



IBM Software Group

Previewing Masala

May 2004

DB2 Information Management Software



Nelson Mattos

Distinguished Engineer and Director, Information Integration

Holly Hayes

Program Director, Information Integration

Kuo-Wei Hwang

Program Director, Information Integration Solutions and Deployment

 business on demand software

Agenda



DB2 Masala Technology Preview

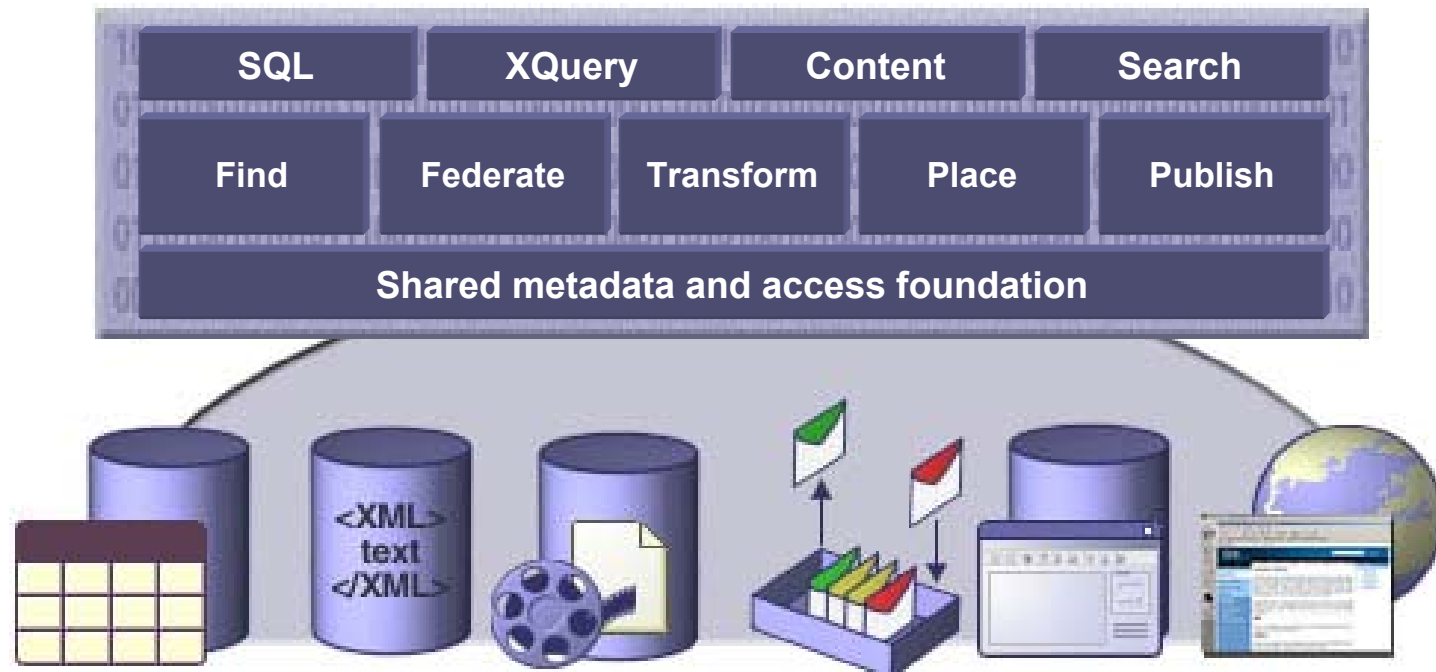
- DB2 Information Integrator
 - ▶ Vision
 - ▶ Function
 - ▶ Value

- New capabilities coming in Masala
 - ▶ Finding the most relevant corporate information
 - ▶ Speeding time to market with federated access capabilities
 - ▶ Adding value to the business integration environment

- Delivering tangible value from partnerships
 - ▶ Additive value and support for client's investments
 - ▶ Product integration and testing
 - ▶ Skilled solution deployment

DB2 Information Integrator Vision

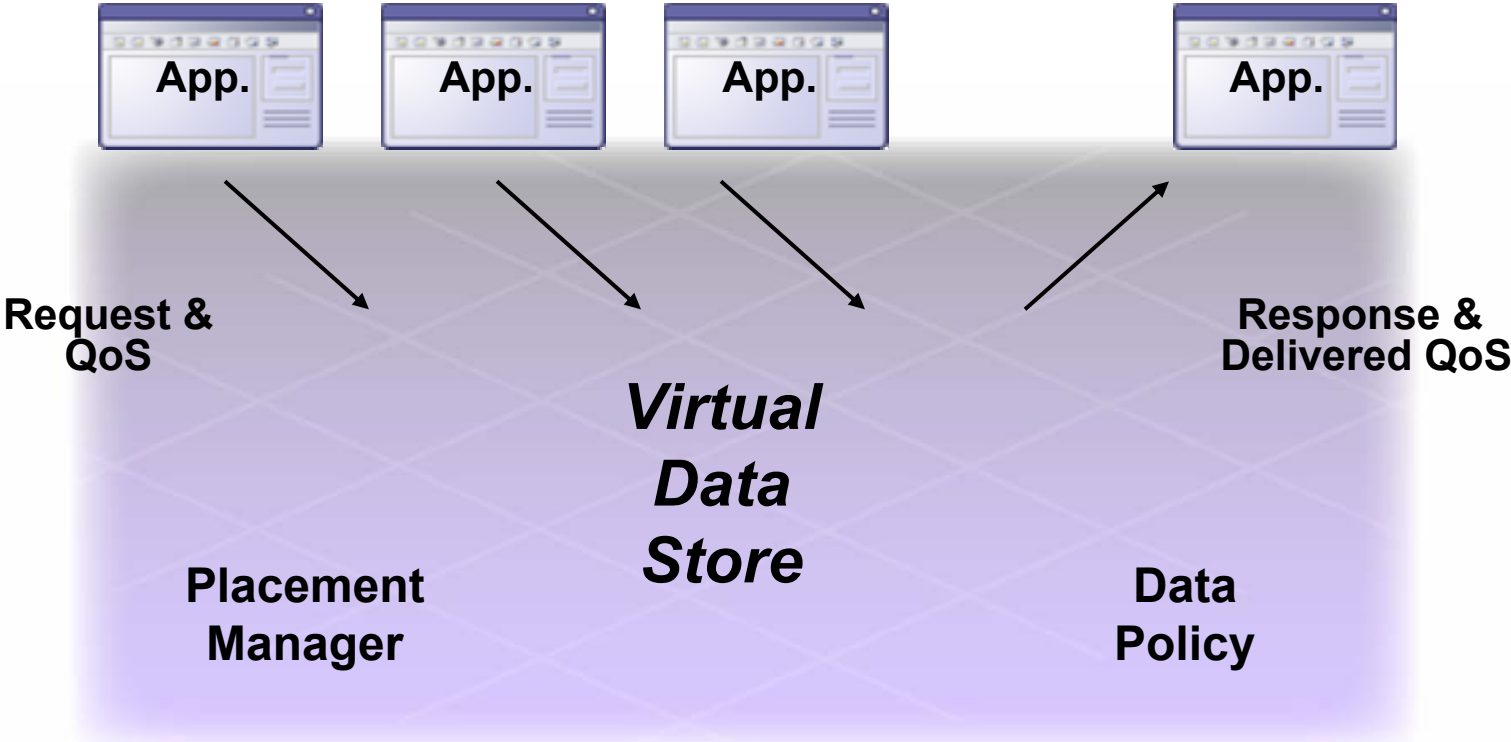
Any data -- Multiple access paradigms -- Multiple integration disciplines



DB2 Information Integrator Vision

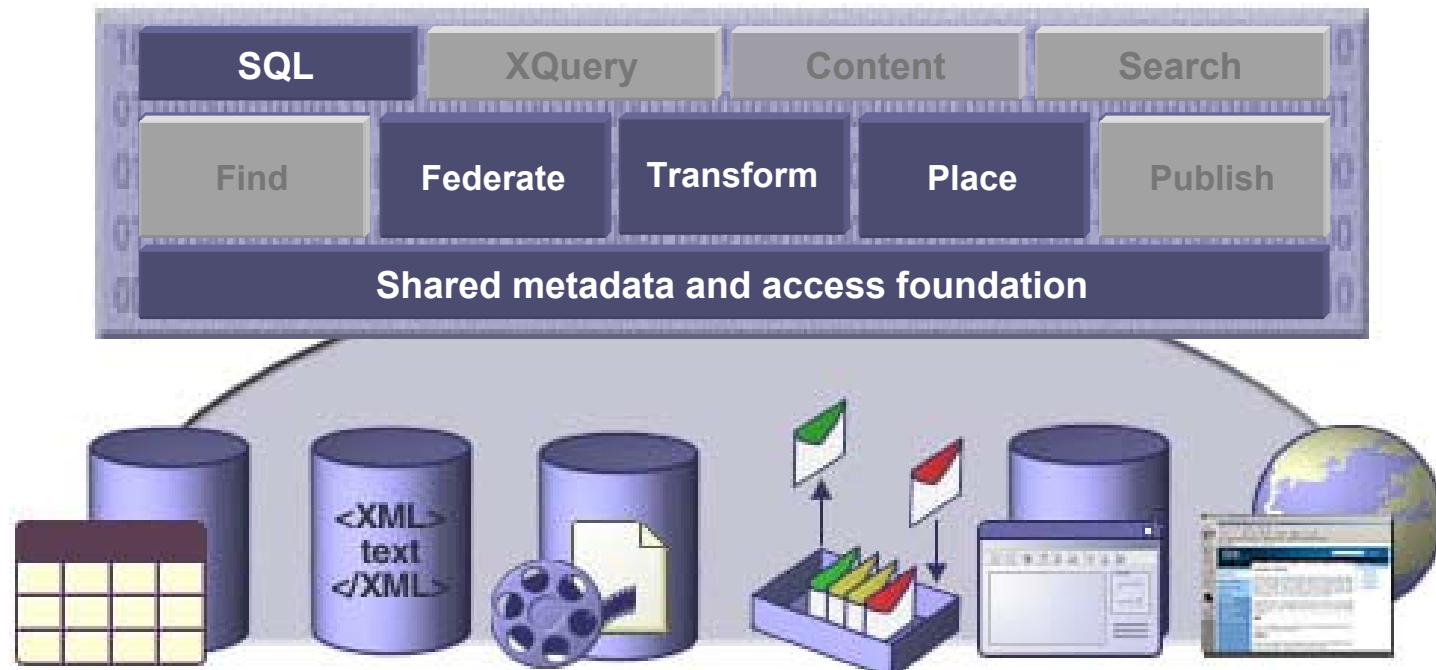
Any data -- Multiple access paradigms -- Multiple integration disciplines

Data virtualization



DB2 Information Integrator (DB2 II) 8.1

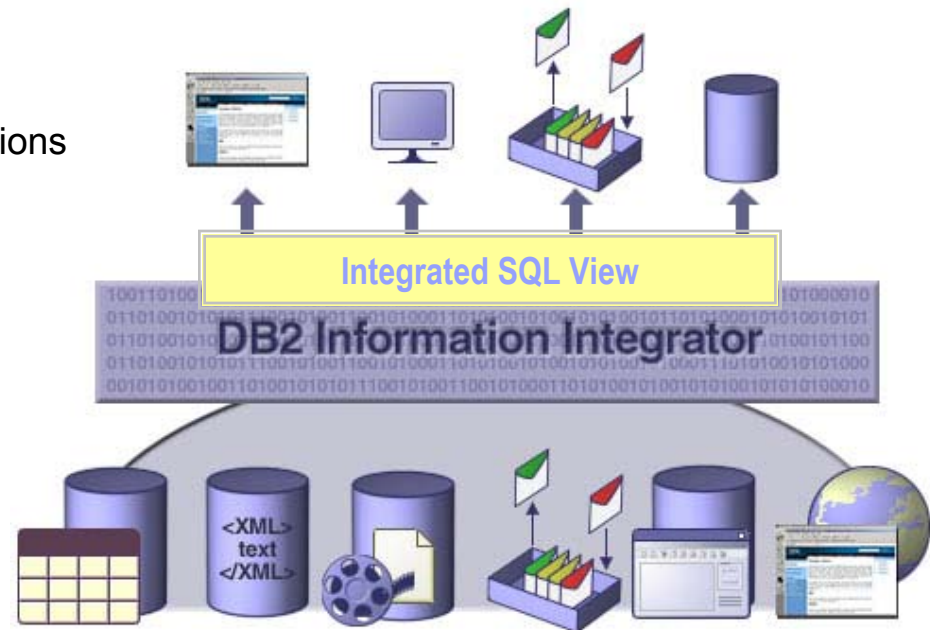
- Integrates structured and unstructured data
 - ▶ to provide real-time read and write access,
 - ▶ to transform data for business analysis and data interchange, and
 - ▶ to manage data placement for performance, currency, and availability



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, Informix, 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

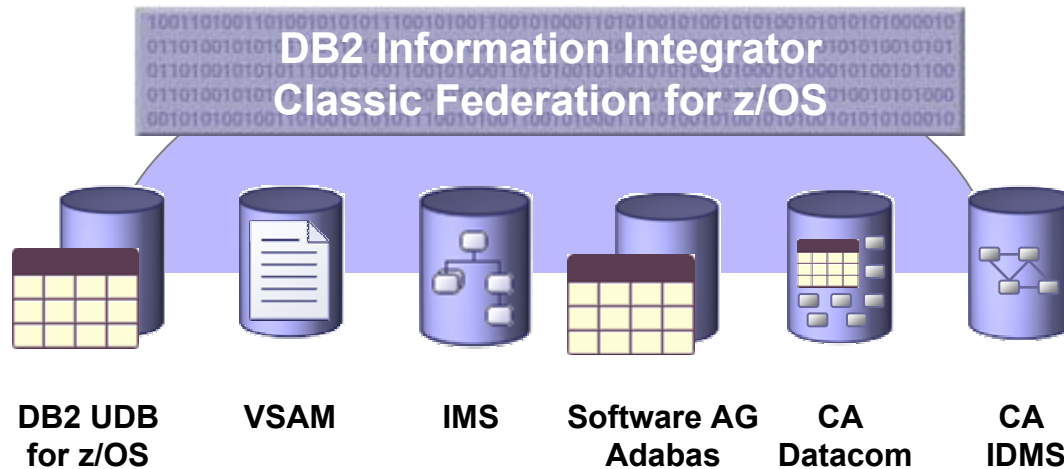
A Replication Server – Manage placement for performance and availability

- Distribute data among relational databases
 - ▶ DB2, Informix, Microsoft, Oracle, Sybase, Teradata
- Support flexible topologies
 - ▶ Distribution: One to many
 - ▶ Consolidation: Many to one
 - ▶ Column and row subsetting
 - ▶ Multiple flavors, e.g. history tables, key changes
- 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, event-driven, continuous
- Apply transformations in-line
 - ▶ Standard SQL expressions or stored procedure execution.

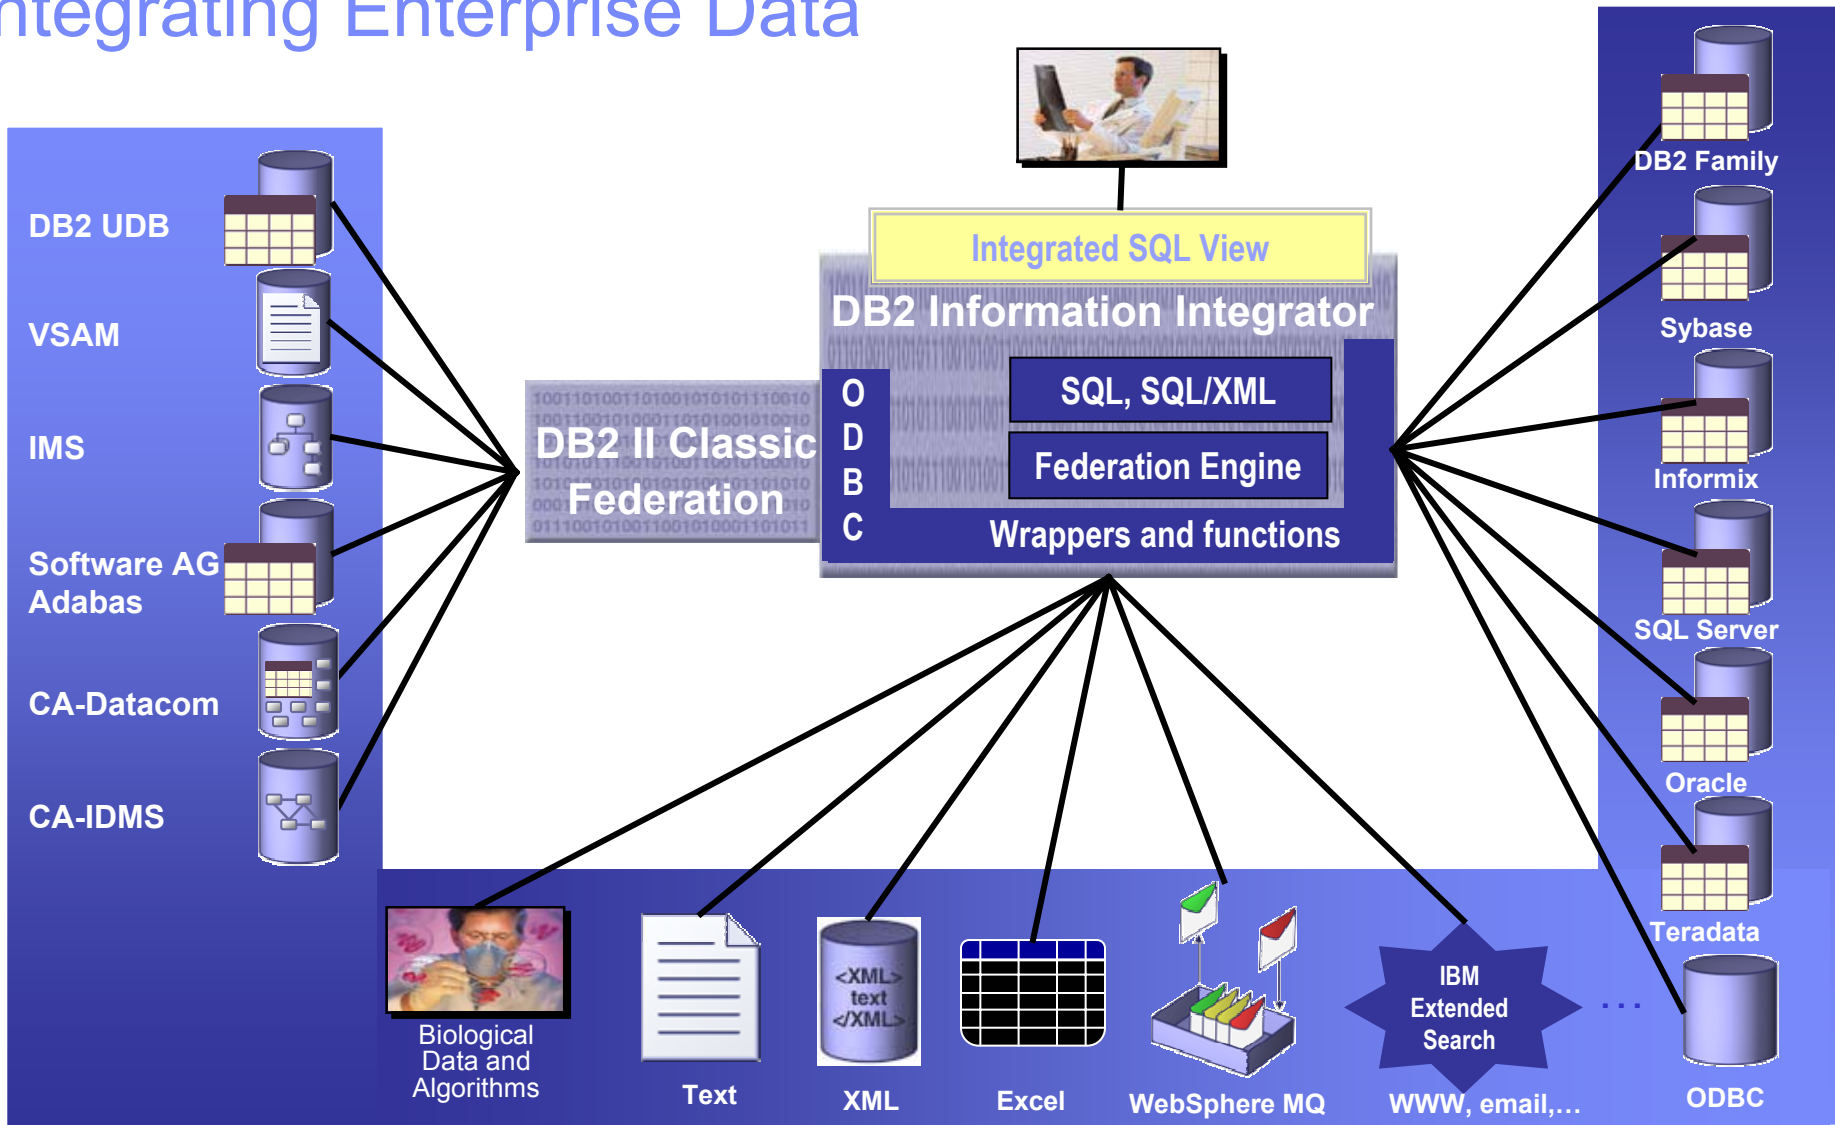


DB2 Information Integrator Classic Federation

- SQL-based read and write access to mainframe data sources
- Metadata-driven for easy configuration and maintenance
 - ▶ No mainframe programming required
- Multi-threaded with native drivers for scalable performance
- Designed for your existing mainframe, tools, and application infrastructure



Integrating Enterprise Data



Enhancing Business Intelligence

- Extending the Warehouse
 - ▶ Extend reporting with real-time data, unstructured content, or cross-warehouse views
 - ▶ Prototyping changes or new requirements
- Regulatory Compliance and Monitoring
 - ▶ Basel II Solution - Enable continuous monitoring
 - ▶ Sarbanes Oxley - Check exceptions in real-time
 - ▶ Risk reporting and analysis for Financial Markets - Access to real-time results, contracts and agreements
- Mergers and Acquisitions
 - ▶ Bridge existing data marts
 - ▶ Staged migration to common warehouse
- Warehouse Maintenance
 - ▶ Trickle-feed for lower latency snapshots, continuous online access, or managing network constraints



Extending Business Integration

- Customer Data Integration
 - ▶ Feed the ODS from diverse data sources
 - ▶ Augment customer views with additional attributes that are seldom needed, must-be-current, or very large
- Business Performance Management
 - ▶ Access and integrate KPIs, planned targets, historical trends, and current values
- Application Integration for Data Consistency
 - ▶ Replicate data between application databases
 - ▶ Support consolidation and distribution between HQ and branch/store locations
- Portal Development
 - ▶ Access disparate sources through unified views for Web deployment




Value to Clients (prioritized by Leadership Board)

1. Speed time to market for new applications
 - Reduce hand-coding and development time by half
2. Extend the value of existing investments
 - In mainframe and distributed, structured and unstructured assets
 - In business intelligence and business integration infrastructures
 - In skills and tools that understand the SQL programming model
3. Gain new insights
 - Leverage real-time data
 - Combine data and content in new ways
3. Control costs
 - Increase productivity and reduce redundant data
5. Deliver tailored views
 - Specific to the application or security requirements
6. Build on a standards-based, strategic integration platform
 - Consolidate data or access distributed data as if it were a single data source



Agenda

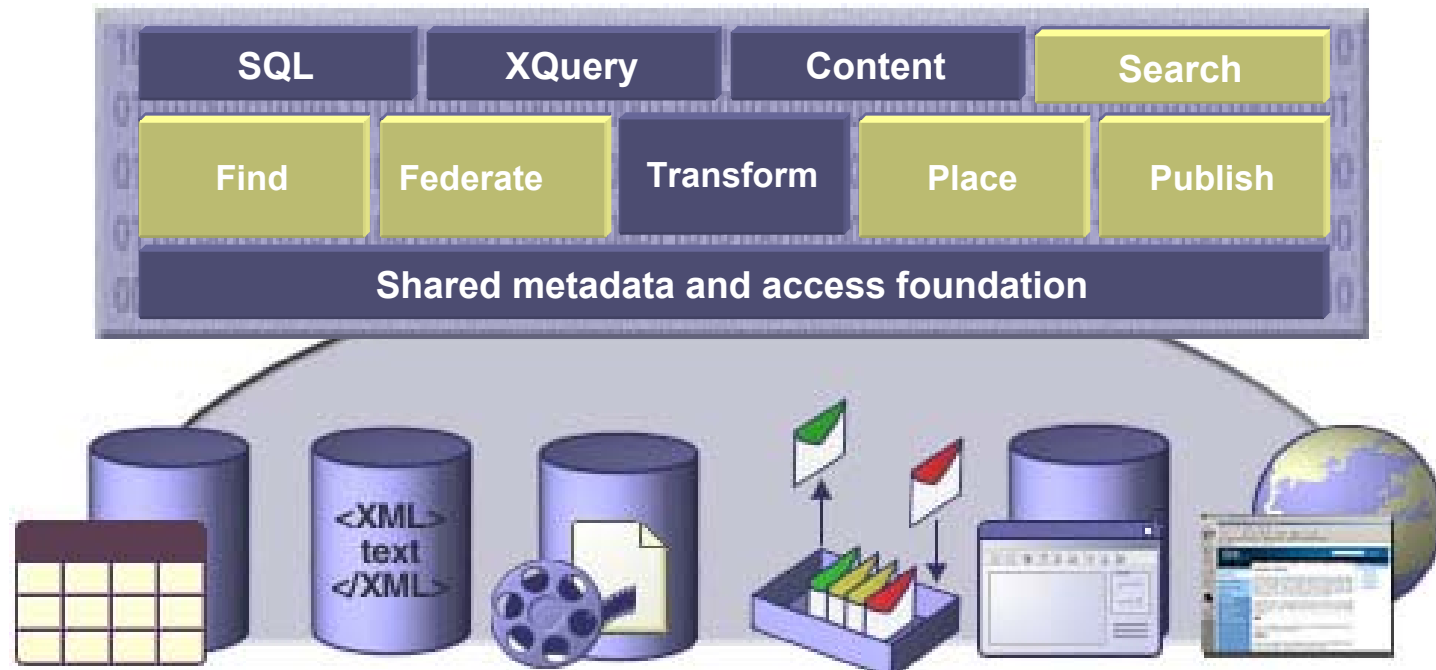


DB2 Masala Technology Preview

- DB2 Information Integrator
 - ▶ Vision
 - ▶ Function
 - ▶ Value
- New capabilities coming in Masala
 - ▶ Finding the most relevant corporate information
 - ▶ Speeding time to market with federated access capabilities
 - ▶ Adding value to the business integration environment
- Delivering tangible value from partnerships
 - ▶ Additive value and support for client's investments
 - ▶ Product integration and testing
 - ▶ Skilled solution deployment

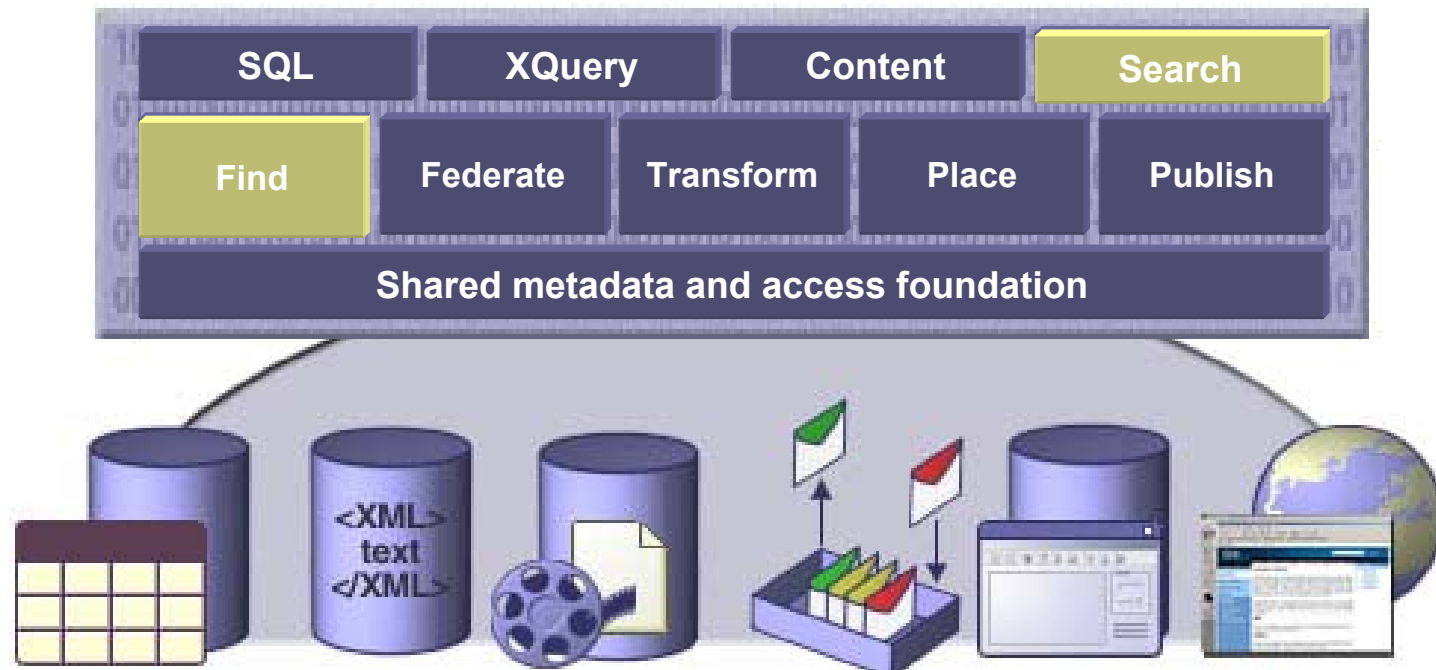
What's new in Masala

- Finding the most relevant corporate information
- Speeding time to market with federated access capabilities
- Adding value to the business integration environment



What's new in Masala

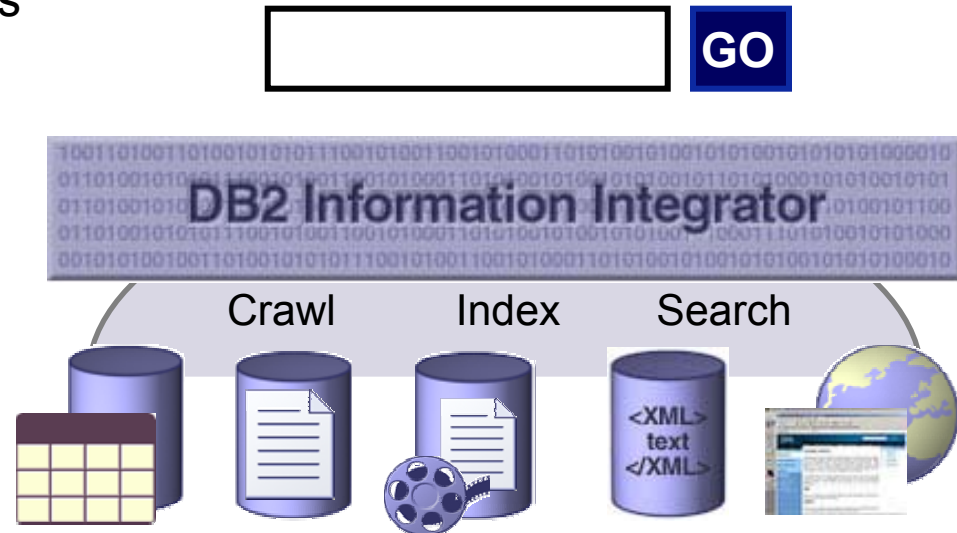
- Finding the most relevant corporate information



Enterprise Search Middleware

Masala powers intranets, extranets, and corporate portals

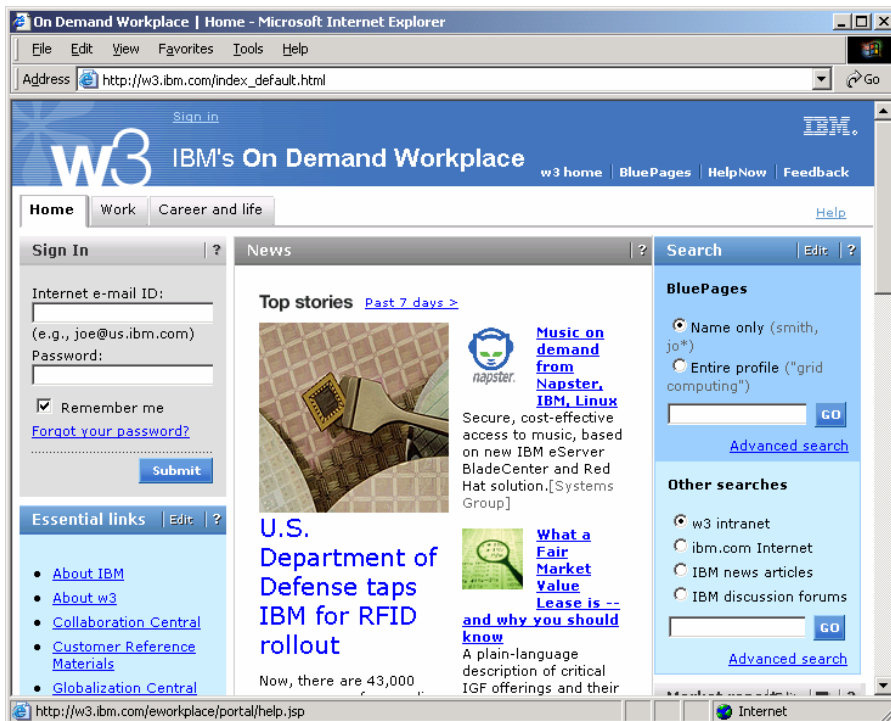
- Delivers the best results with sub-second response
 - ▶ Sophisticated relevancy algorithms for corporate content
- Scales for large collections or enterprises
 - ▶ 500K documents and above
 - ▶ 1000s of concurrent users
- Fits easily into enterprise applications
 - ▶ Java APIs
 - ▶ Document level security
- Eases administration and maintenance
 - ▶ Analysis features all under-the-covers



HTTP/HTTPS, News groups (NNTP), File systems, Domino databases, Microsoft Exchange public folders, DB2 Content Manager, DB2 UDB, Informix, Oracle

Proven Quality, Scale and Robustness on IBM Intranet

A test bed to bring the right product to market



- Quality – preferred 2:1 over prior technology
- Scale -- 80K queries/day with sub-second response over 7M pages
 - ▶ Indexes
 - 7 M unique pages
 - 10,000 websites
 - 20 K per document
 - ▶ Processes
 - 80 K queries/day
 - 7 K queries/hour peak
 - Stressed to 10x higher
- Robustness -- 99.9% availability since Sept.
 - ▶ 24x7 operation

Extending focus for information integration

Become the preferred provider for search services in enterprise applications

- Focus on search and text analysis as a core competency
 - ▶ Strategic importance for many IBM products, services, and new opportunities
 - ▶ Key extension to information integration capabilities

- Embed basic features and upgrade to richer technology
 - ▶ Embedded in Email, Knowledge Management, Portals, Databases, Content Management, E-Commerce, Enterprise Search
 - ▶ Upgrade for broader enterprise content access, greater scalability, and rich analytics

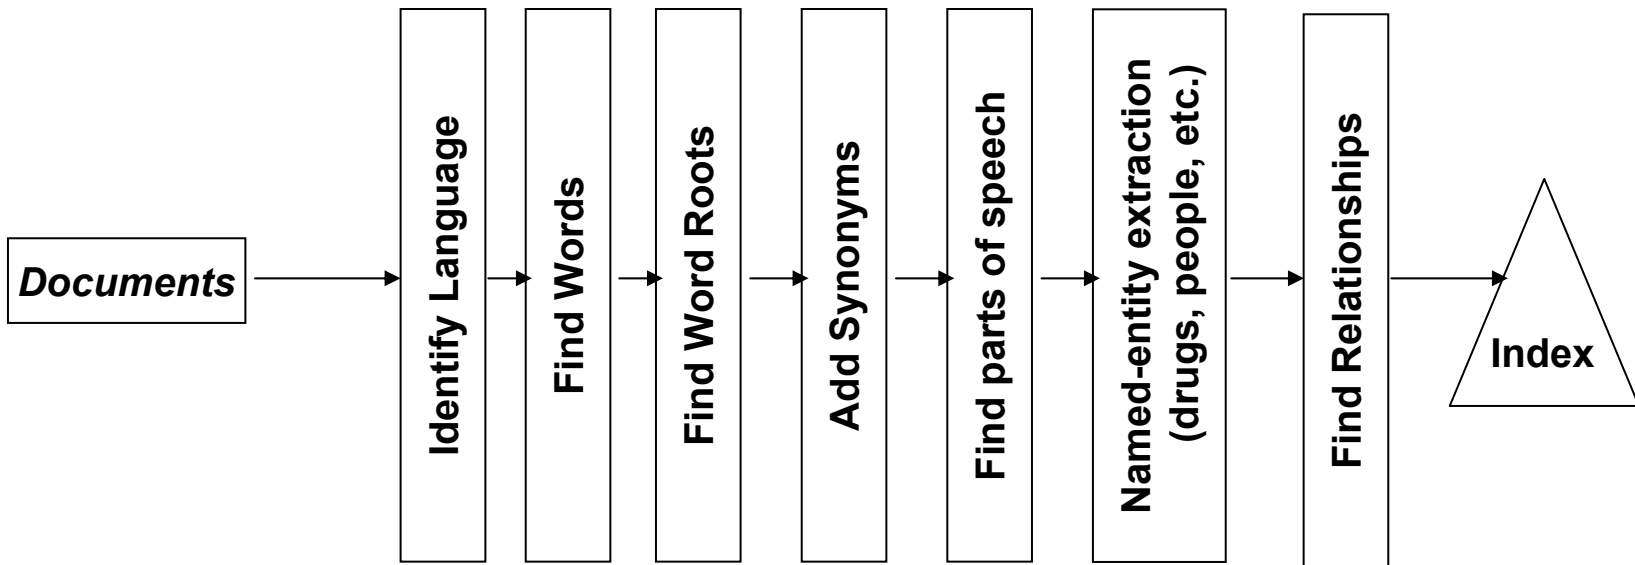
- Provide an architecture to help businesses fully exploit their content assets
 - ▶ Fostering a marketplace around UIMA that delivers new analytical functions
 - ▶ Accelerating custom solution delivery through plug and play application composition
 - ▶ Populating structured sources with extracted semantic content



Unstructured Information Management Architecture (UIMA)

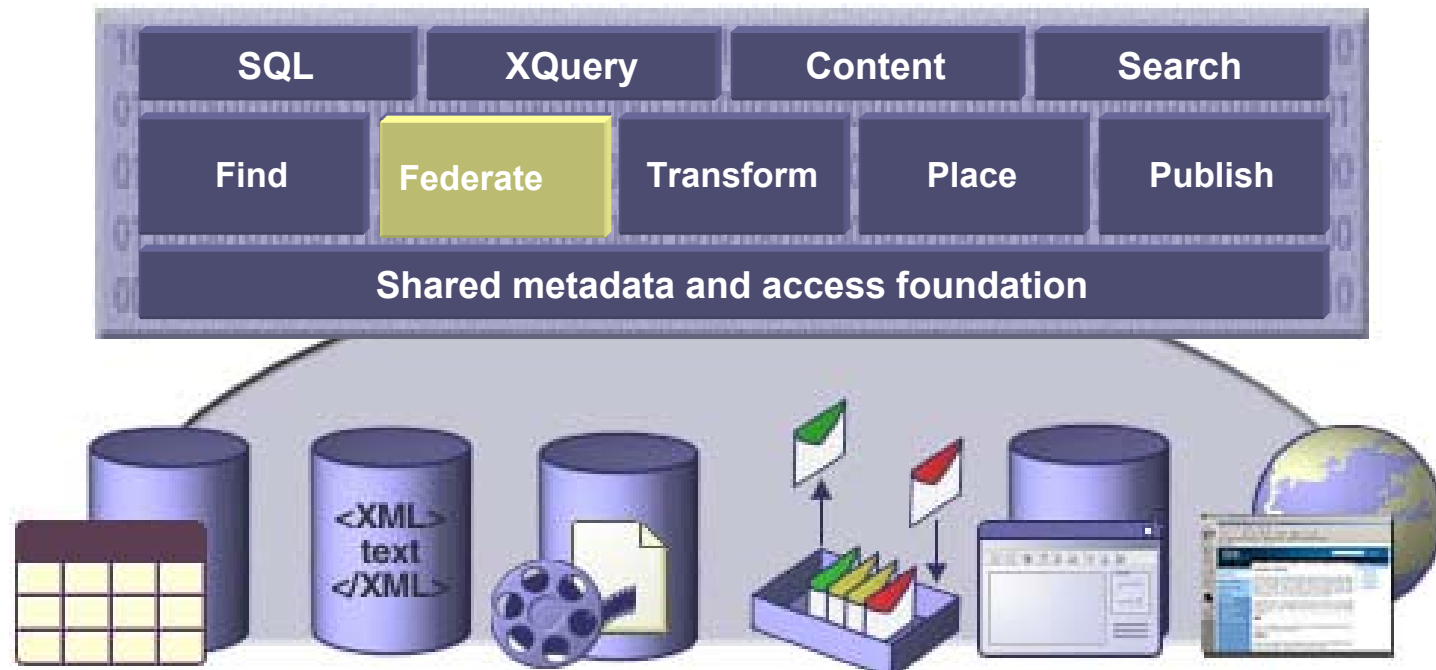
A well-disciplined architecture for natural-language content analysis and integration

- Lets highly skilled professionals focus on analytical algorithms
- Implemented as a Java UIMA framework
- Simplifies application development and deployment through plug and play composition
- Already deployed in WebFountain, WebSphere Portal, and other offerings
- Provides the basis for IBM's new Masala search capability



What's new in Masala

- Speeding time to market with federated access capabilities



Speeding time to market for more projects

Whether business intelligence or business integration projects are your focal point, DB2 Information Integrator helps you get to market faster.

- Extends access to more data sources
 - ▶ SAP, PeopleSoft, Siebel, and support for complex Web services results, such as those that might be returned from content repositories or legacy applications.
 - ▶ Java software development kit for developing new connections
 - ▶ New partners enable access to VSE environments (GT Software) and wide array of content and workflow systems (Venetica), plus much more
- Delivers more plug-ins for .NET, eclipse, and Rational tools
 - ▶ Inherits rich capabilities from DB2 UDB Stinger features
 - ▶ Validated for federated Masala access
 - ▶ Development repository for XML artifacts to facilitate collaboration
- Offers more partner tools and solutions from which to choose
 - ▶ Over 40 partners worldwide have tested with and gained skills in DB2 Information Integrator



Improving performance for key usage scenarios

Masala provides relevant and significant performance boosts to typical usage scenarios.

- Extend your enterprise warehouse with real-time data
 - ▶ Masala exploits parallel architectures typical in DB2 UDB warehouse implementations, thus analyses combining remote data with your warehouse can be distributed across multiple processors to speed results.

- Build a cross enterprise view of your financials or your customers
 - ▶ Masala more efficiently processes unions and outer joins typically required for building enterprise views across geographies or business units as well as for building unified views of customers or products.

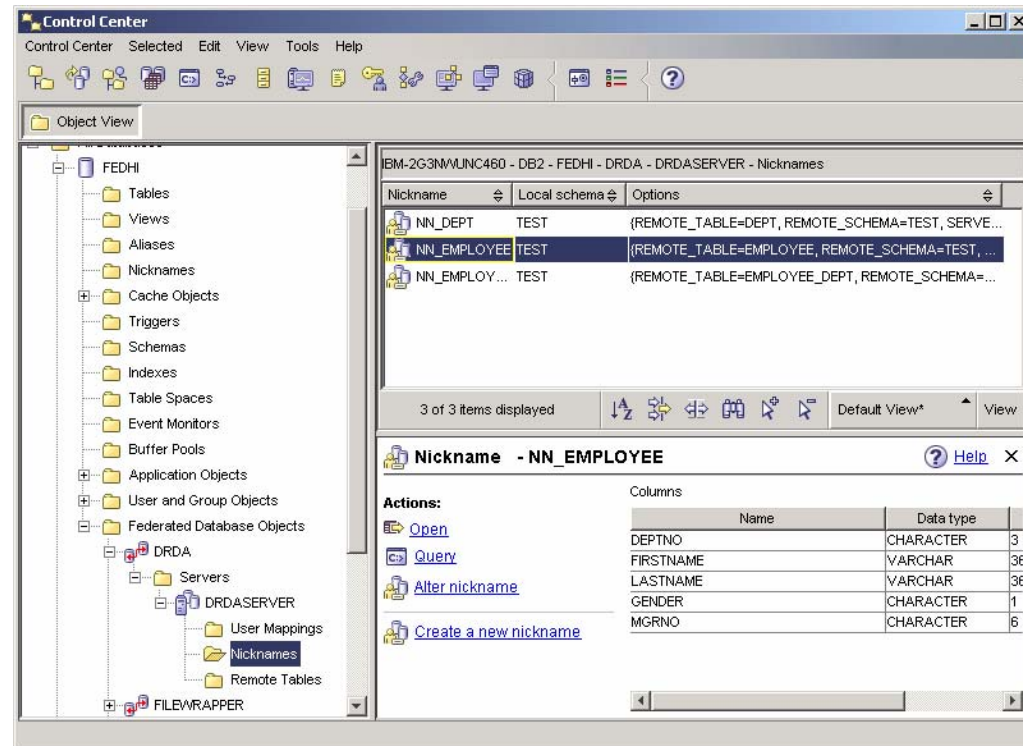
- Serve up information through corporate portals
 - ▶ Masala offers more options for defining and managing local caches



Easing set up and maintenance in dynamic environments

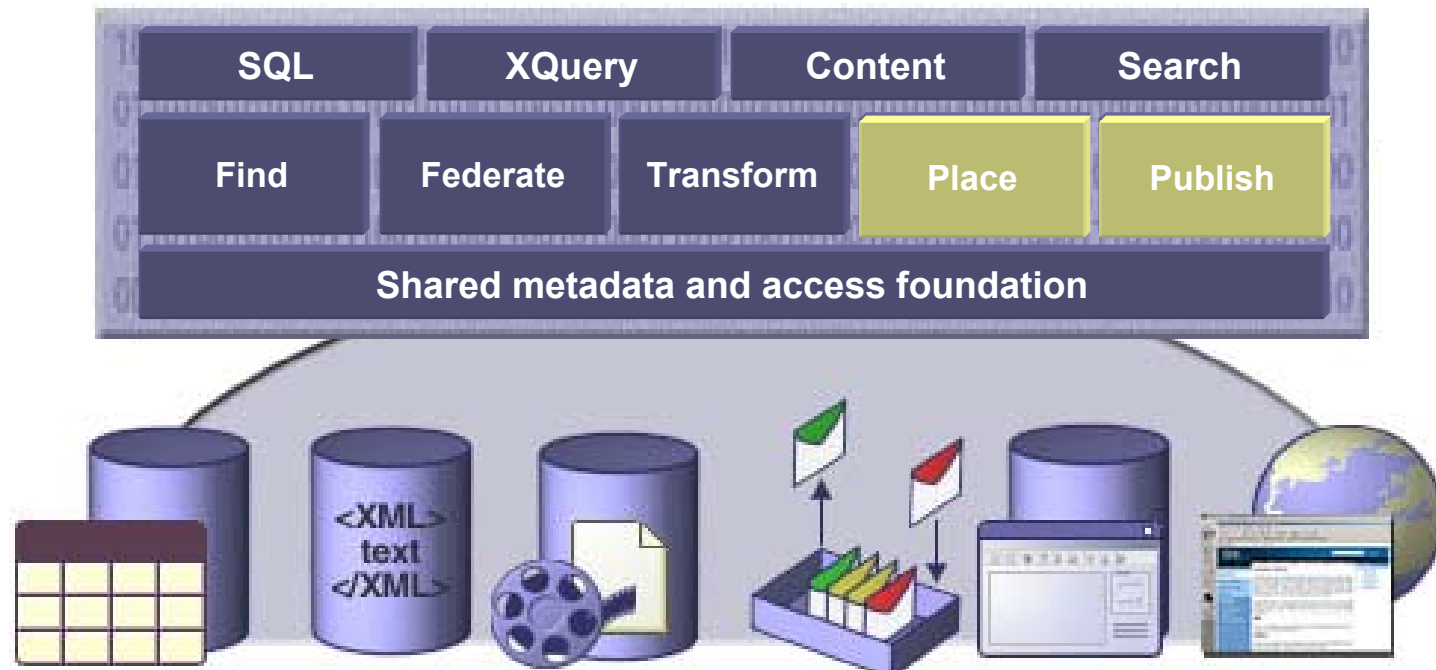
Masala introduces extensive tooling to easily configure, monitor, and maintain your integration hub in top running order

- Speed configuration with enhanced object discovery and new configuration wizard
- Set up cache and cache refresh to speed query responsiveness with new cache table wizard
- Move from test to production or replicate integration hubs with new deploy function
- Deal with environment changes by automating statistics refresh and detecting server or schema changes.
- Use DB2 Query Patroller to monitor, prioritize, analyze, and govern query workload to minimize disruption to critical systems.



What's new in Masala

- Adding value to the business integration environment



Publishing data events to facilitate business integration

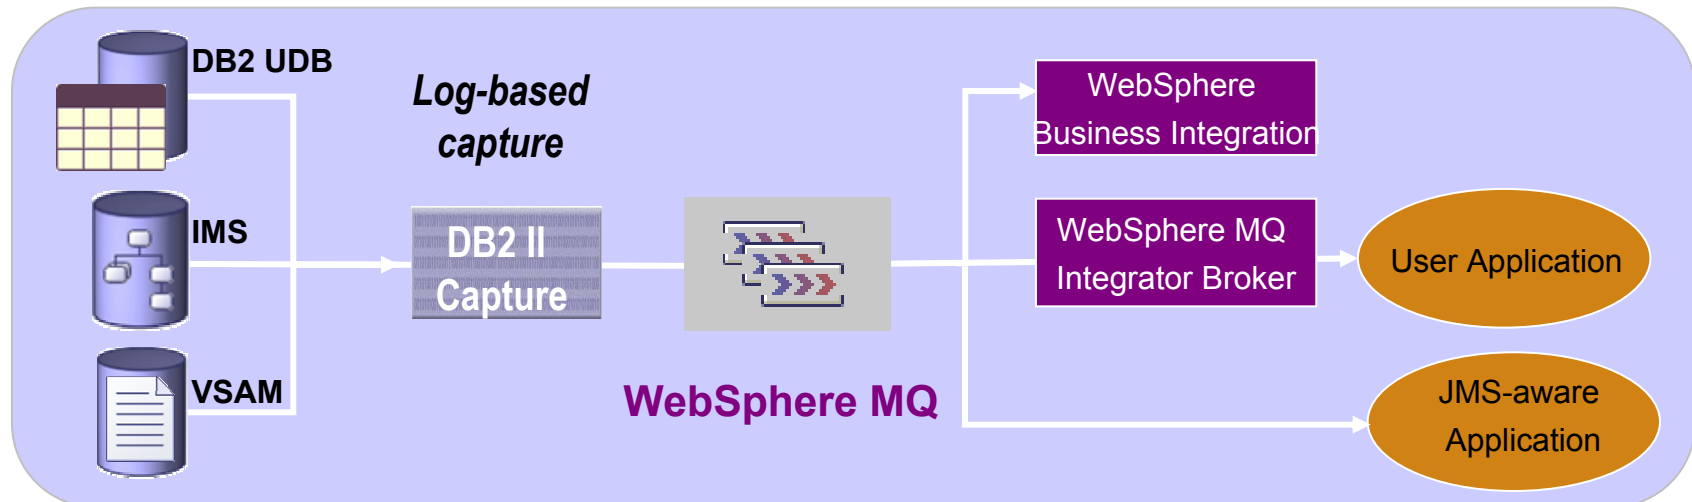
Masala captures database changes as XML messages and publishes them to WebSphere MQ

Function

- Publish events to a message queue
- XML self-describing format
- Wizard-driven configuration

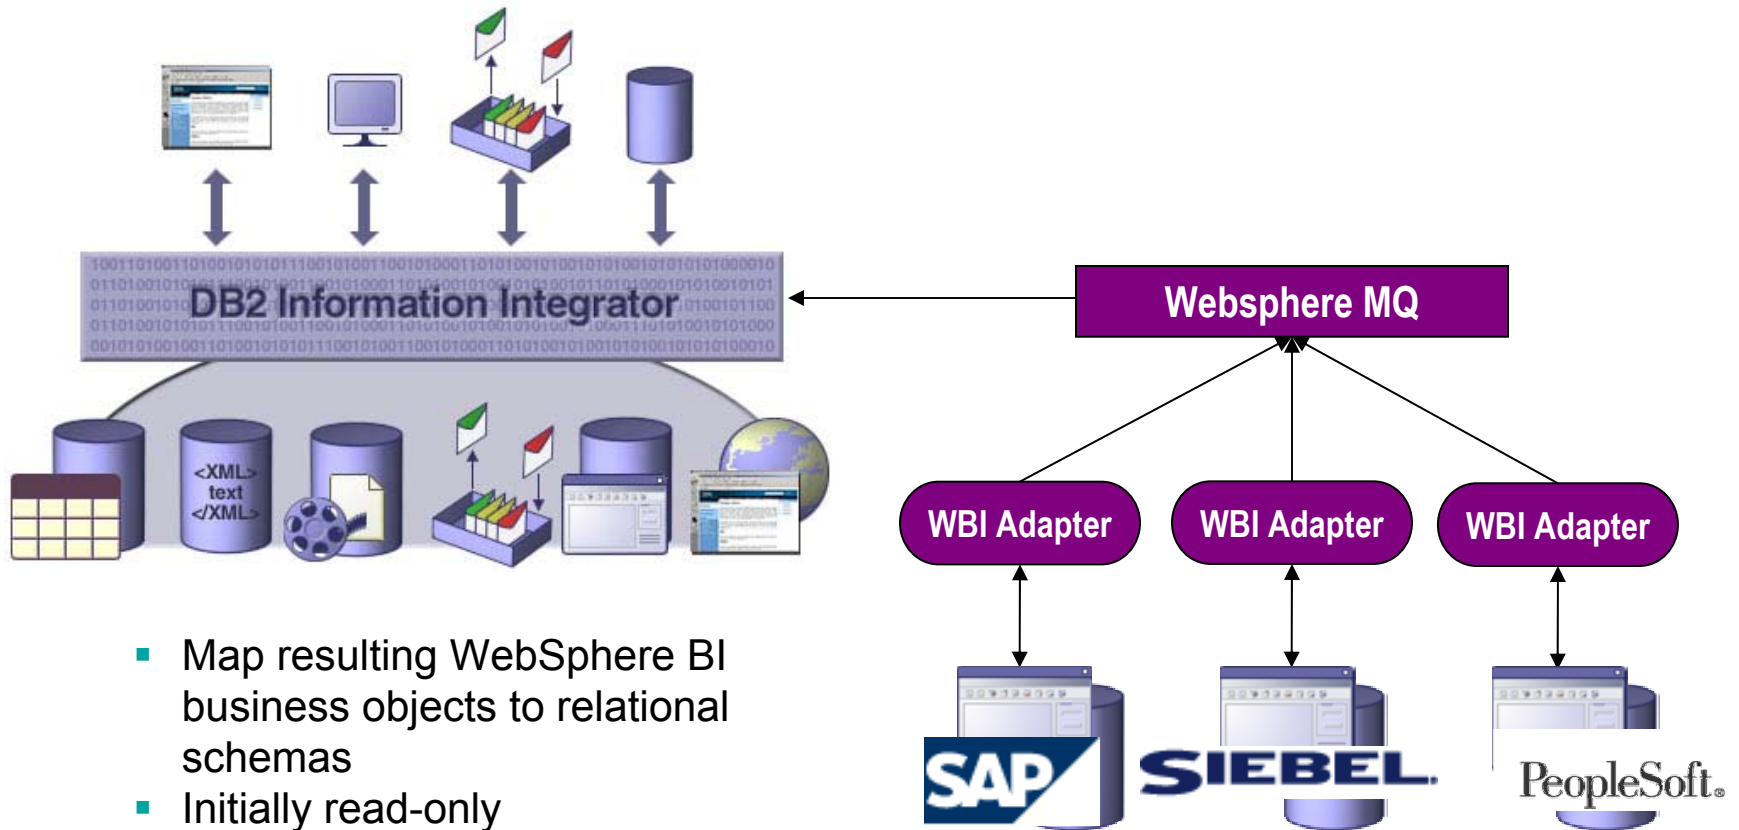
Usage

- Application to application messaging
- Event streaming
- Source for ETL tool



Using WebSphere Business Integration adapter portfolio

Masala takes advantage of WebSphere investments in packaged application access



- Map resulting WebSphere BI business objects to relational schemas
- Initially read-only



Builds on WebSphere MQ speed and resiliency

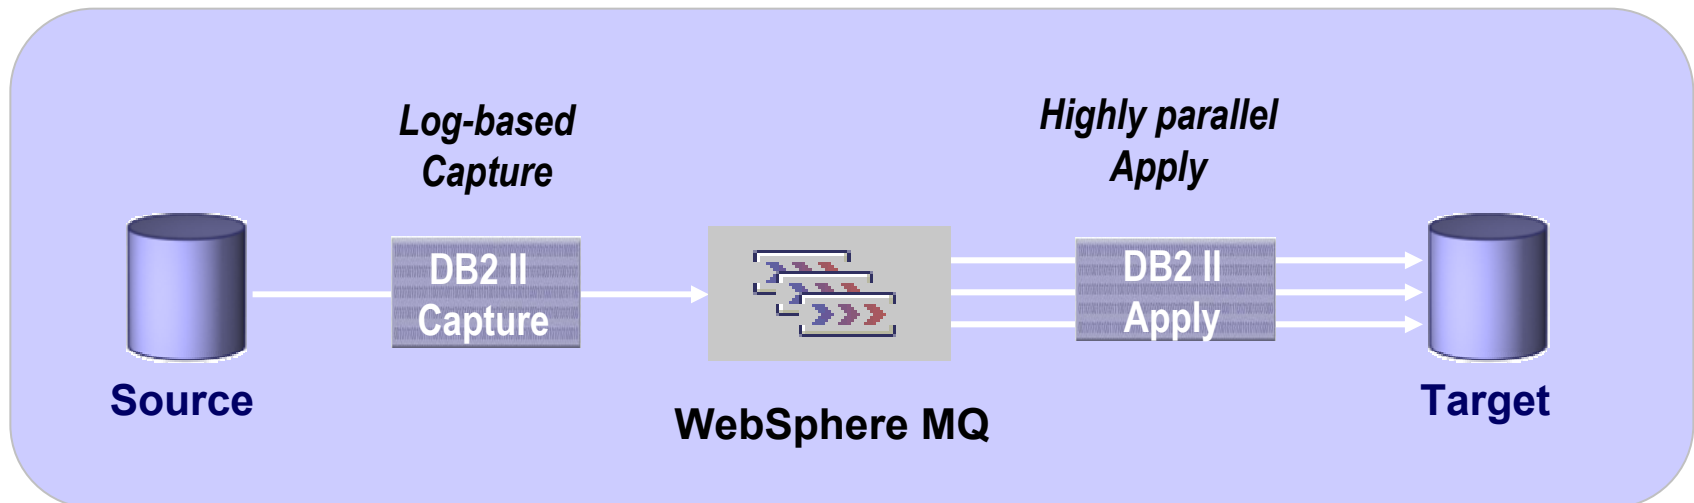
Masala introduces new replication architecture for delivering extremely low latency replication

Function

- Each message represents a transaction
- Apply simulates transaction environment
- Conflict detection and resolution
- Wizard-driven configuration

Usage

- High availability
- Workload distribution
- Application integration



Masala release will be for DB2 UDB only

Agenda



DB2 Masala Technology Preview

- DB2 Information Integrator
 - ▶ Vision
 - ▶ Function
 - ▶ Value

- New capabilities coming in Masala
 - ▶ Finding the most relevant corporate information
 - ▶ Speeding time to market with federated access capabilities
 - ▶ Adding value to the business integration environment

- Delivering tangible value from partnerships
 - ▶ Additive value and support for client's investments
 - ▶ Product integration and testing
 - ▶ Skilled solution deployment

World-Class Partner Ecosystem

Over 40 partners worldwide


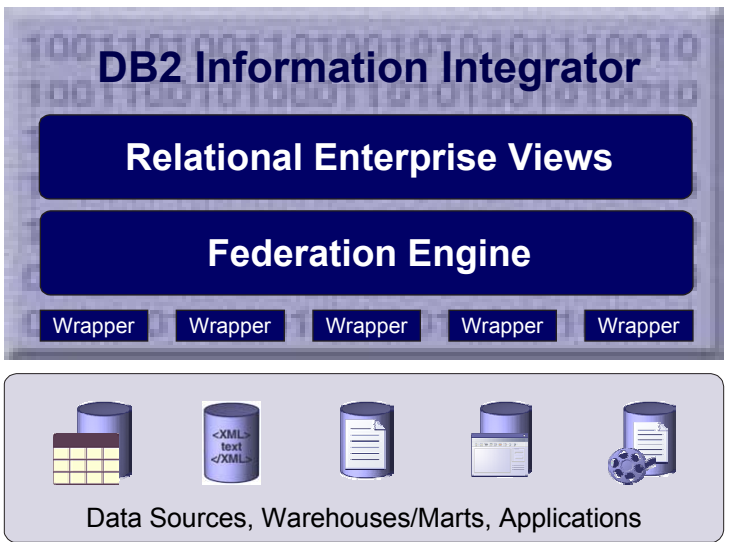
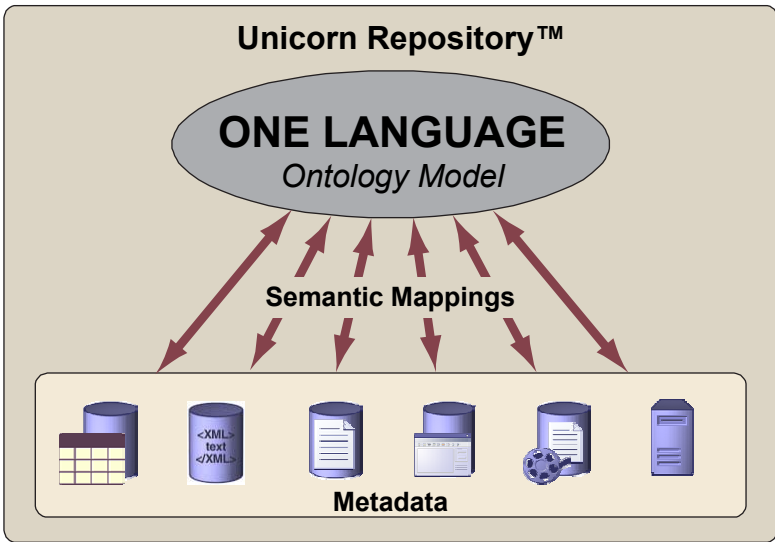

- Enhance analysis and reporting
 - ▶ Certive, Hinge Software, Viador...
- Enable applications and application development
 - ▶ Agilent Technologies, deCODE genetics, Scitegic, Spotfire, UFSOFT,...
- Enrich DB2 II function
 - ▶ Unicorn, Sypherlink,...
- Extend reach
 - ▶ Venetica, GT Software,...
- Expedite deployment worldwide
 - ▶ BeaconIT (Japan), Bearing Point , CDE Solutions (UK), TCS (India MSI),...



Not a complete list of partner or logos



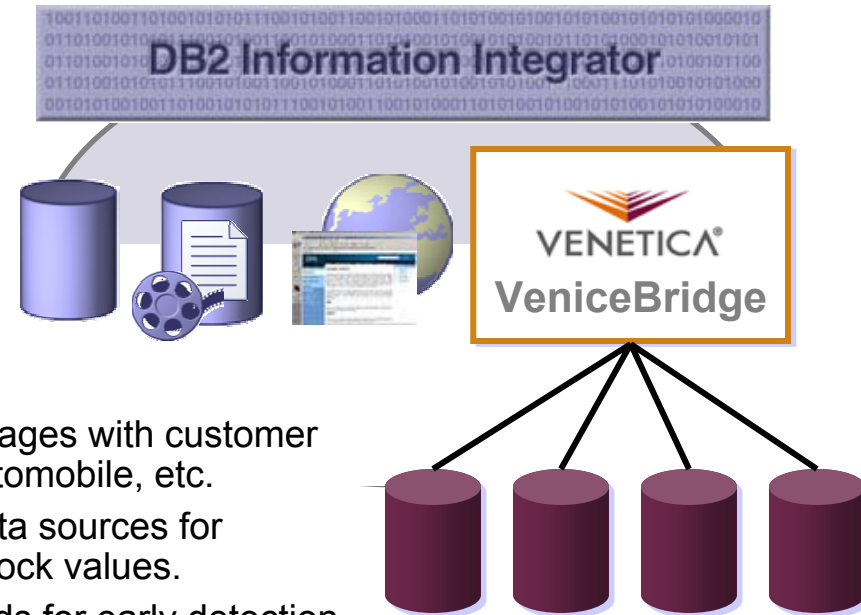
Unicorn and DB2 Information Integrator



Venetica

Content and data integration requirements emerging

- Regulatory compliance – combining loan applications, credit rating, loan default records, etc. to assess risk and justify actions
- Law enforcement – combining diverse information (criminal records, immigration records, travel records, DNA records...) to assist criminal investigations
- Pharmaceutical research – research collaboration across biological and chemical research domains
- Medical research – analyzing imaging data together with treatment plans, prognosis, and actual results
- Claims adjustment – combining current claim forms and images with customer account detail and historical trends re: the claimant, the automobile, etc.
- Finance – analyzing news feeds and other unstructured data sources for specific kinds of information about companies that affect stock values.
- Call Center– analyzing call center and other problem records for early detection of problem trends



Documentum, FileNet Content Services, FileNet Image Services, FileNet Report Manager, FileNet, eProcess, FileNet P8 BPM, FileNet P8 Content Mgr, Google, IBM DB2 Content Manager, IBM DB2 CM OnDemand, IBM ImagePlus, IBM WebSphere MQ Workflow, Interwoven TeamSite, Lotus Domino.Doc, Lotus Notes, Microsoft Index Server, Microsoft Site Server, Open Text Livelink, Stellent, Verity K2, UNIX and Windows File Systems

Systems Integrators

Partners bring DB2 Information Integrator to clients in a solution context



Enition In Joint Software Solution Sale With IBM

April 6, 2004: 3:22 a.m. EST



From the editors of CNN and MONEY magazine

Enition said Tuesday that CDE Solutions, its principal trading subsidiary, has concluded a joint software solution sale with IBM to Shimadzu Corporation, a manufacturer of Life Sciences analytical instrumentation....Commenting, Ray Dutton, Chief Executive said: "Several other companies are at an advanced stage of trials with our innovative technology and I hope to be able to report further successes by CDE Solutions in the near future."

- Allied Technologies
 - ▶ Taiwan
- Azure Technologies
 - ▶ Singapore
- BAE Systems
- Beacon IT
 - ▶ Japan
- Bearing Point
- Booz Allen Hamilton
- CDE Solutions
 - ▶ UK
- CGE&Y
- Computer Sciences Corp
- Deloitte Consulting
- Fillmore Group
- IndigeTech
- iQDynamics
 - ▶ Singapore
- Knowledge Dynamics
 - ▶ Singapore
- Lenovo
 - ▶ China
- Lightyear consulting
- Resolvo Systems
 - ▶ Singapore
- Sysware
 - ▶ Taiwan
- TCS
 - ▶ India

DB2 Information Integrator Services Opportunities

- Consulting and solution design
 - ▶ Establishing an information on demand infrastructure
 - E.g., Leveraging real time federation and low latency replication
 - ▶ Integrating or re-purposing heterogeneous legacy application data
 - E.g., IDMS, Adabas, IMS, VSAM, ...
 - ▶ Integration of structured and unstructured data for industry applications
 - E.g., Life Sciences, Government, Media, ...
 - ▶ Establishing enterprise wide search and common data access layer
- Education
 - ▶ Federated DBMS Design and Implementation with DB2 Information Integrator
 - ▶ DB2 Information Integrator data placement technologies
 - ▶ Performance tuning and considerations
- Integration and implementation Services




Gold Consultant Support

- Get evaluation copies of DB2 Information Integrator
 - ▶ Contact Gil Lee, gklee@us.ibm.com, (408)463-3860
- Participate in betas
 - ▶ Contact Doreen Fogle, dfogle@us.ibm.com, (877)508-3457
- Technical sales assistance in customer engagements
 - ▶ Contact Ravi Ravindranath, ravindr@us.ibm.com (214)233-3410
- Solution assurance assistance
 - ▶ Contact Monique Mol-Kobeh, mmonique@us.ibm.com (408) 463-4431
- Discuss additional support requirements
 - ▶ Contact Atul Chadha, achadha@us.ibm.com , (408)463-3706



Summary

Masala extends IBM's leadership in Enterprise Information Integration

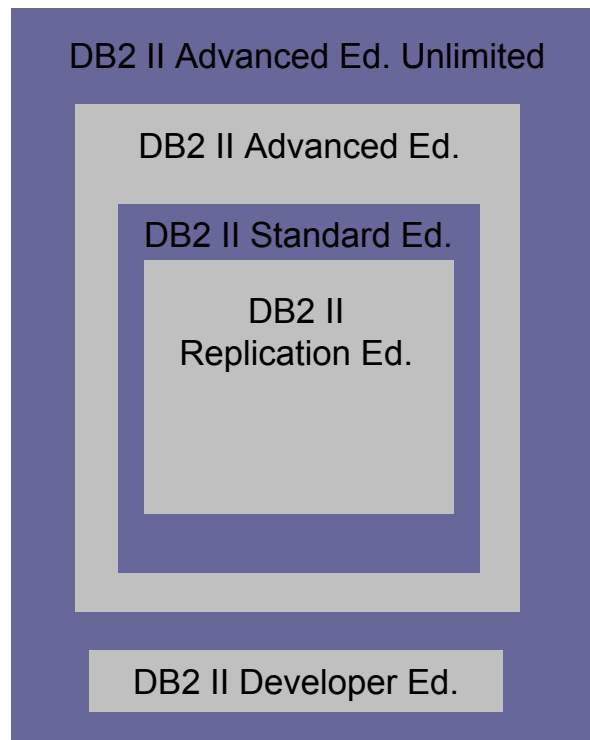


DB2 Masala Technology Preview

- Finds the most relevant corporate information
 - ▶ Search is a fundamental aspect of information integration technology
 - ▶ Masala delivers high quality, scalable, secure enterprise search
 - ▶ UIMA help businesses get the most out of their content assets
- Strengthens DB2 Information Integrator core value
 - ▶ Speeding time to market for more applications
 - ▶ Boosting performance for key usage scenarios
 - ▶ Making it easy to setup and maintain in dynamic environments
- Adds value to business integration environment
 - ▶ Publishes data events into business integration environment
 - ▶ Leverages investments in WebSphere BI adaptor portfolio
 - ▶ Delivers low latency replication for centered around DB2 UDB
- Delivers tangible value from partnerships
 - ▶ Deliver the industries broadest EII capabilities
 - ▶ Provide synergistic value through integration and testing
 - ▶ Evangelize the technology and drive more usage

DB2 Information Integrator (DB2 II) in 2003

Linux, Unix, Windows



z/OS

