WebSphere Information Integrator Success Stories

James Reed, Product Marketing Mgr
Information Integration Solutions

The Information Challenge

- **Where is my information?**
- **How do I get it when I need it?**
- **What does it mean?**
- **Can I trust it?**
- **How do I get it in the form I need?**
- **How do I get it where it needs to go?**
- **How do I control it?**

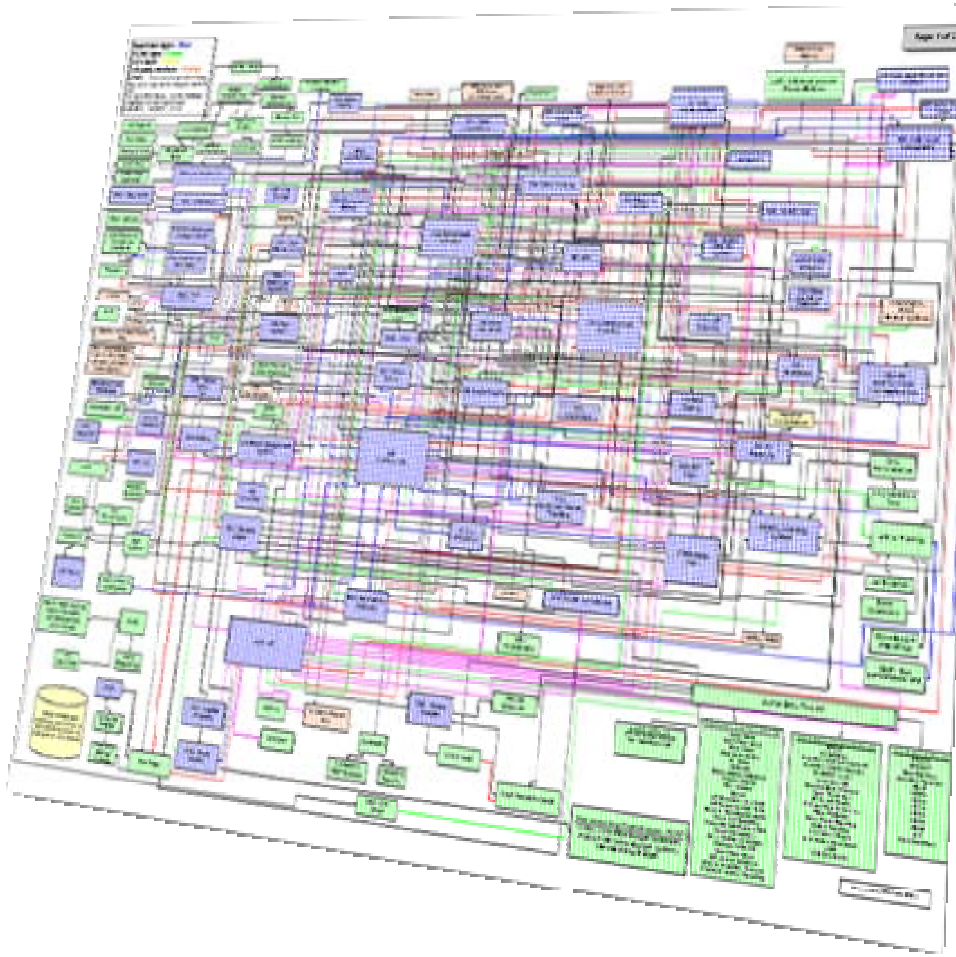

Market Dynamics

- Convergence continues
- Organizational structure and buying patterns changing
- New requirements bringing structured and unstructured worlds together
- An overall integration strategy is now a necessity

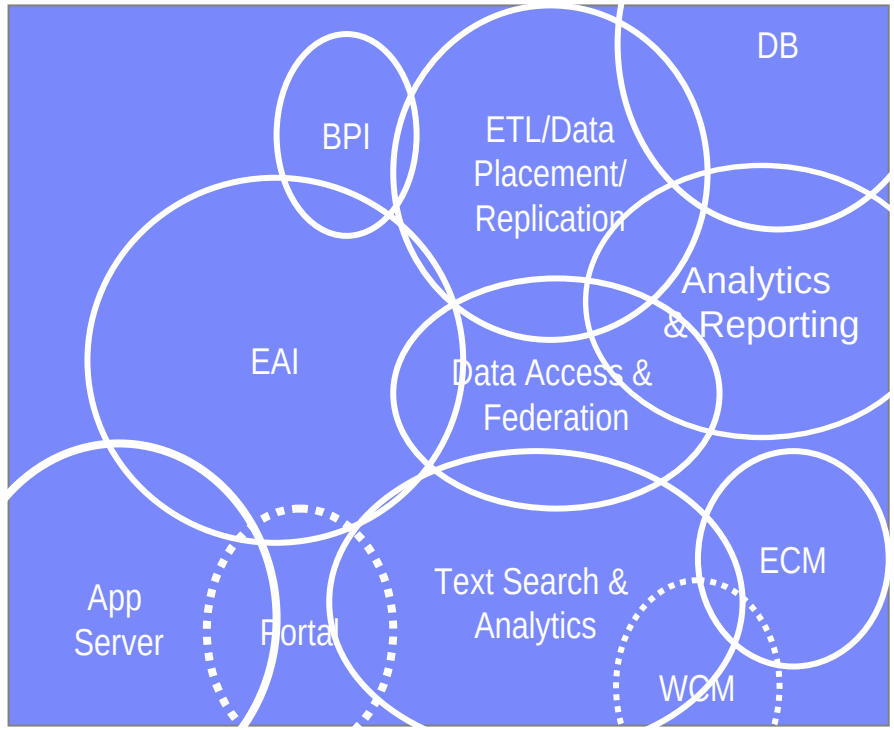

Business Integration Market

Meta Group Opinion:

"ETL, EAI, and EII Converge in the Data Exchange Facility"

Mark Beyer, Janelle Hill, 14 February 2005

"META Trend: Data management will increasingly be viewed as a part of overall integration architectures and designs (2004/05). This will cause organizations to consolidate these efforts organizationally (i.e., COE for integration services), to take advantage of common technologies and skills. Distinct data integration technologies (e.g., EII, ETL, EAI) will converge (2007), ultimately surviving only as various subsets of intermediary capabilities in the service-oriented architecture (2009)."

The IBM WebSphere Information Integration Platform

Delivering information you can trust

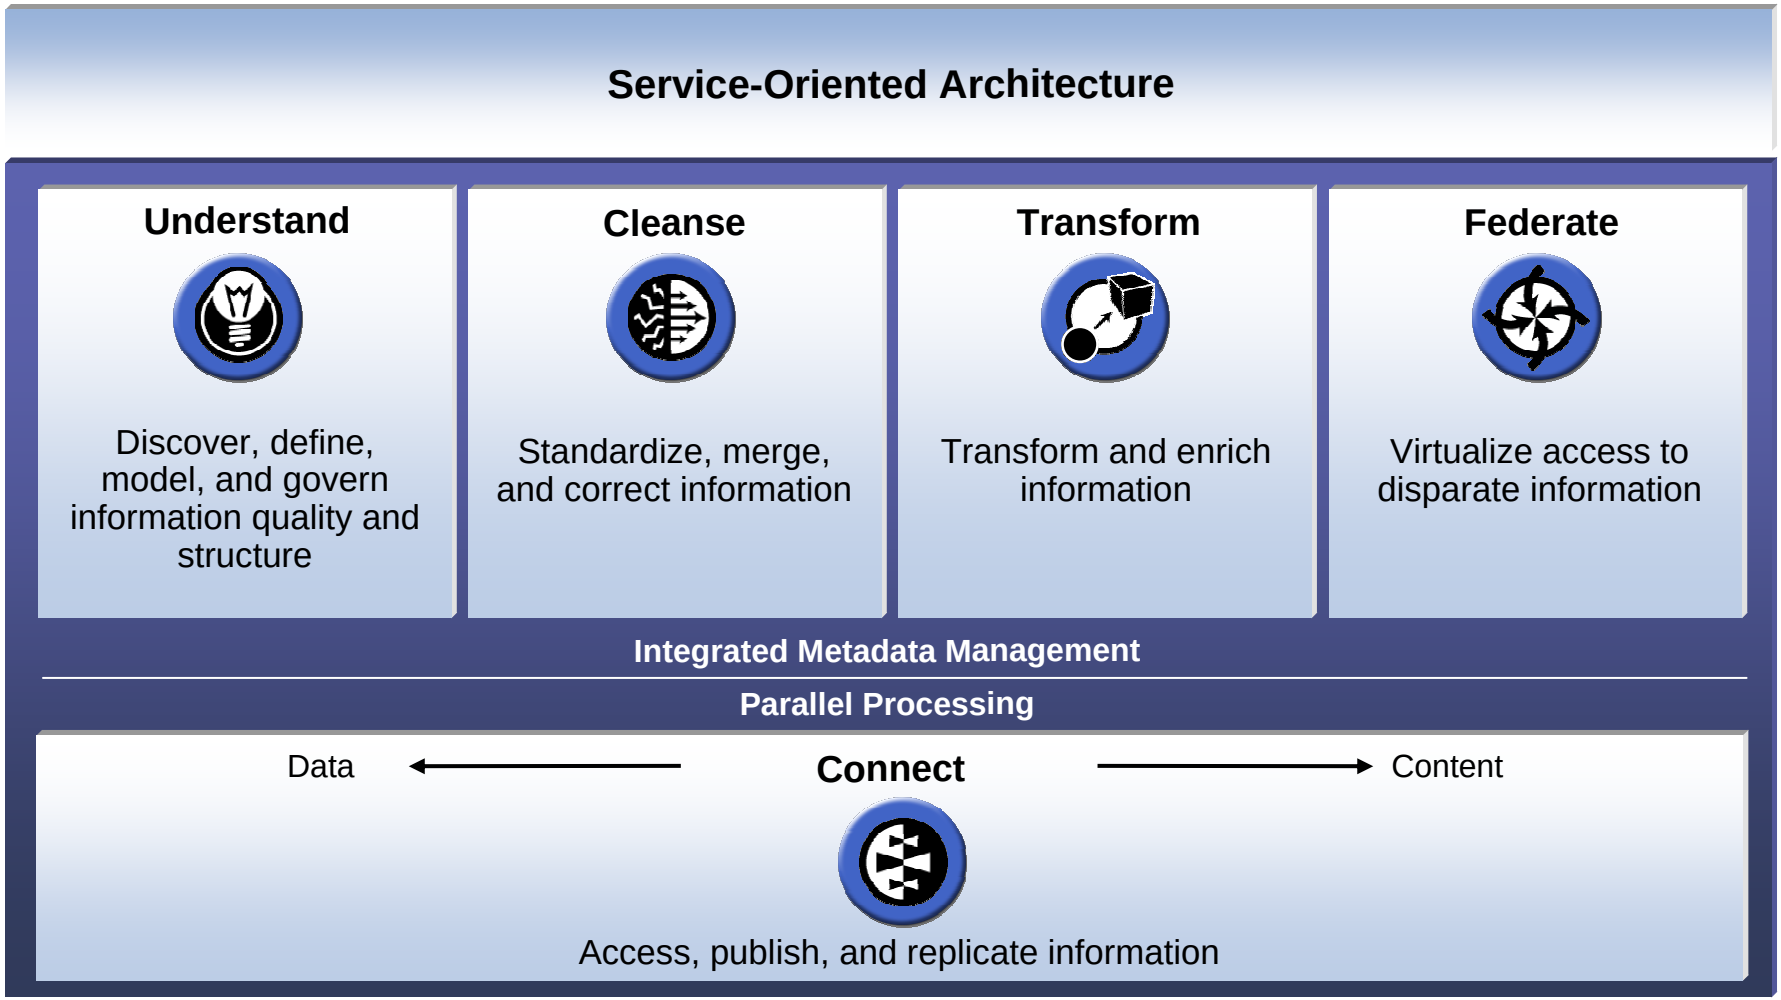

The IBM WebSphere Information Integration Portfolio

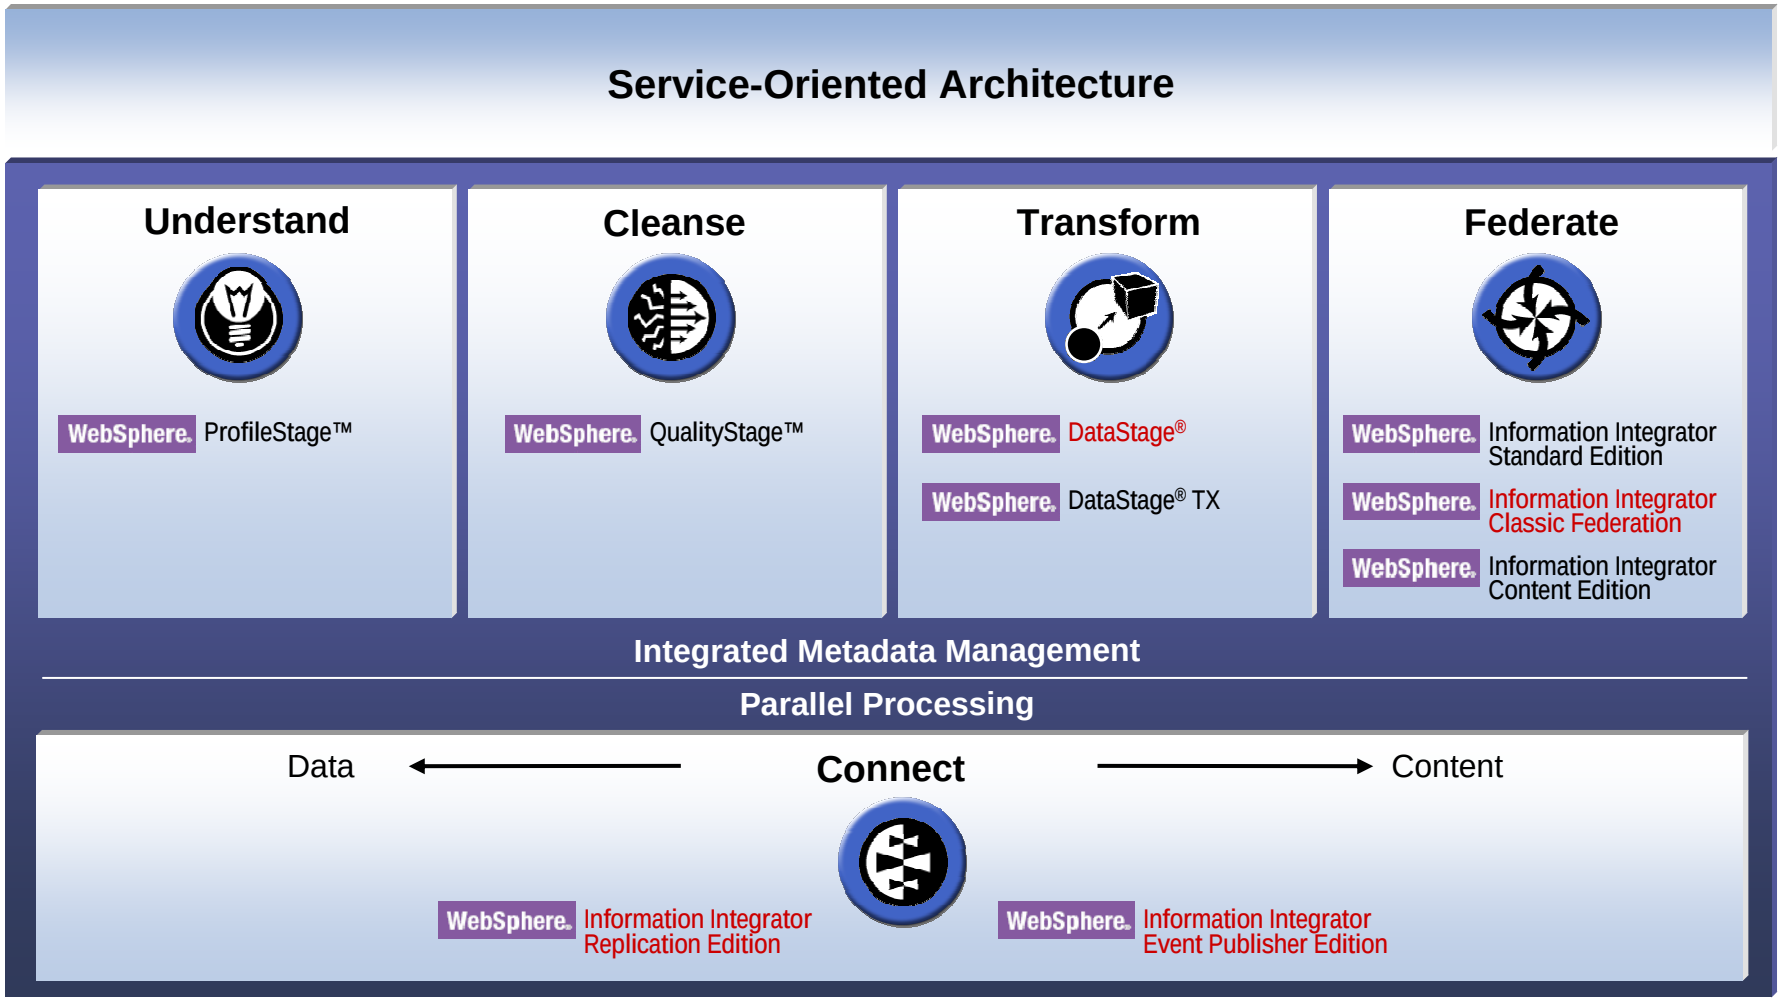

Information Integration: Federate

Architects

IT Admin

Developers

Testers

Federate

WebSphere. Information Integrator Classic Federation

SQL-based federation across mainframe databases and files

WebSphere. Information Integrator Standard Edition

SQL-based federation across databases, files, services, and packaged applications

WebSphere. Information Integrator Content Edition

Content-optimized federation across repositories, files, and collaboration and workflow systems

Information Integration: Connect

Developers

Testers

IT Admin

Connect

WebSphere. Information Integration Portfolio

Direct, native access to relevant sources bundled with each product in the portfolio

WebSphere. Information Integrator Replication Edition

Consolidation, synchronization, and distribution across disparate databases

WebSphere. Information Integrator Event Publisher Edition

Changed data and event-based publishing of data

For example:

Leveraging zSeries Assets with WebSphere II

Capability	Value
Provide bidirectional access to mainframe data using standard ODBC and JDBC, no mainframe programming	Speed project deployment, reduce development and maintenance costs, and reduce project risk
Implement low-latency replication for high volume workloads with active conflict detection and resolution	Maintain high availability for critical data or distribute workloads across distributed sites.
Replicate data to and from DB2 UDB for zSeries and distributed open systems databases	Leverage z-Series investment for new business applications
Publish mainframe database events to drive processes or link applications	Minimize application impact and maintenance

Agenda

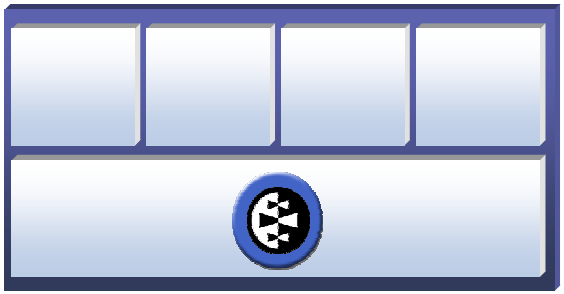

- **IBM Information Integration Portfolio**
- **WebSphere II Replication**
- **WebSphere II Mainframe access & federation**
- **Data Event Publishing**
- **Wrap-up & Additional Questions and Answers**

Slide 2 of 128 more to go

SQL Replication --
Flexible transformation and scheduling

- **Function**
  - Filter and transform
  - Apply by table or by transaction
  - Choose latency by schedule, interval, event or continuous
  - Replicate point-to-point for distribution or for consolidation
  - Maintain snapshots, simple copies, histories or aggregates

Usage
- Business intelligence
- Distribution
- Consolidation
- Application integration

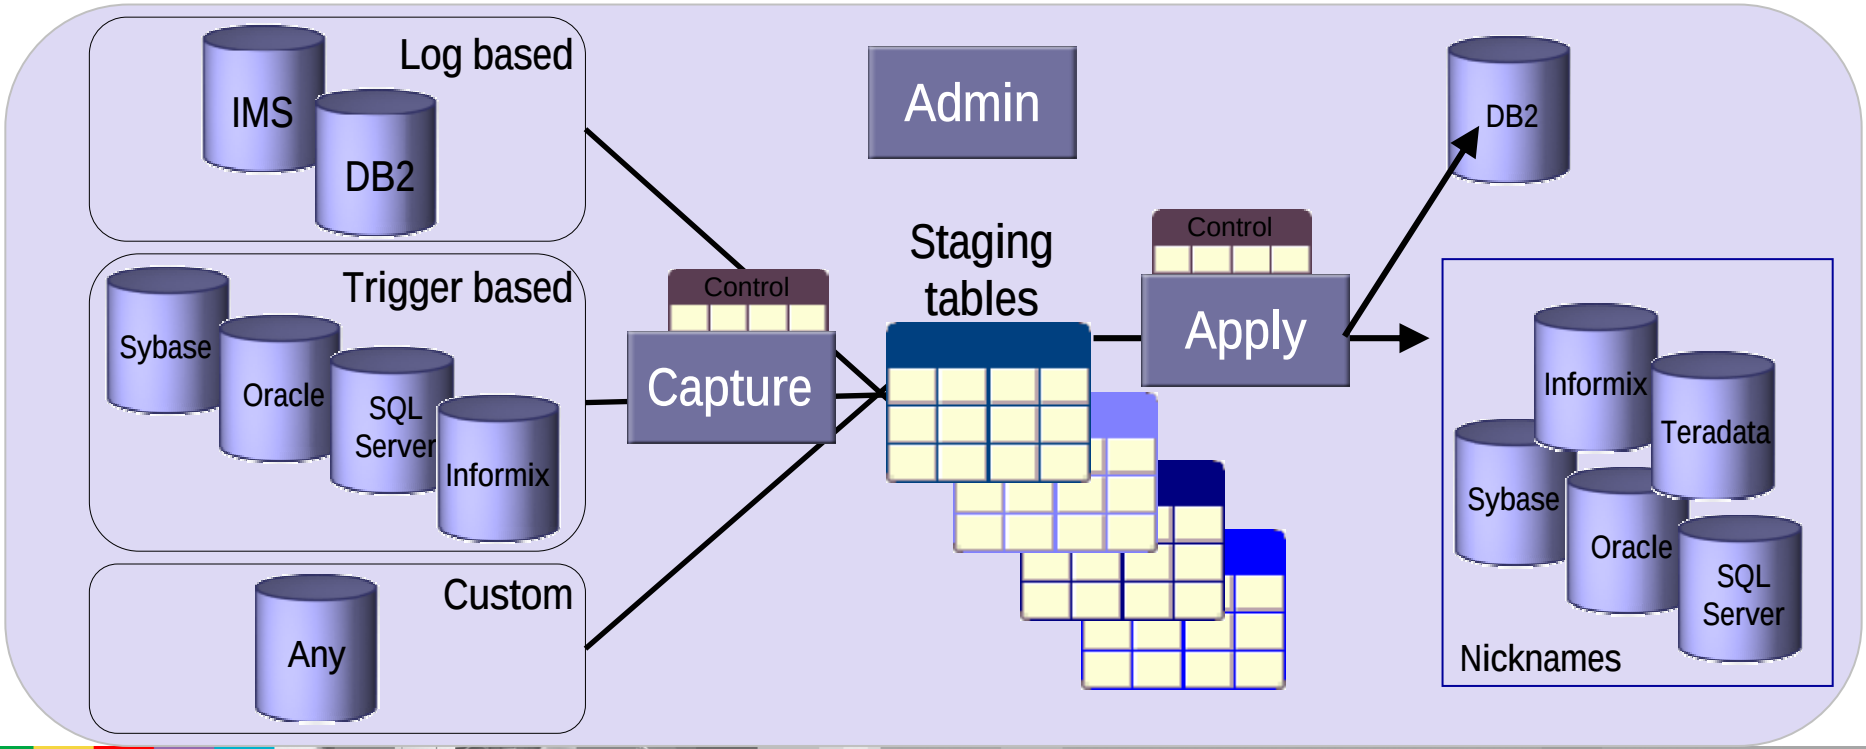

Q Replication -- low latency for peer-to-peer environments

Function

- Replicate rows or transactions
- Filter and transform data
- Detect and resolve conflict
- Configure and monitor environment

Usage

- High availability
- Workload distribution
- Application integration

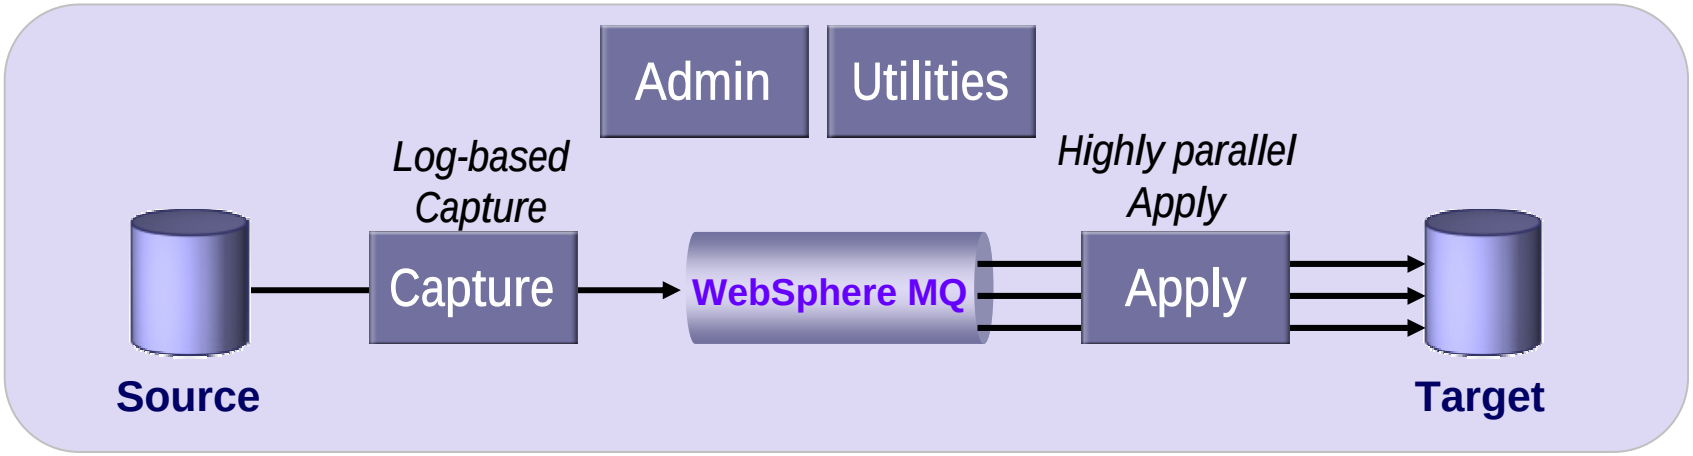

- DB2 UDB source on z/OS, Linux, UNIX, and Windows
- DB2 UDB native apply target on z/OS, Linux, UNIX, and Windows
- Oracle & Sybase apply targets now avail; SQL Server avail Q3

International provider of financial & investment services

Challenge

- Replication of critical order processing details for core business functionality
- Corporate initiative to provide customers perform with real-time queries across multiple sites.

Solution

- Q Replication for high speed movement of up to 10 Million transactions to secondary site several thousand miles away. Current implementation is Uni-Directional with peer-to-peer plans.

Business benefits

- Replicating 10 Million transactions with less than 2 seconds latency.
- More efficient and cost-effective resource utilization
- Secondary platform services reporting and business intelligence queries and acts as backup to primary

Technology benefits

- Real-time back up of secondary system provides results in increased capacity for peek workloads.

Online Trading – A case for very high speed replication

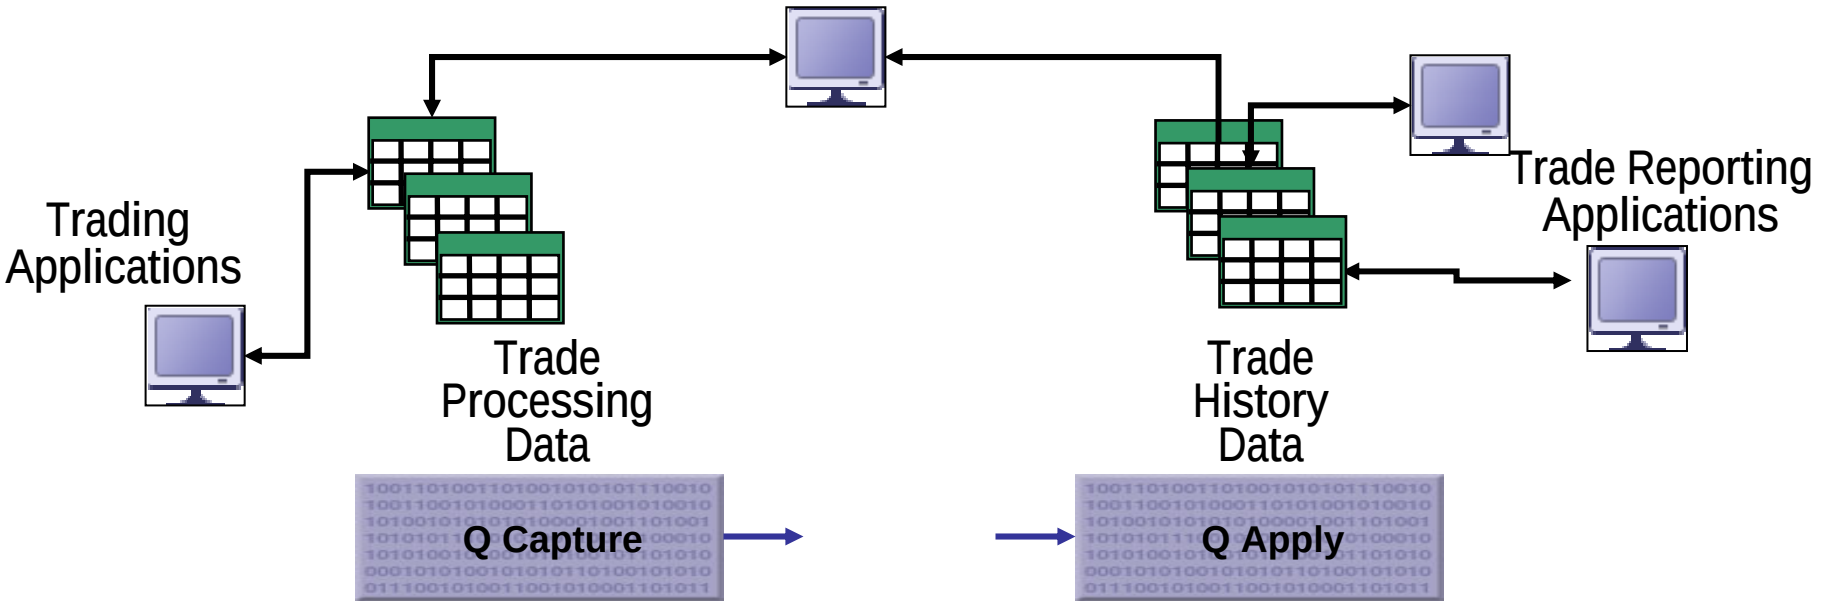

Primary System

- Supports trade (write) processing
  - Speed is critical
- Optimized for transactions
- Capacity for peek update workloads

Secondary System

- Back-up operational platform
  - 1.5 second average lag
- Services business intelligence queries
- Reduces demand on primary system

More efficient and cost-effective resource utilization
Secondary platform services "real work" and acts as backup to primary

CitiStreet

Challenge

- **Support single sign-on access through both Web and IVR applications ensuring 24x7 portal access for plan participants and sponsors**

Solution

- **Support redundant, active single sign-on applications for failover processing replicating profile changes between them in real time.**

"Since nearly 10 million of CitiStreet customers are offered 24-hour access to their retirement accounts, the company can't afford downtime and must be able to replicate data changes when they happen. We fully replicate our database over redundancy data lines, so to us the stability and speed of that asynchronous replication is strategic for us."

Barry Strasnick , CIO
CitiStreet

Overview

- CitiStreet is one of the largest and most experienced global benefits providers servicing over 9 million plan participants across all markets. CitiStreet was formed in partnership between subsidiaries of State Street Corporation and Citigroup

Business benefits

- Ensure application availability for plan participants and sponsors
- The new solutions from IBM will improve data integrity with a reduced level of maintenance

Technology benefits

- Maintain bi-directional synchronization of profile updates (approx 175,000 updates daily) in real time

Mazda North America

- **Vision**
  - Mazda North American Operations (MNAO) is responsible for the sales and marketing, customer service and parts support of Mazda vehicles in the United States. Headquartered in Irvine, Calif., MNAO has more than 700 dealerships nationwide.
- **Challenge**
  - Dealers were having trouble matching customer demand with available inventory due to latency of data in dealer portal
  - Full refresh strategy limited portal data currency to every 15 minutes due to network bandwidth
  - Dealers wanted more current data to track achievement of period-end sales goals
- **Solution**
  - Using WebSphere Information Integrator, Mazda replicates sales and inventory information from DB2 UDB on MVS to Microsoft SQL server (the portal data server) every minute.
- **Business Value**
  - Increase auto sales, increase dealer satisfaction, use existing application infrastructure
- **Competitive Value**

Within 5 weeks of receiving the WebSphere Information Integrator product we were able to implement it in our Test, QA and Production environments. What we accomplished in this short amount of time without any prior knowledge of the product is a testimony to how easy the product can installed, configured and utilized in a production environment. It now provides us with up to the minute sales activity. This is significant improvement over our previous process which we were doing every 15 minutes.

– Joe Neria, Software Consultant, Mazda

Mazda North America

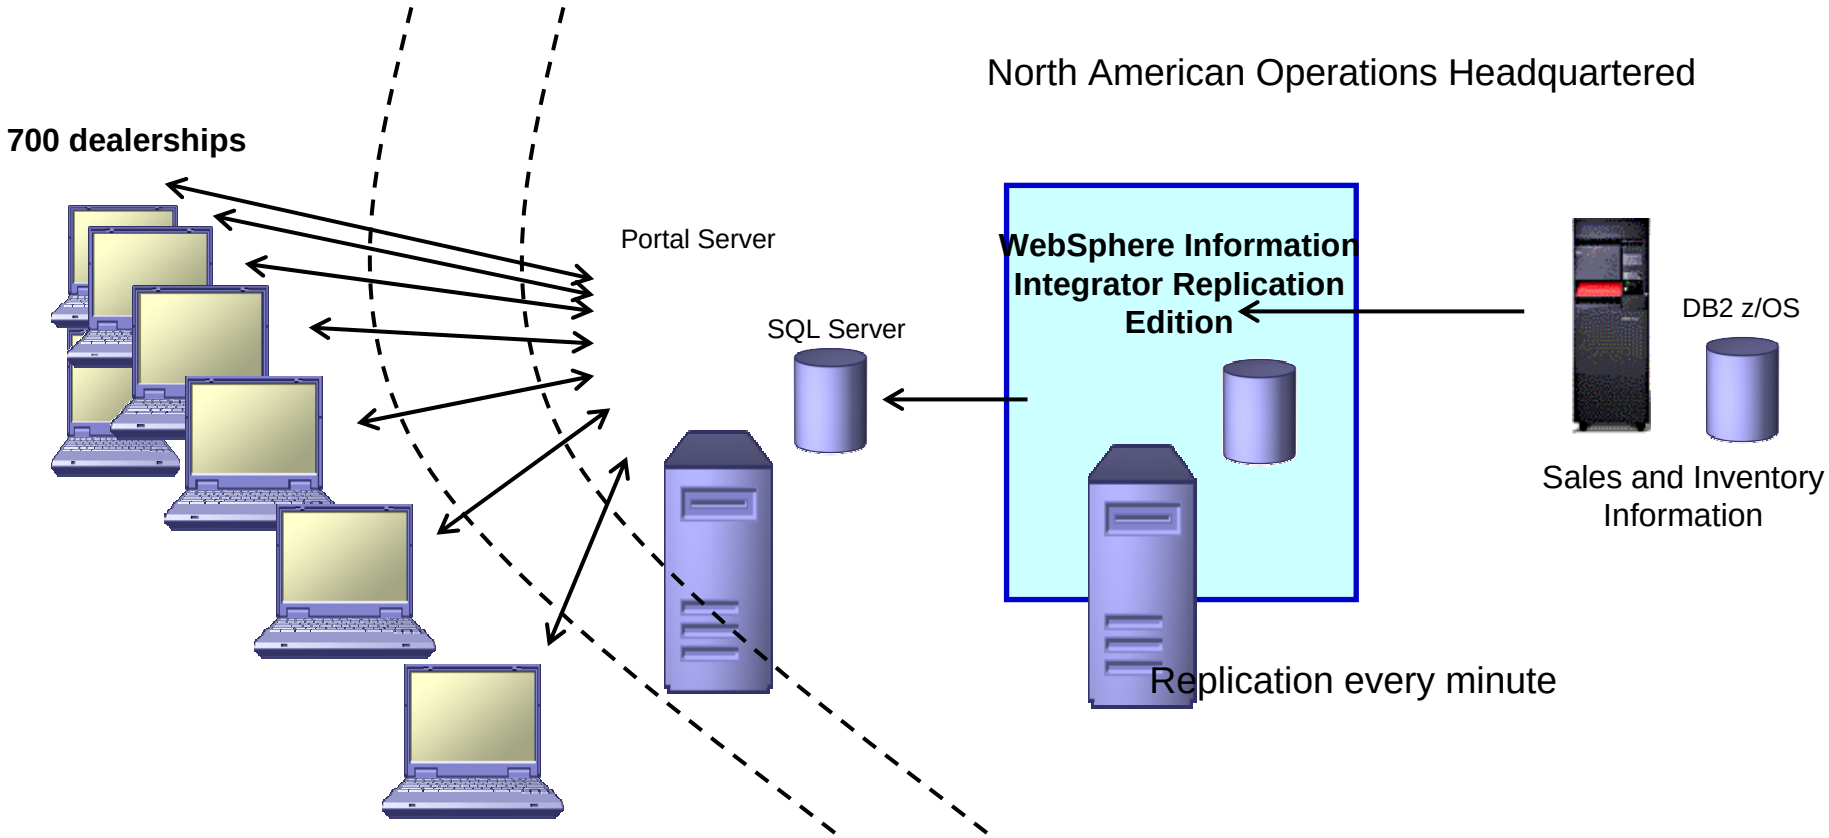

Fondiaria Insurance - Italy

FONDIARIA - SAI

Vision

- Fondiaria – SAI has grown to be a major player in the Italian insurance community after the consolidation of SAI, Fondiaria and Milano Assicurazioni.

Challenge

- After the implementation of a new Oracle Financials accounting application, the company struggled to keep information consistent with their DB2 OS/390 database.

Solution

- Rather than replace their existing DB2 solution which was an expensive proposition, they looked for a way to make the Oracle data accessible from OS/390 environment.
- Fondiaria is now using WebSphere Information Integrator to read from Oracle and write to the DB2 OS/390 environment.

Value

- Timely and cost effective solution to keep critical information sources accurate

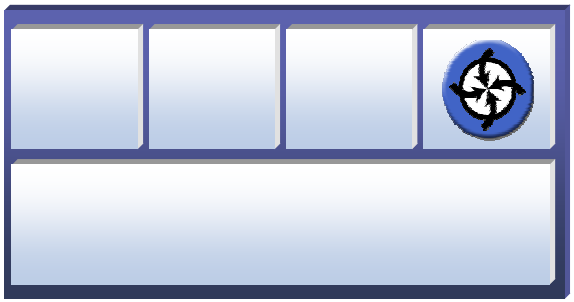

- IBM Information Integration Portfolio
- WebSphere II Replication
- **WebSphere II Mainframe access & federation**
- Data Event Publishing
- Wrap-up & Additional Questions and Answers

SQL Federation for the Mainframe

Integrate Mainframe Data Assets

- SQL-based read and write access to mainframe data sources
  - Standard ODBC and JDBC
- Multi-threaded with native drivers for scalable performance
- Metadata-driven for easy configuration and maintenance
  - No mainframe programming required
  - Fast installation & configuration
  - Easy maintenance
- Works with existing
  - Mainframe infrastructure
  - Application infrastructure
  - Tools infrastructure

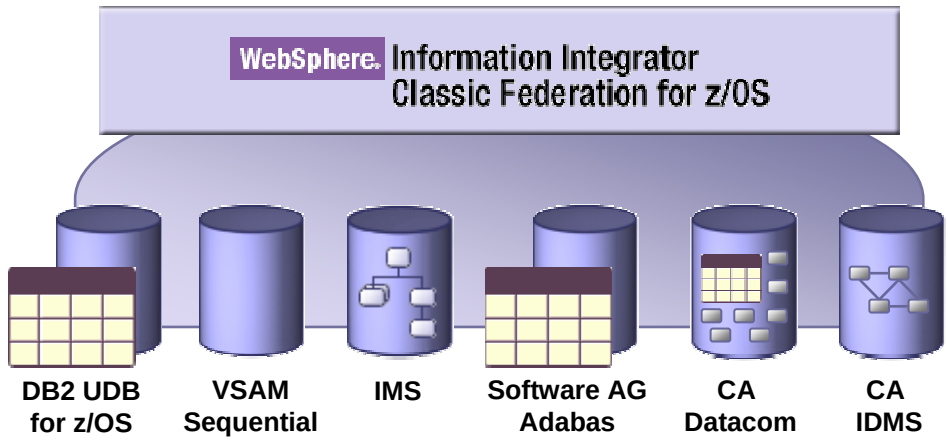

WebSphere II Classic Federation Implementation

- Create relational description of mainframe data sources by mapping the physical data definitions to logical tables and views
- Mainframe Server and components act as a relational database engine
- JDBC and/or ODBC drivers provide standardized interface for tools and applications

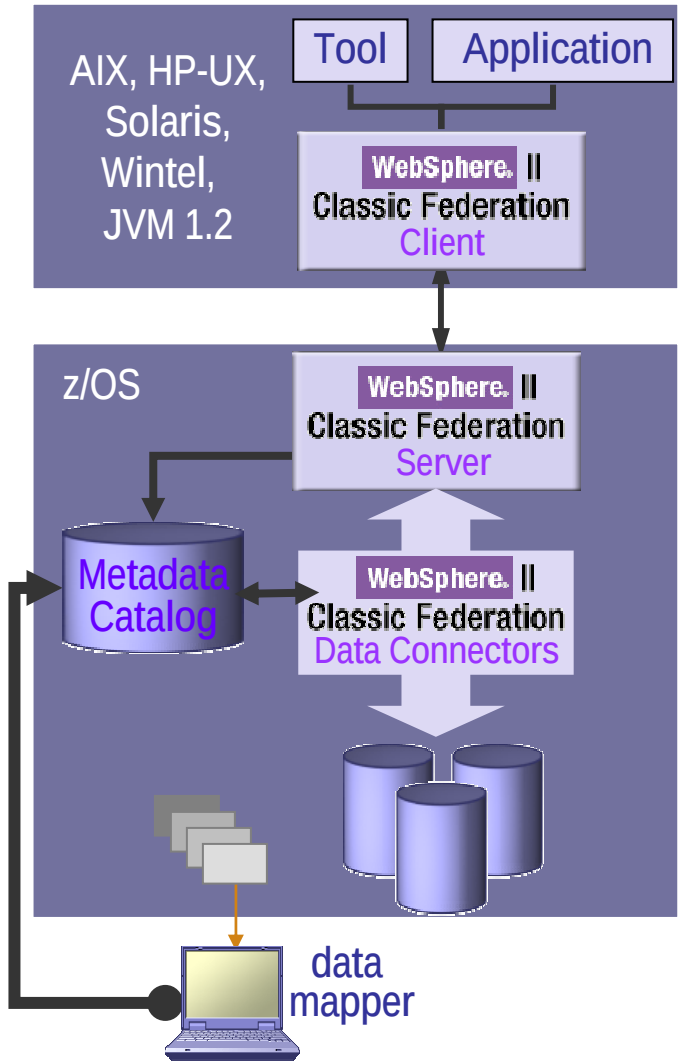

Usage scenarios

- e-Business
  - Deliver mainframe data to
    - Self-service portals (real-time account details)
    - e-commerce solutions (real-time inventory)
    - Employee portals (real-time claims detail)
  - Web developers become productive with no mainframe skills
  - Eliminates data latency business issues caused by copied data
- Business intelligence
  - Integrates seamlessly with
    - Reporting and analytical tools, e.g. Business Objects
    - Portals, e.g. WebSphere Portal
    - ETL, e.g. WebSphere DataStage
- Scenarios

1. Empower self-service environments with key operational data – IVRs, Web sites, Portals, etc.
2. Single-source mission critical information – Inventory-value, account-balance, available-credit, etc.
3. Feed operational data to business intelligence initiatives – Source for data marts, data warehouses, operational data stores

Kawasaki Motors Corporation

Kawasaki
Let the good times roll.

Vision

- Increase competitive position with over 8000 dealers. This includes outdoor power products (generators, lawn mower engines) and power sports products (motorcycles, ATV's, etc.).

Challenge

- Improve delivery of parts & accessories business to dealers from their regional warehouses with minimal disruption to existing applications & processes.

Solution

- WebSphere Information Integrator Classic Federation for z/OS (formerly CrossAccess XDI) provides the ability to join mainframe data sources with Microsoft SQL Server.

Value

- Cost effective, flexible and high performance solution
- Reduced requirements for ETL resources
- High level of operational reliability
- ***"We have saved one full year of development time using WebSphere Information Integrator Classic Federation for z/OS to provide up-to-the-minute information into our supply chain warehouse. "***

Victor Martinez

Order Fulfillment

Kawasaki
Let the good times roll.™

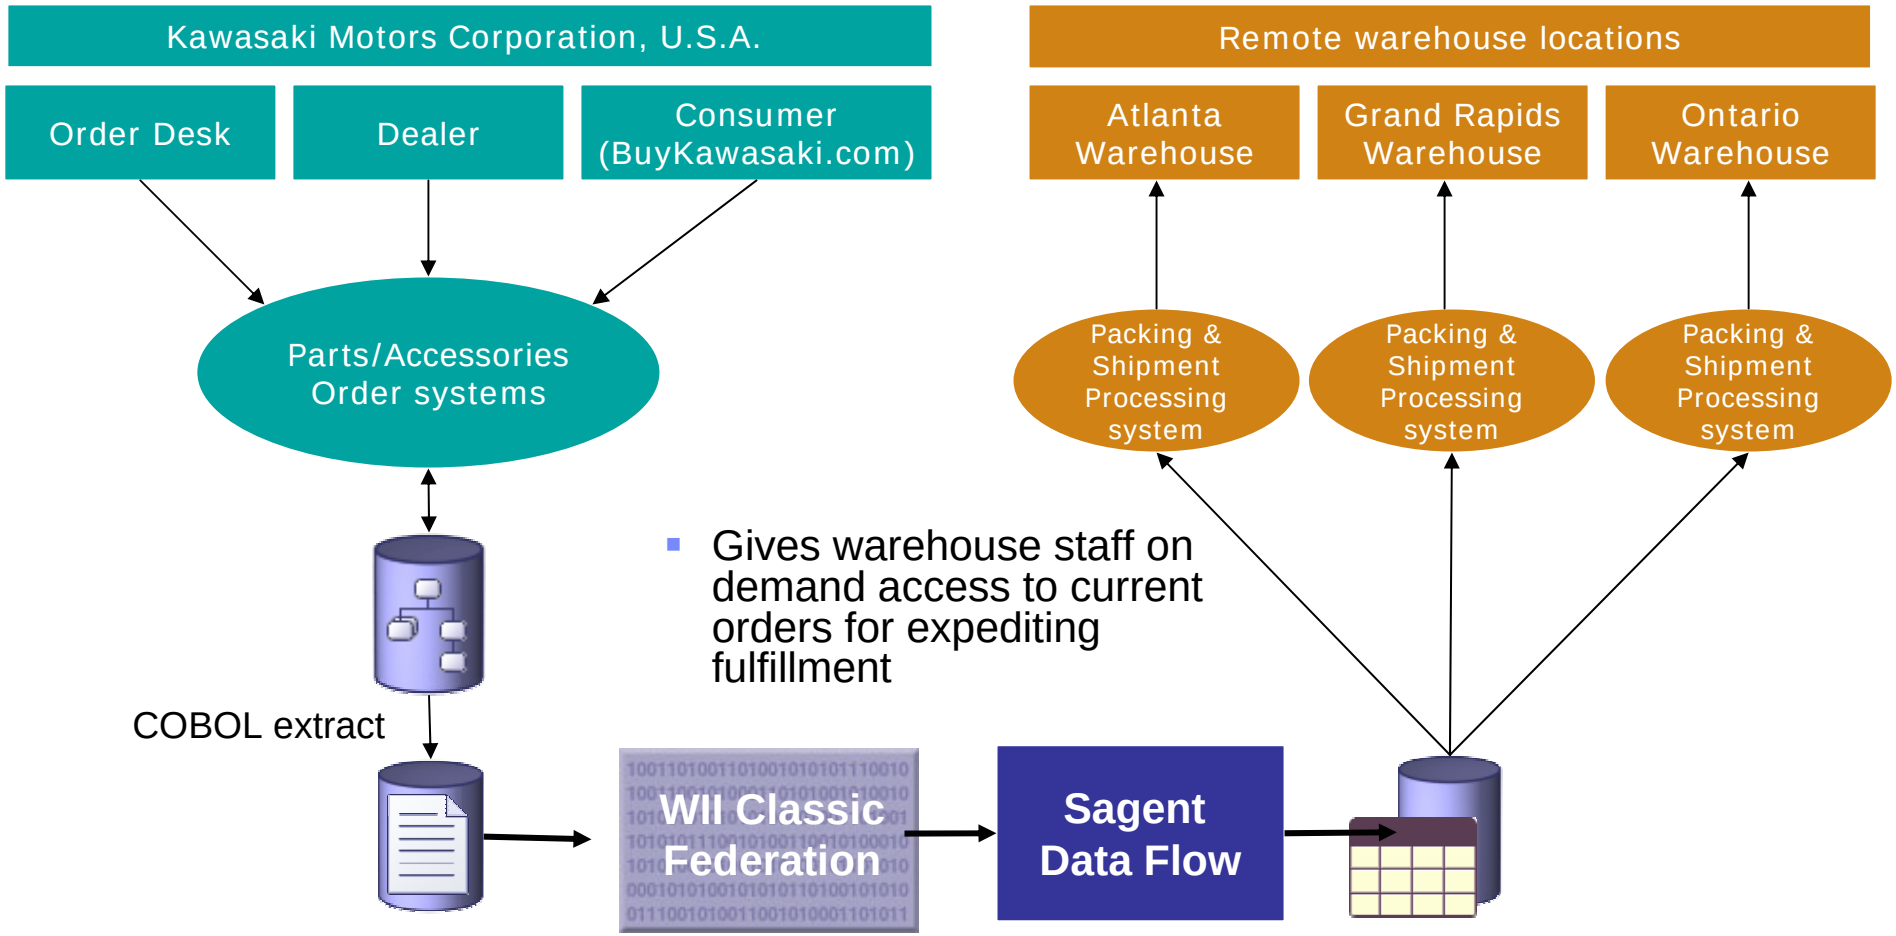

- Gives warehouse staff on demand access to current orders for expediting fulfillment

Banque Populaire
French bank Implements Basel II Solution

Challenge

- Sixth largest banking group in France needed to improve risk management across all member institutions and meet deadline compliance with Basel II.
- The bank wanted one solution which would encompass all of the new regulations, not just a few areas.
- They needed one solution to access information in 23 different retail systems from over 2500 branch offices.

Solution

- Data warehousing, business intelligence and integration solution including DB2 UDB, Intelligent Miner, WebSphere Information Integrator, WebSphere Business with pSeries, xSeries, and zSeries. Implemented with Premier Business Partner Computer Sciences Corporation

Business Benefits:

- Enhanced risk management and increased efficiency of data collection for Base II required data.
- Ability to view data in operational systems spread across the enterprise including third party information.
- No disruption to retail banking systems.

Technology Benefits:

- Best of breed solution for crucial data management functionality.

Neckermann (Netherlands) (Mainframe Database Connectivity)

- **Challenge:** As one of Europe's largest mail order organizations, Neckermann needed to be able to support increasing orders from their online store.

  All orders from the web were being re-entered into their mainframe IDMS-based order application.

- **Solution:** WebSphere Information Integrator Classic Federation provides connectivity to web application developed using WebSphere Application Server

- **Result:** Develop and support new channel without disrupting existing infrastructure. The new Web-based application has grown 60 times in just a few years. Labor savings in from not re-entering these online orders and all without requiring any dedicated IDMS resources or skills

Integration in Action

neckermann SHOPPING

- Seamlessly share order processing data and algorithms between legacy call-center systems & new Web-based “online store”

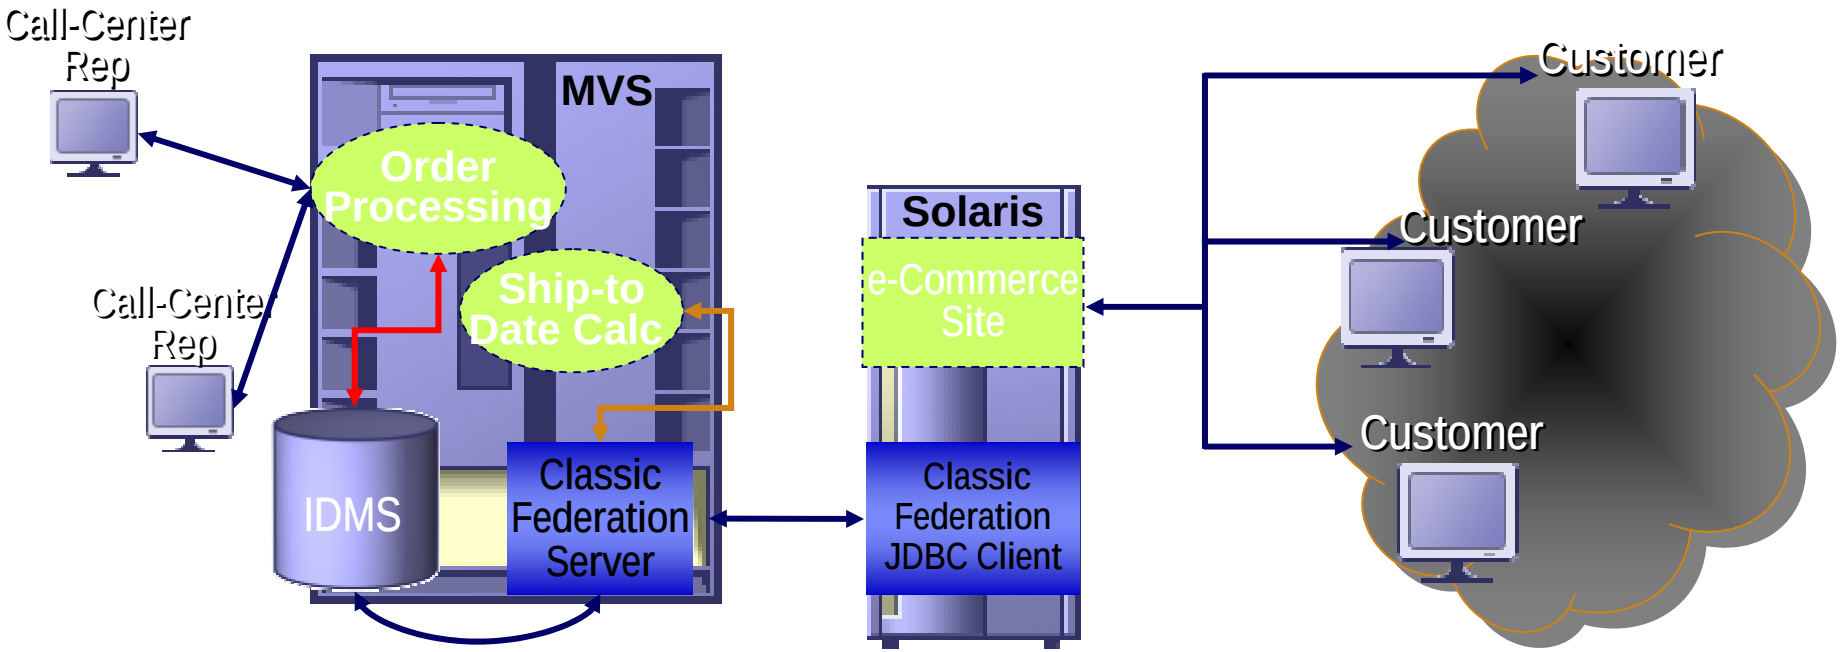

TSYS
technology provider for electronic payment transactions

Challenge

- Corporate wide migration toward more robust web-based reporting systems has created a need to access various incompatible heterogeneous data stores across the corporate infrastructure
- Key data is contained in mainframe sources (VSAM, QSAM, & IMS) is not readily accessible to users or existing distributed business applications

Solution

- The combination of WebSphere II Classic Federation (IMS, VSAM, and DB2), WebSphere II Advanced Unlimited, and DB2 Connect will provide access to data for application development efforts with .Net and Cold Fusion technologies which require access to enterprise data as if it were on one data source, thus simplifying application development and providing productivity

Business Benefits

- Getting accurate information from all of their systems in a timely manner without adding additional load to those systems.
- Save them capacity and money over existing 3rd party solution
- Improve data quality and data access from cross platform processes.

Technology Benefits:

- Successfully join information across VSAM, QSAM, Lotus Notes, SQL Server, and DB2 data sources

TSYS

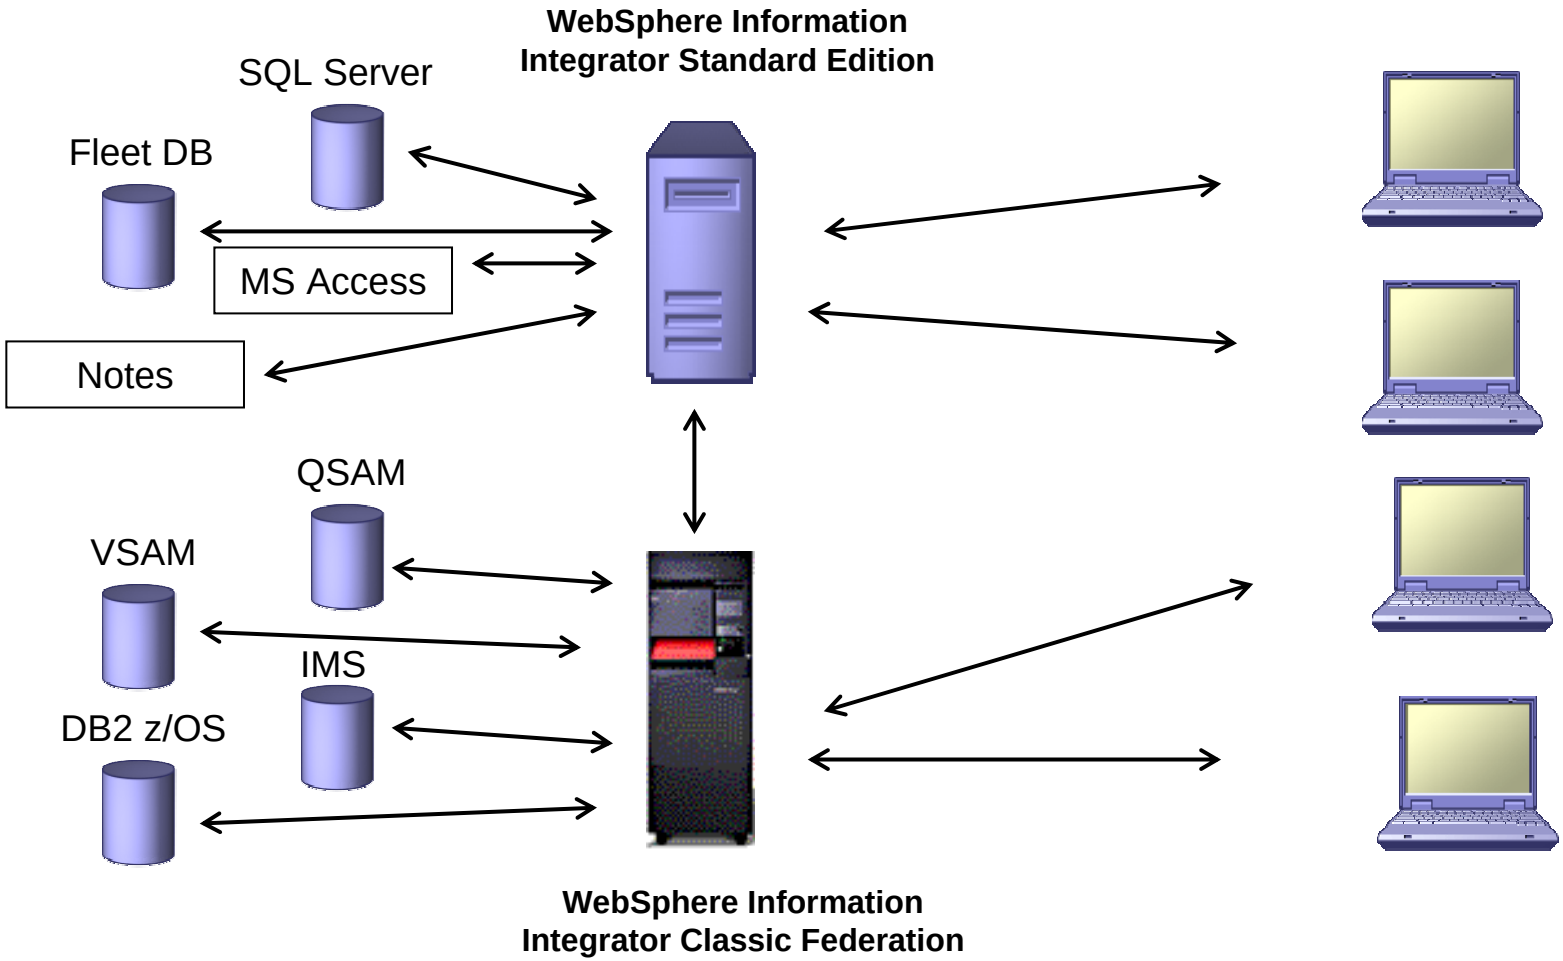
WebSphere Information
Integrator Standard Edition
SQL Server
Fleet DB
MS Access
Notes
QSAM
VSAM
IMS
DB2 z/OS
WebSphere Information
Integrator Classic Federation

Agenda

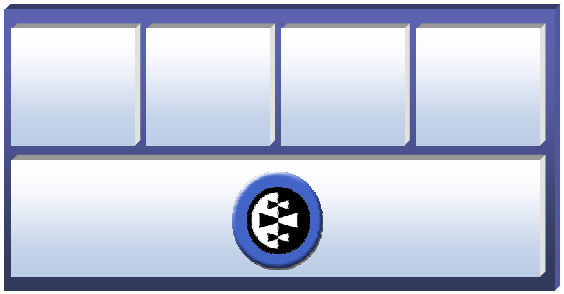

- IBM Information Integration Portfolio
- WebSphere II Replication
- WebSphere II Mainframe access & federation
- **Data Event Publishing**
- Wrap-up & Additional Questions and Answers

Slide 2 of 128 more to go

Data Event Publishing

Capture database changes as XML messages and publish them to WebSphere MQ

Function

- Publish events to a message queue
- XML self-describing format
- Active (real time) or recovery logs

Usage

- Application to application messaging
- Initiate business processes
- Source for ETL tool

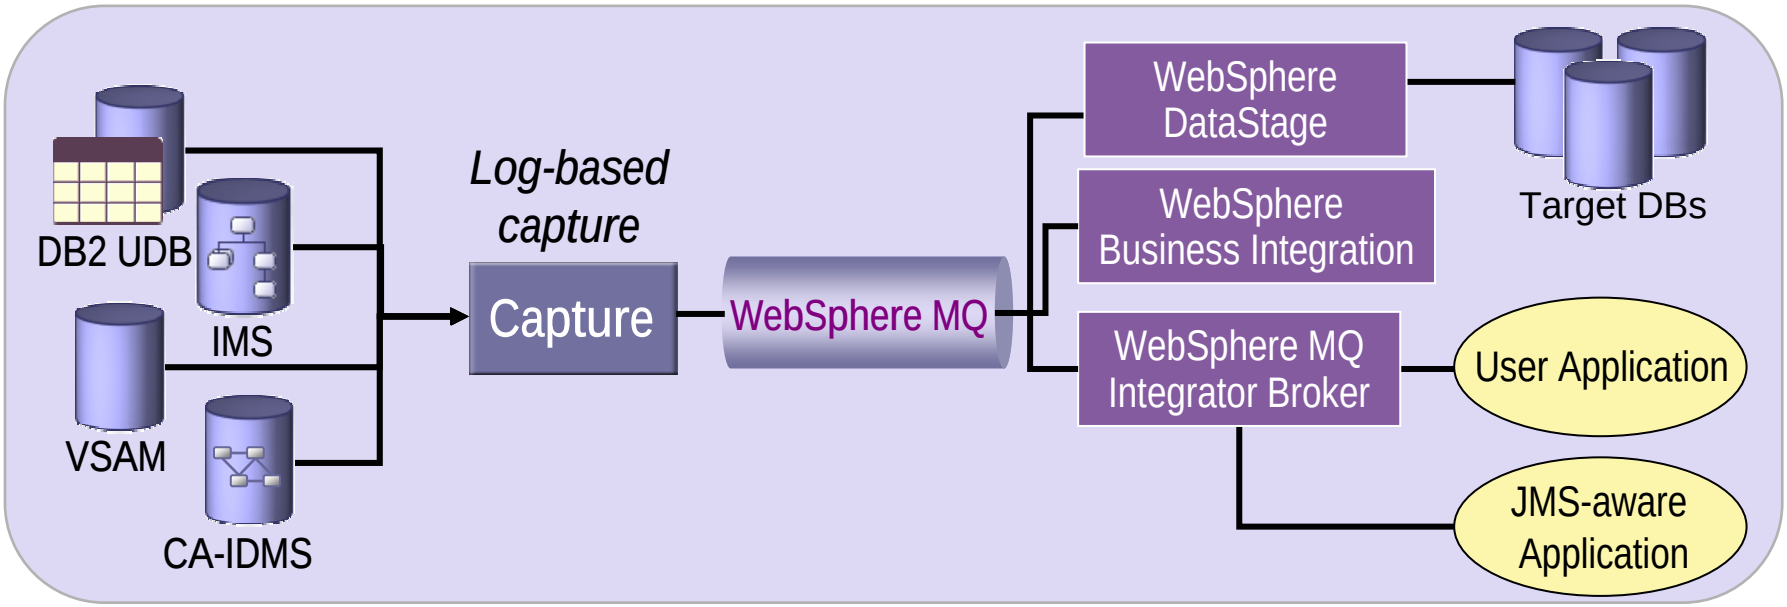

Why data event publishing?

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

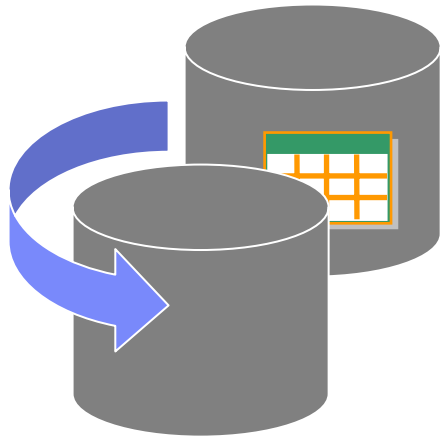

“Traditional” process integration can be complex

- **Business integration using application hooks**
  - **Complex**
    One hook for each process involved
    Many processes can impact the same data
  - **Maintenance intensive**
    Application changes typically impact integration

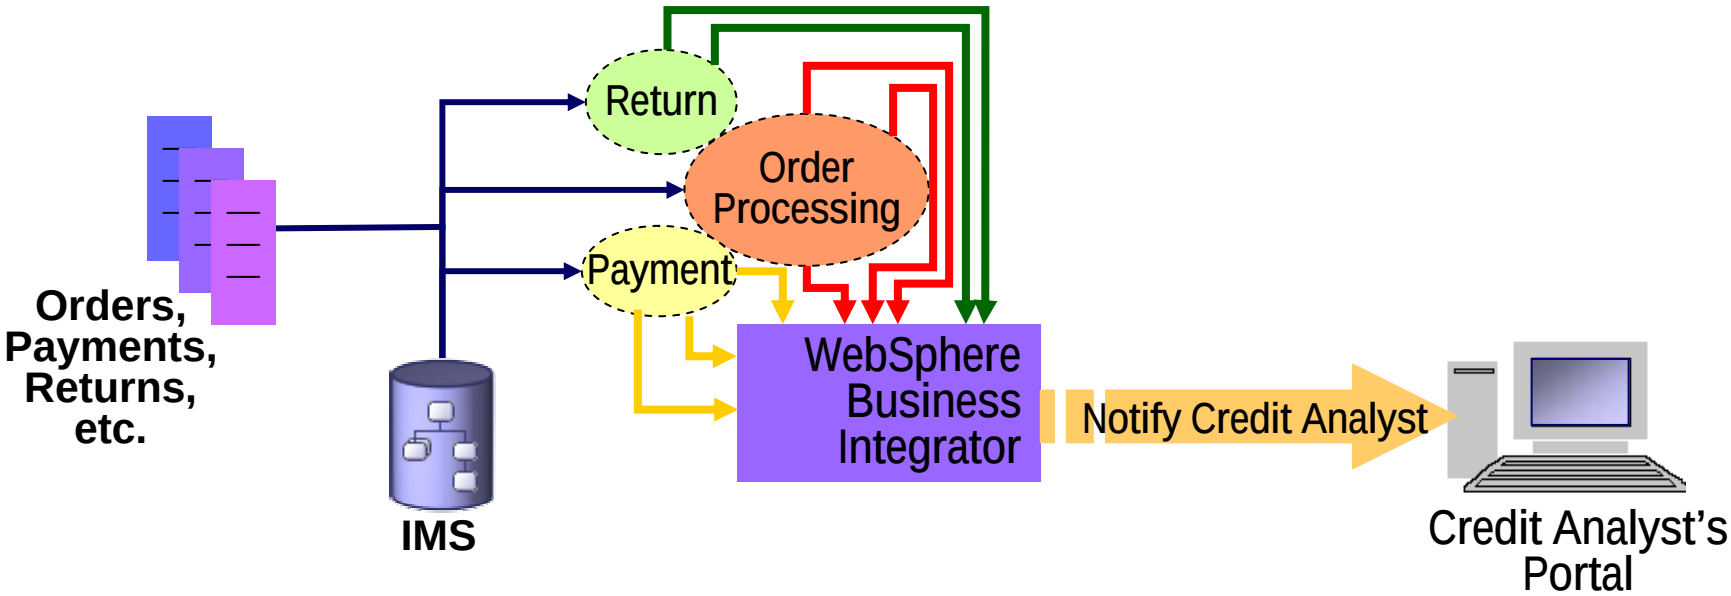

Data event publishing simplifies some integration
e.g. cross-application synchronization

- **Near real-time cross-silo data synchronization**
  e.g. New order data is automatically pushed to a CRM application
  e.g. VSAM employee data updates are pushed to SAP payroll
  - Loosely coupled integration
  - Minimizes development effort
  - Simplifies maintenance

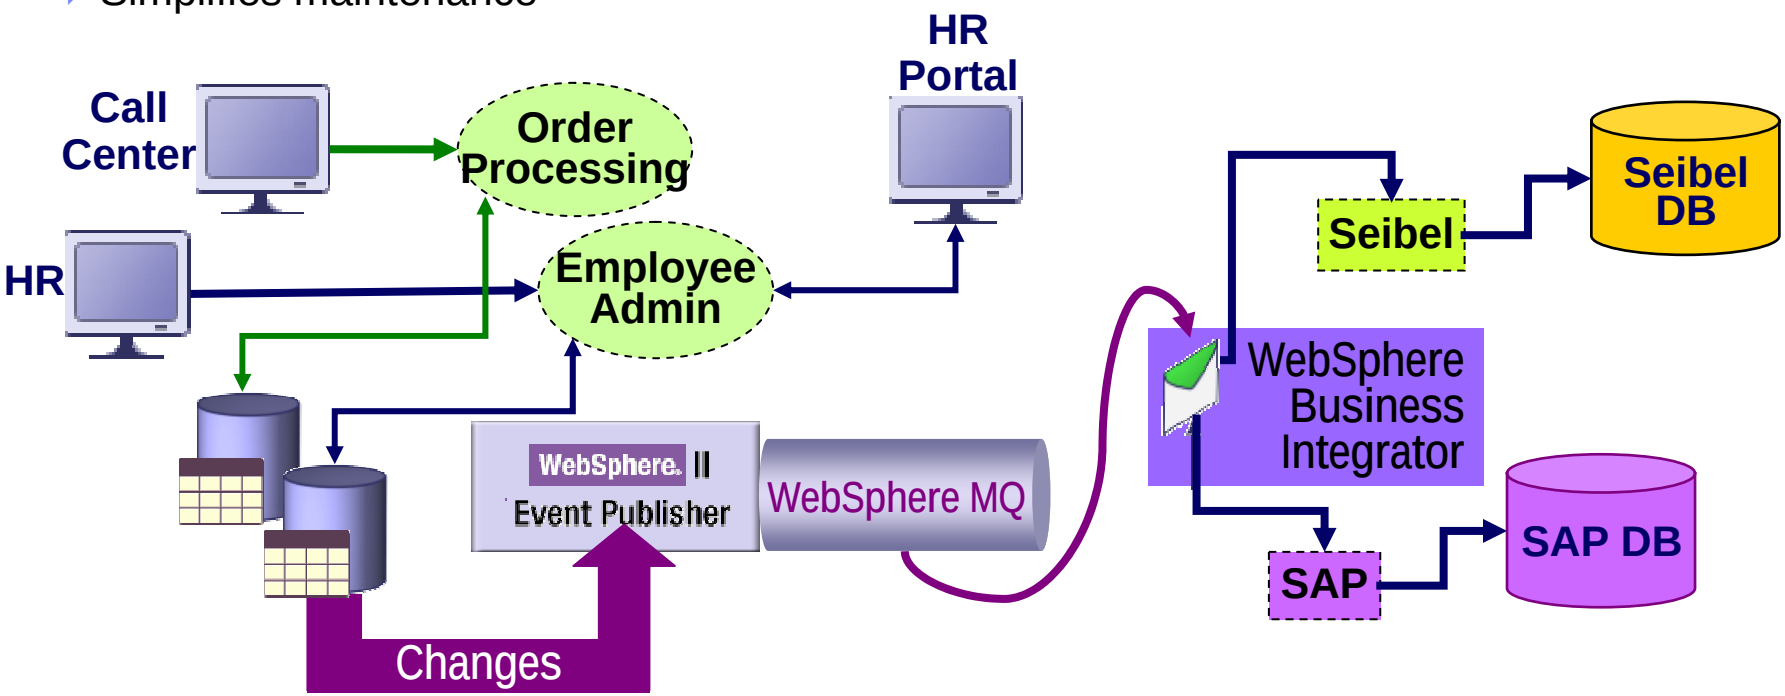

Data event publishing simplifies some integration
e.g. Customer Profile Management

- Basic customer profile management
  - WebSphere “listener application” picks up changes and initiates update
    - Directly to database
    - Using a database stored procedure
- Complex customer profile management
  - Push customer updates to WebSphere Business Integrator Event Broker
  - Appropriate updates or transactions initiated on other systems

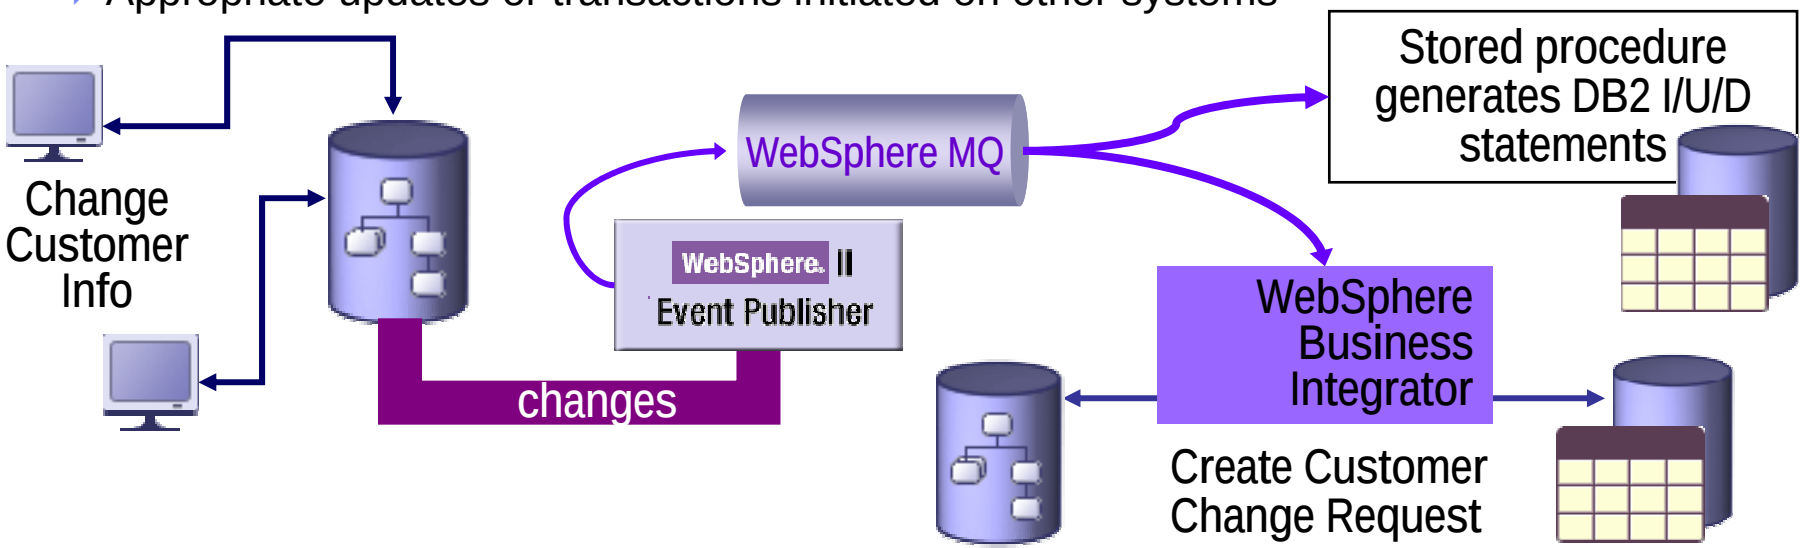

Why changed-data events versus traditional ETL?

- **ETL environments can't keep up with the data**
  - Shrinking batch windows demand ever larger "pipes" – no time for errors
  - Full data pulls are too large
- **Difficult to find only "the changes"**
  - Legacy data stores may not have date/time stamps
    - Wasted machine cycles searching legacy data stores
    - Wasted man-hours building legacy application hooks

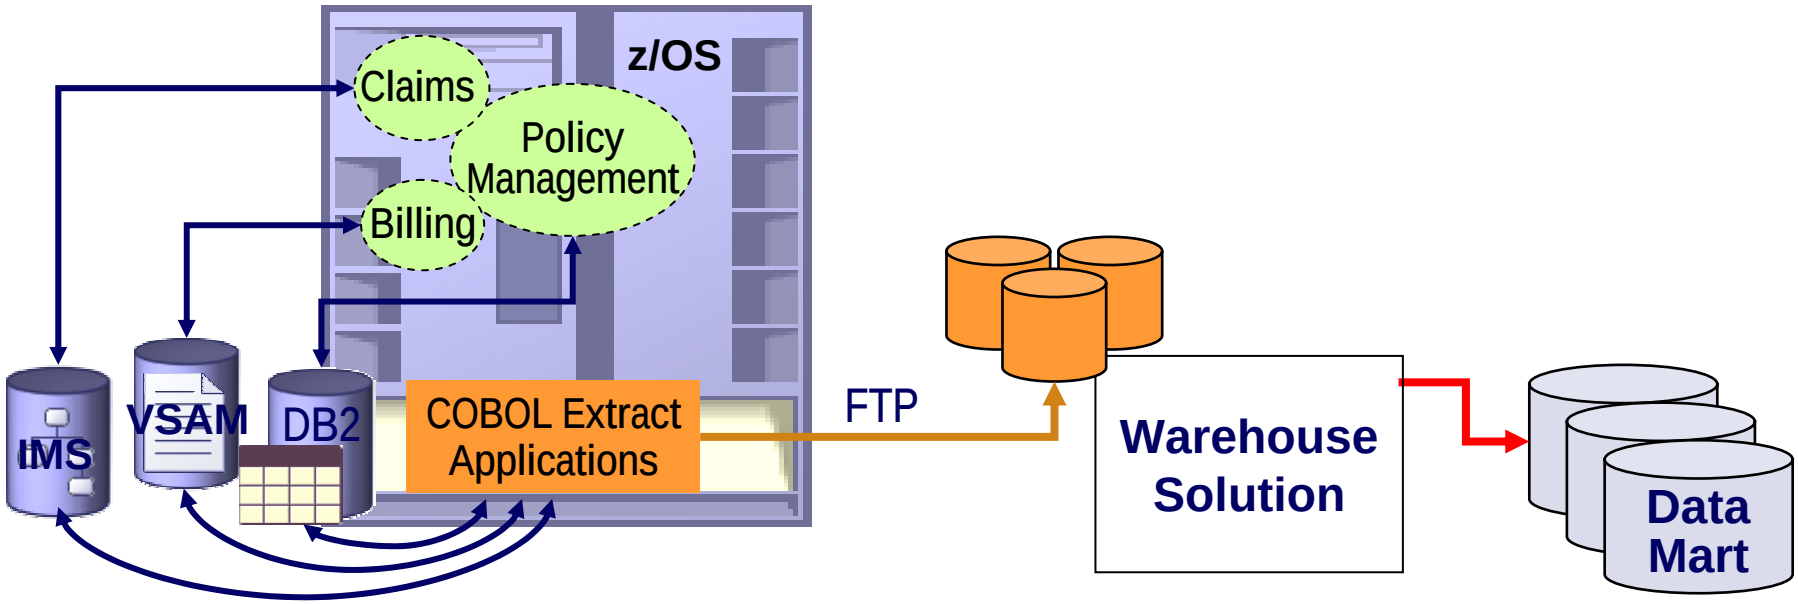

Feed operational data – Real Time "Push"

- **Dynamic, changed-data feed**
  - Maximize data currency while minimizing & stabilizing bandwidth utilization
- **Reliable and recoverable**
  - Recoverability built-in and WebSphere assures delivery

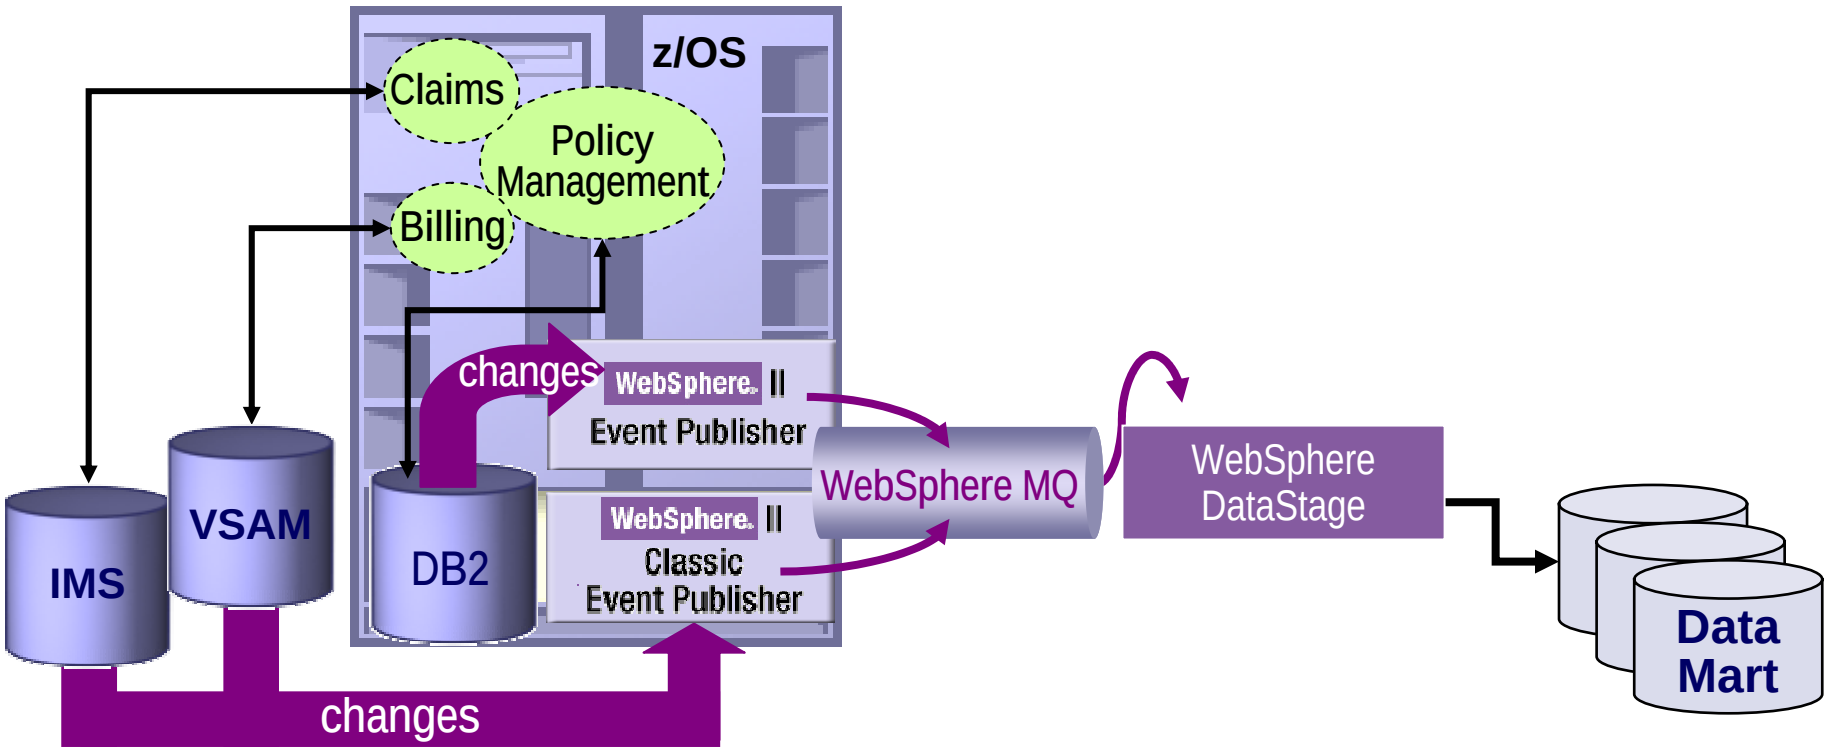

Value to the Business

- **Accelerate time-to-value of integration projects**
  - Fast and easy to implement
  - Reduces integration complexity
  - Leverages IBM & 3rd party tooling through WebSphere MQ
- **Extend the value of existing mainframe investments**
  - One consistent published data format regardless of source
  - Leverage operational data to drive business processes
  - Loosely coupled integration maintains application independence
- **Complements traditional ETL approaches**
  - Feeds data integration & ETL solutions from IBM & 3rd parties like Informatica
  - Minimize bandwidth requirements
  - Eliminates dependence on “shrinking” batch window
  - Enables near real-time updating when this is needed

Summary

- Q Replication delivers low latency, high throughput, completely reliable replication
  - Commercial WebSphere MQ transport
  - Highly parallel apply
  - Peer-to-peer conflict detection and resolution
  - Consider Q Replication for your High Availability needs
- Classic Federation provides transactional mainframe data access and federation for application modernization
  - Leverage your mainframe data investments as you evolve and grow
  - Minimize demand for specialized skills
  - Simplify maintenance while maximizing efficiency of resource utilization
- Event Publishing provides high speed linkage of applications, without change or impact to those applications
  - First DB2, IMS, CICS VSAM and CA-IDMS - then more!
  - Consider event publishing to establish or add to an existing Service Oriented Architecture

More information…

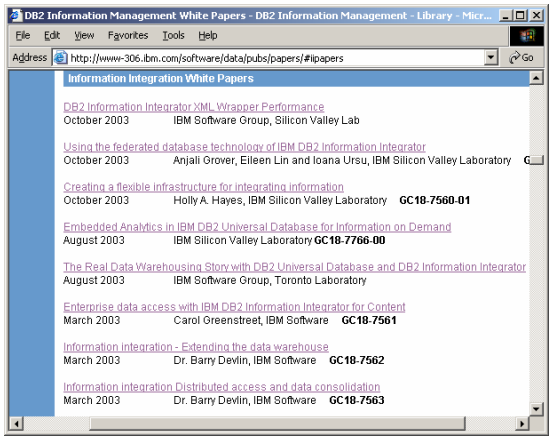

White papers

www.ibm.com/software/data/pubs/papers/

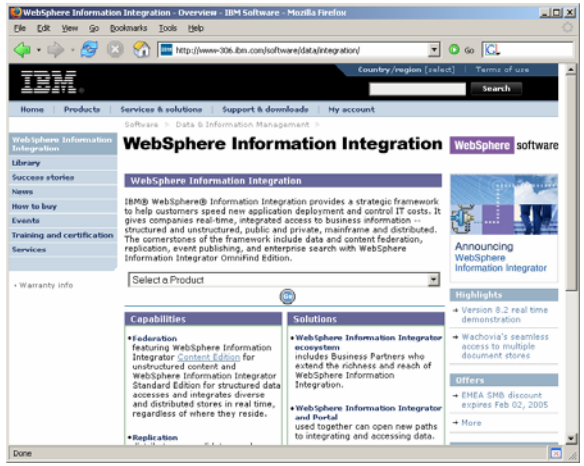

Home page

www.ibm.com/software/data/integration

James Reed
Product Mktg Mgr
Info Integration Solutions
jwreed@us.ibm.com


Thank
You