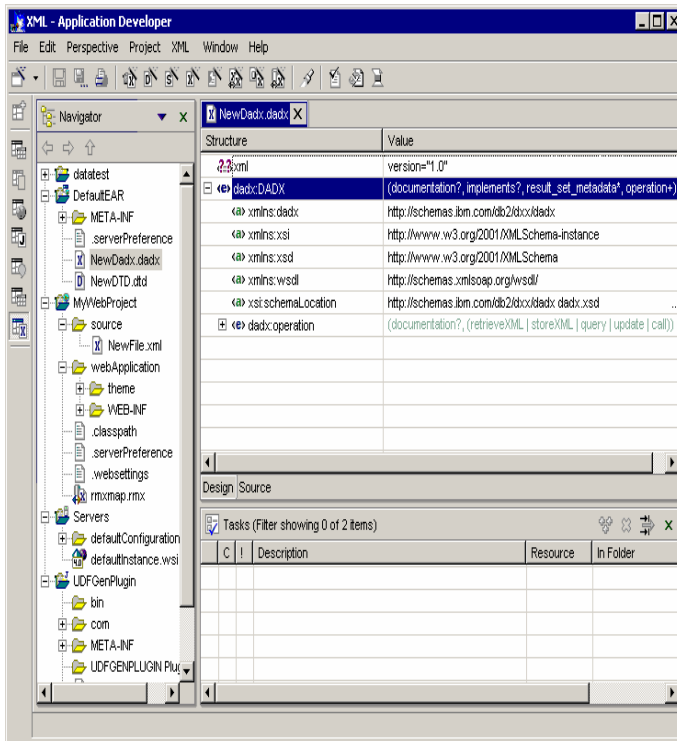
- *"..exceptionally fast and efficient federated querying"*
- *"..completely transparent.."*
- *"..opens new opportunities.."*
- *"..allows users access to important data that was often forgotten or used infrequently.."*
- *"..providing integrated on-demand analysis and execution.."*
- *"..an open, extensible, and consistent information integration framework.."*



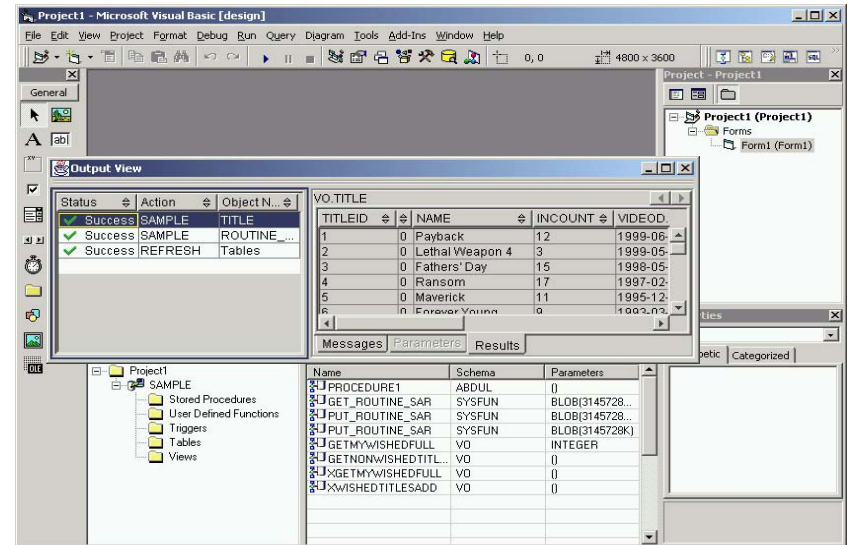
Enable Application Development

Developer productivity for both worlds: Java and .NET

WebSphere Studio



Microsoft Visual Studio



- Reduce hand coding 40%-65%
- Maintain performance

IBM In House Experiment

- Understand trade-offs of building Web components spanning multiple data sources
 - ▶ Build servlets, session EJBs, CMP entity EJBs
 - With DB2 Information Integrator for data access
 - With hand-coded direct data access
- Application scenario
 - ▶ Corporate merger of multiple distributors
 - ▶ Access to 2 RDBMSs (DB2, Oracle), 1 non-RDBMS (Excel)
 - UNION ALL views, nicknames defined in DB2 II
 - Data derived from TPC-H
 - ▶ Multiple queries (some derived from TPC-W, TPC-H)
 - Catalog search
 - Post-merger analysis
 - ▶ Some overlapping info across sources (e.g., same customer placing different orders through different distributors)



Findings

- DB2 II simplified server-side Java development
 - ▶ Reduced coding efforts by 40 - 65%
 - ▶ Reduced skills requirement
 - No need to manually decompose queries
 - No need to manage multiple connections
 - No need to worry about different APIs, data types
 - ▶ Maintained reasonable performance
- Other benefits of DB2 II
 - ▶ Extends reach of WebSphere App Server (e.g., CMP EJBs for "unusual" data sources)
 - ▶ Supports new object models (e.g., CMP EJBs whose attributes span multiple databases)
 - ▶ Minimizes DBMS client software installation on Java IDE platforms, WebSphere App Server



Enrich Function



Profit from Intelligent Information™

- Complementary ETL, Data Profiling, Transformation, Data Quality, Metadata Management



- Complementary ETL, Data Profiling, Transformation, Data Quality, Metadata Management



- Link disparate representations of customer information in multiple data sources
- Enhance DB2 II capabilities to include Customer Data Integration



- Construct and layer XML views over federated data with easy-to-use GUI
- XML-based query
- Support for caching XML views
- Reconciliation of keys in disparate sources



Extend Reach



- CrossAccess eXadas enables read/write access VSAM, IMS, ADABAS, CA-IDMS, or CA-Datacom/DB



- Information Builder's I-Way adaptors enable read/write access to more than 250 application and data server sources



- Neon Shadow enables read access to IMS, VSAM, ADABAS



- Striva DETAIL enables read access to VSAM, IMS, ADABAS, CA-IDMS, or CA-Datacom/DB, HP Turbo Image

Agenda



Announcing: DB2 Information Integrator real time data integration

- Information Integration in on demand computing
 - ▶ A fundamental underpinning
- DB2 Information Integrator 8.1
 - ▶ Federated query server and replication server combined
 - ▶ General availability June 13, 2003
- Pricing and packaging details
 - ▶ Five licensing options
- World-Class Partner Support
 - ▶ Broad range of tool interoperability
- ▶ Customer Experiences and Value
 - ▶ Cross-industry and geography



ITERGO



Vision

- ▶ **ITERGO** is part of the **ERGO** insurance group, the second largest insurance company in Germany. They want to provide better service to their customers by improving their IT and business intelligence applications.

Challenge

- ▶ Simplify access to disparate data for inventory management and customer care.

Solution

- ▶ Federated access over Oracle inventory scan database and DB2 employee and inventory basic and planning database for integrated view of asset and contact information.
- ▶ Federated access across disparate warehouses and marts for unified client view .

Value

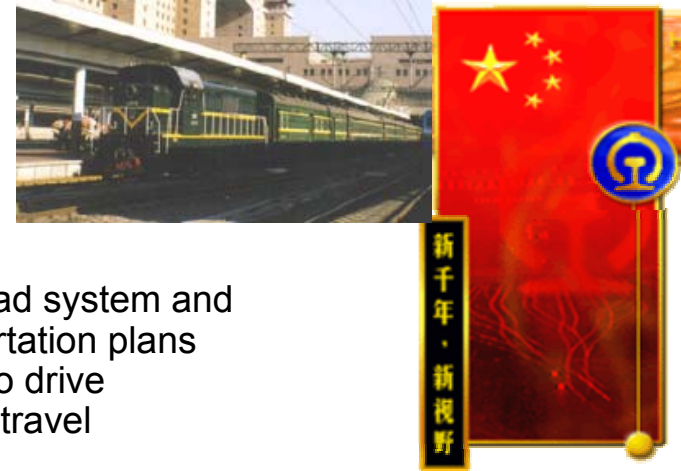
- ▶ Simplify asset management and comparison between target/actual inventory.
- ▶ Improve customer services without extensive changes to their existing infrastructure.

“DB2 Information Integrator is a strategic product for all cases where different database systems needs to be federated. ”

Thomas Dalewski, **ITERGO**



Ministry of Railways (China)



Vision

- ▶ Ministry of railways manages the Chinese the national railroad system and is responsible for defining and implementing railway transportation plans
- ▶ Want to provide an organization-wide intelligence platform to drive decision making that will increase freight and passenger rail travel

Challenge

- ▶ Data needed for effective decision-making is spread across multiple databases
- ▶ Begin solution rollout with limited skill and resource

Solution

- ▶ Give business managers access to summarized historical data in a DB2 based warehouse with federated access to real-time DB2 and Oracle freight transportation operational data and passenger data in Sybase and SQL Server.

Value

- ▶ Improve the quality of railway transportation management
- ▶ Respond more quickly to market demand by improving freight forecasting
- ▶ Optimize carrying resource assignment

"With its real time federation capability DB2 Information Integrator significantly enhances the power of the data warehouse supporting our intelligent transportation management systems and other enterprise applications. The XML and unstructured data support of this product will help us fulfill the reporting requirements to the states administrative bureaus."

Mr. Sun Yuan-Yun, Chief Engineer, Ministry of Railways, People's Republic of China

Quad Graphics



Vision

- ▶ It is Quad Graphics mission to be the very best at putting ink on paper based on achieving the best fit with client needs.

Challenge

- ▶ Incrementally migrate from legacy Sybase infrastructure to DB2 infrastructure.

Solution

- ▶ Provide transparent access to Sybase and DB2 databases for migrating applications using DB2 Information Integrator

Value

- ▶ Enable incremental application migration
- ▶ Works seamlessly with established PowerBuilder tools
- ▶ Anticipated strategic value in accessing heterogeneous data environment



Summary

- Information integration addresses a key pain point for customers
- IBM leads the industry in integrating information for on demand computing
- DB2 Information Integrator represents a distinct difference in approach between IBM and Oracle and Microsoft
 - ▶ Commitment to open architectures that extend current enterprise investments
 - ▶ Strategic vision for information on demand
- IBM is developing a world-class ecosystem of partners that enhance and exploit its capabilities
- Customers see real value in their deployments across diverse scenarios
 - ▶ Control costs
 - ▶ Speed time to value
 - ▶ Extend current investments



More Information

The screenshot shows the IBM website for the DB2 Information Integrator Product Family. The browser window title is "DB2 Information Integrator Product Family - Family Overview - IBM Software - Microsoft Internet Explorer". The address bar shows the URL "http://www-3.ibm.com/software/data/integration/". The page features the IBM logo, a search bar, and navigation links for Home, Products & services, Support & downloads, and My account. The main content area is titled "Information Integration" and includes a description of the product family, a "Select a Product" dropdown menu, and sections for Products, Solutions, and Highlights. The Products section lists "DB2 DataJoiner" and "DB2 Warehouse Manager". The Solutions section lists "IBM DiscoveryLink" and "WebSphere Business Integration". The Highlights section includes a link to "What customers need to know about existing products".

DB2 Information Integrator Product Family - Family Overview - IBM Software - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Favorites History Size Print Print Preview

Address http://www-3.ibm.com/software/data/integration/ Go Links >>

Google Search Web Search Site New! PageRank Page Info Up Highlight

IBM Search

Home Products & services Support & downloads My account

→ Select a country

Software > Database and Data Management >

Information Integration

DB2 Information Integrator is a new family of products, designed to address customer requirements for integrating structured, semistructured and unstructured information effectively and efficiently. These products are based on the previously disclosed Xperanto project. They provide the foundation for a strategic information integration framework to help access, manipulate, and integrate diverse, distributed and real-time data.

Select a Product

Announcing:
DB2 Information Integrator
real time data integration

Products

- DB2 DataJoiner**
makes diverse and distributed relational databases appear to be a single database and supports replication among mixed relational databases

Solutions

- IBM DiscoveryLink**
Query access to existing databases, applications, and search engines
- WebSphere Business Integration**
Build an integration infrastructure to

Highlights

INFORMATION INTEGRATION
Technology Preview

[What customers need to know about existing products](#)

DB2 Information

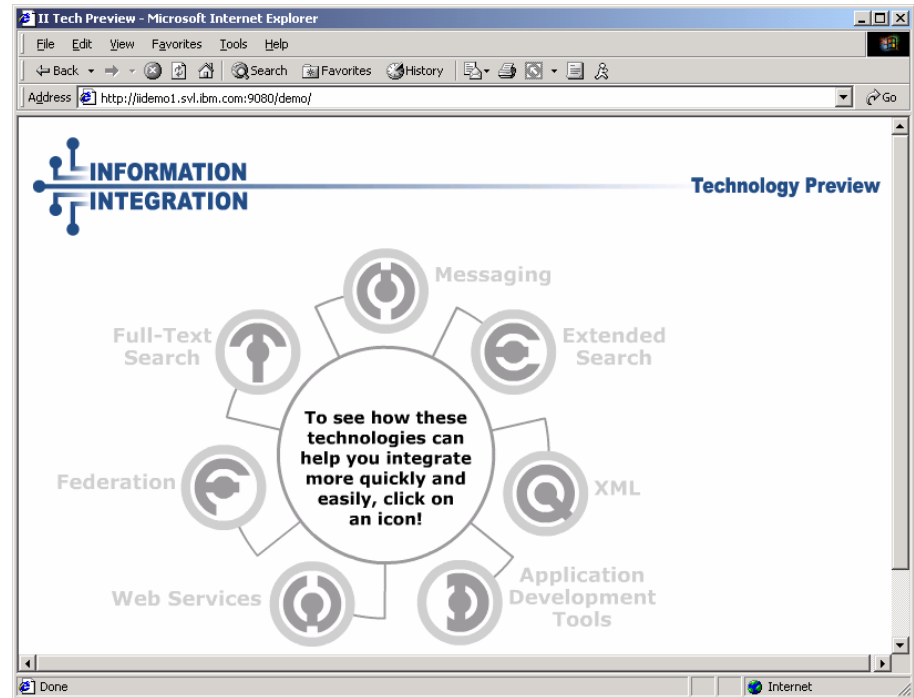
Internet

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

More information



<http://db2ii.dfw.ibm.com/demo/>



<http://www.research.ibm.com/journal/sj41-4.html>

