



Extending Your Mainframe For More Business Value

Add New Workload –
Data Warehouse On System Z

Get More Business Results Out of Your Data

Our branch offices have separate databases.
Each branch is analyzing customers and sales on their own.



**Service Oriented Finance
Marketing**

Looking at data in isolation can miss larger trends and opportunities




IBM

Service Oriented Finance Needs A Data Warehouse To Make Optimal Business Decisions

- Each branch is responsible for its own marketing campaign
- Corporate marketing gets reports from each of the branches based on local results
- Corporate marketing needs to spot trends to know what campaigns are most effective region-wide
- A corporate data warehouse would give marketing the data to easily do comparisons between the branches and promote best practices


Deliver Business Insight With A Data Warehouse



Compare Scalability Achievements: Winter Corporation's "Top Ten" Awards

- “The highest performing transaction processing system in the 2005 program, a [DB2] z/OS implementation, executed **over one billion** SQL statements in an hour. The average for operational systems was **35 million** SQL statements or database operations per hour”
 - The study lists the largest known peak workload on **Oracle RAC** to be **8.6 million** SQL statements per hour
-
- “The largest transaction processing [database] in the program, **23 TB**, was hosted on [DB2] z/OS, as in the last program”
 - The study lists the largest transaction processing database on **Oracle RAC** to be **9.6 TB**

Leading OLTP System,
DB2 for z/OS, Processed
Over One *Billion* SQL
Statements per Hour



http://www.wintercorp.com/VLDB/2005_TopTen_Survey/TopTenWinners_2005.asp
http://www.wintercorp.com/WhitePapers/WC_TopTenWP.pdf

IBM InfoSphere Information Server For System z Can Load Your Data Warehouse

IBM InfoSphere Information Server for System z

Understand



Information Analyzer
on Linux for System z

Business Glossary
on Linux for System z

Rational Data Architect

Cleanse



QualityStage on
Linux for System z

Transform



DataStage on Linux
on for System z
DataStage for z/OS

DataStage MVS

Deliver



Classic Fed. Server
Classic Replication
Replication Server
Change Data Capture

Platform Services

Parallel Processing Services



Connectivity Services



Metadata Services

Metadata Server

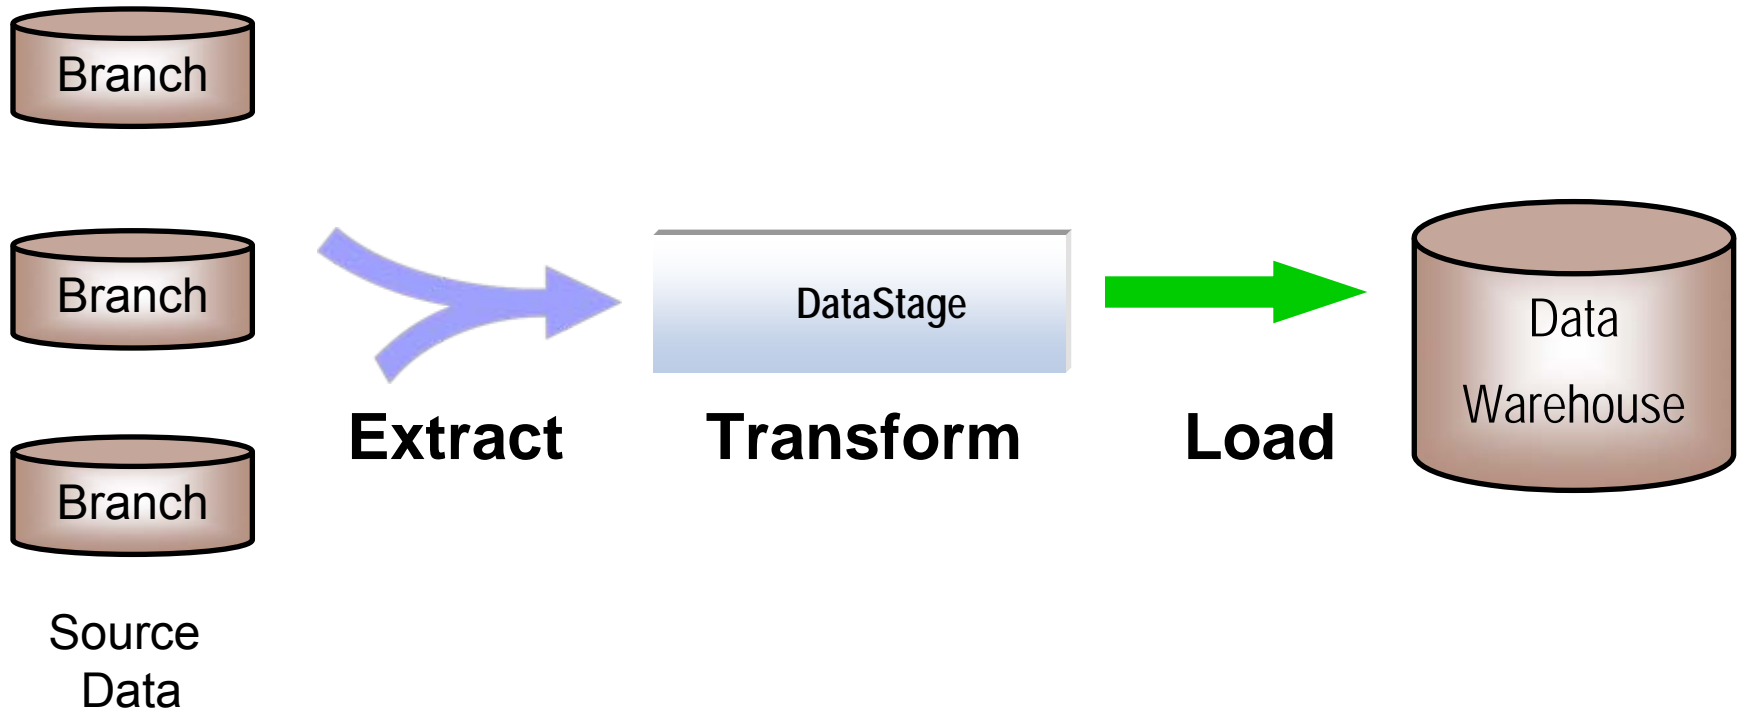
Administration Services



Deployment Services

Information
Services Director

Load Your Data Warehouse With DataStage



Data Stage Transforms Data On The Fly

- Different field names
- Different field order
- Add Branch Identifier
- Different currency format



PROD ID	CUST ID	BRANCH ID	QTY	AMT	SALEDATE
000 101	100	01	01	10,000.00	2007-02-28
000 121	100	01	03	500.50	2007-02-28
000 101	101	01	01	20,000.00	2007-03-01

Data Warehouse



Transform

PRODUCT	QTY	CUSTNO	AMOUNT	DATE


Branch Data

DEMO: Extract, Transform, Load

Use DataStage to load sales and customer data from a branch into your data warehouse

- ▶ ODBC (Input) – Branch Sales info from SQL Server
- ▶ DB2 (Update) – Corporate Sales Warehouse

Show how built-in stages make it easy to handle transformation and aggregations



BlueCross BlueShield Of Tennessee

Challenge

- In order to compete effectively for new business in the complex healthcare market, BCBST needed to differentiate itself from competitors with targeted offerings
- They needed a single view of information across their multiple LOBs with business intelligence capabilities

Solution

- BCBST is using **IBM InfoSphere Information Server** with **IBM DB2** to profile, transform, and load data to their enterprise data warehouse
- The solution also provides intelligent search capabilities for unstructured data using **IBM OmniFind** and **IBM DB2 Content Manager**



Business Benefits

- Improved customer and provider satisfaction by enabling them to run their own analytics and better manage their healthcare costs.

IBM InfoSphere Information Server For System Z Has Connectivity And Scalability Suitable For The Large Enterprise

IBM InfoSphere Information Server for System z

Understand



Information Analyzer
on Linux for System z

Business Glossary
on Linux for System z

Rational Data Architect

Cleanse



QualityStage on
Linux for System z

Transform



DataStage on Linux
on for System z

DataStage for z/OS

DataStage MVS

Deliver



Classic Fed. Server

Classic Replication

Replication Server

Change Data Capture

Platform Services

Parallel Processing Services



Connectivity Services



Metadata Services

Metadata Server



Administration Services



Deployment Services

Information Services Director


Data From 3rd Party Systems Can Be Integrated Using Connectivity Services



IBM InfoSphere Information Server Connects To Almost All Sources Of Data


RDBMS

DB2 (on Z, I, P or X series)
Oracle
Informix (IDS and XPS)
Ingres
MySQL
Nettezza
Progress
RDB
RedBrick
SQL/DS
SQL Server
Sybase (ASE and IQ)
Teradata
Universe
UniData
NonStopSQL
And more.....




General Access

Sequential File
Complex Flat File
File / Data Sets
Named Pipe
FTP
Compressed / Encoded Data
External Command Call
Parallel/wrapped 3rd party apps
EMC InfoMover
Web logs
Unstructured: e-mail, docs, etc.
Content Management Systems
Life Sciences




Enterprise Applications

JDE/PeopleSoft EnterpriseOne
Oracle Applications
PeopleSoft Enterprise
SAS
SAP R/3 and BI
SAP XI
Siebel
JDA
Ariba
Manugistics
I2
And more...



Standards and Real Time

WebSphere MQ
Java Messaging Services (JMS)
Java
XML and XSL-T
EBXML
Web Services (SOAP)
Enterprise Java Beans (EJB)
EDI
FIX
SWIFT
HIPAA



CDC / Replication


DB2 (on Z, I, P, X series)
Oracle
SQL Server
Sybase
Informix
IMS
VSAM
ADABAS
IDMS
NonStopSQL
Enscribe

Legacy

Allbase/SQL
C-ISAM
D-ISAM
Datacom/DB
DS Mumps
Enscribe
Essbase
FOCUS
IDMS/SQL
ImageSQL
Infoman
KSAM
M204
MS Analysis
Nomad
Nucleus
RMS S2000
Supra
TOTAL
TurboImage
Unify
And many more....

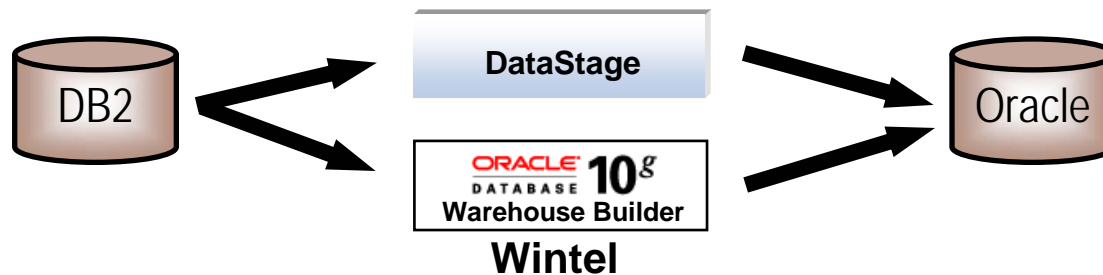


DataStage Utilizes Parallel Processing Services For Extreme Scalability



- Provides automatic dynamic data partitioning, repartitioning, and pipelining for optimal parallel performance
- Design integration processes without concern about underlying hardware architecture or number of processors
 - ▶ Resources defined in a separate configuration file
 - ▶ Allows easy expansion to new hardware
- Benefits from processing capacity, I/O capacity, and Hipersockets on System z


DataStage Outperforms Oracle Warehouse Builder



Results of a comparative study testing a simple ETL scenario involving database to database operations using Intel based servers

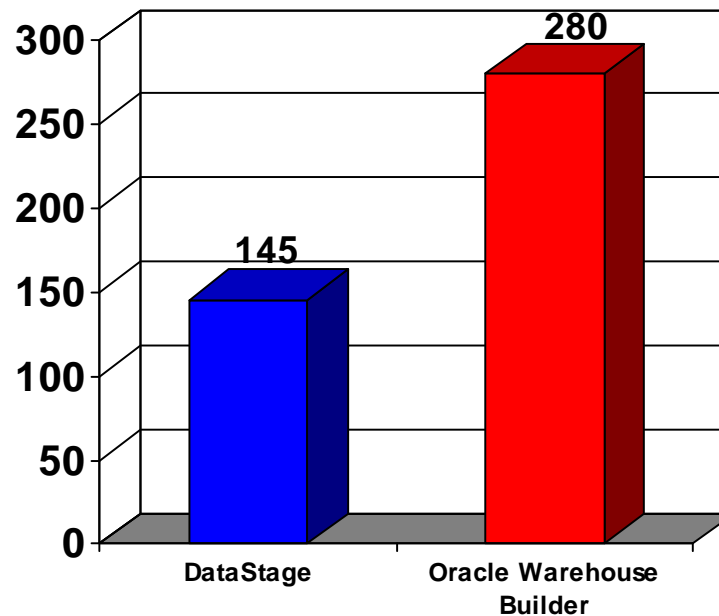
ETL Performance

Records/Sec (bigger is better)

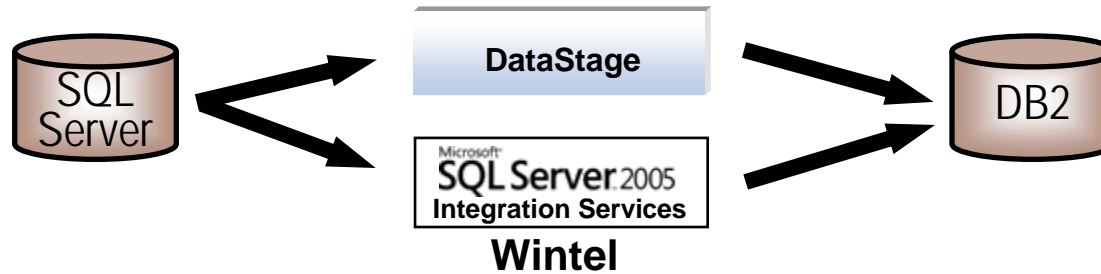


Development Time

Minutes (smaller is better)

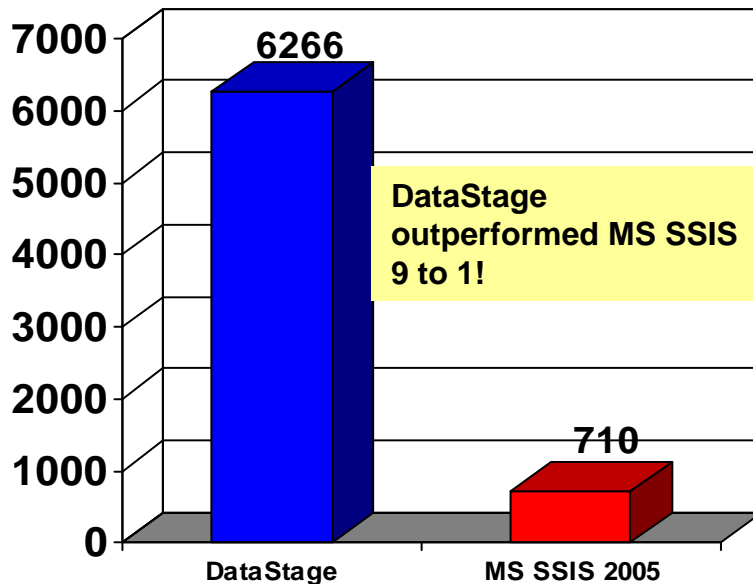


DataStage Outperforms Microsoft SQL Server 2005 Integration Services

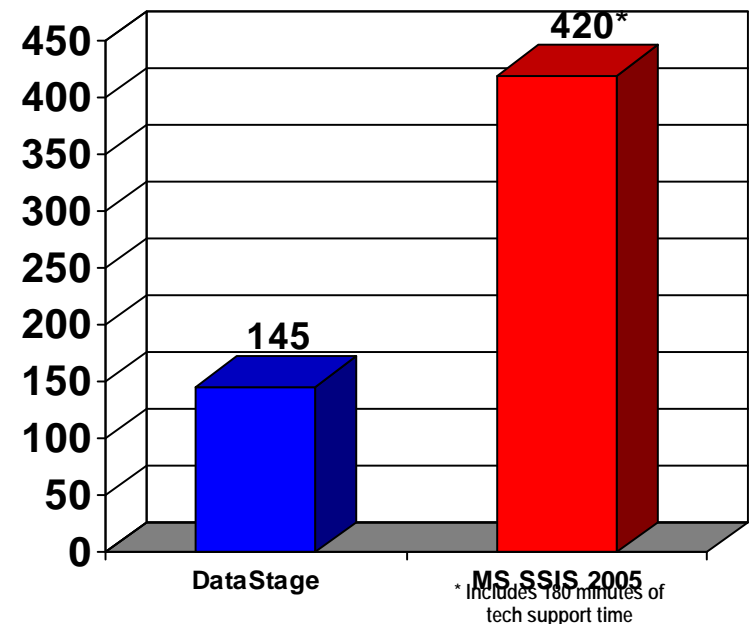


Results of a comparative study testing a simple ETL scenario involving database to database operations using Intel based servers

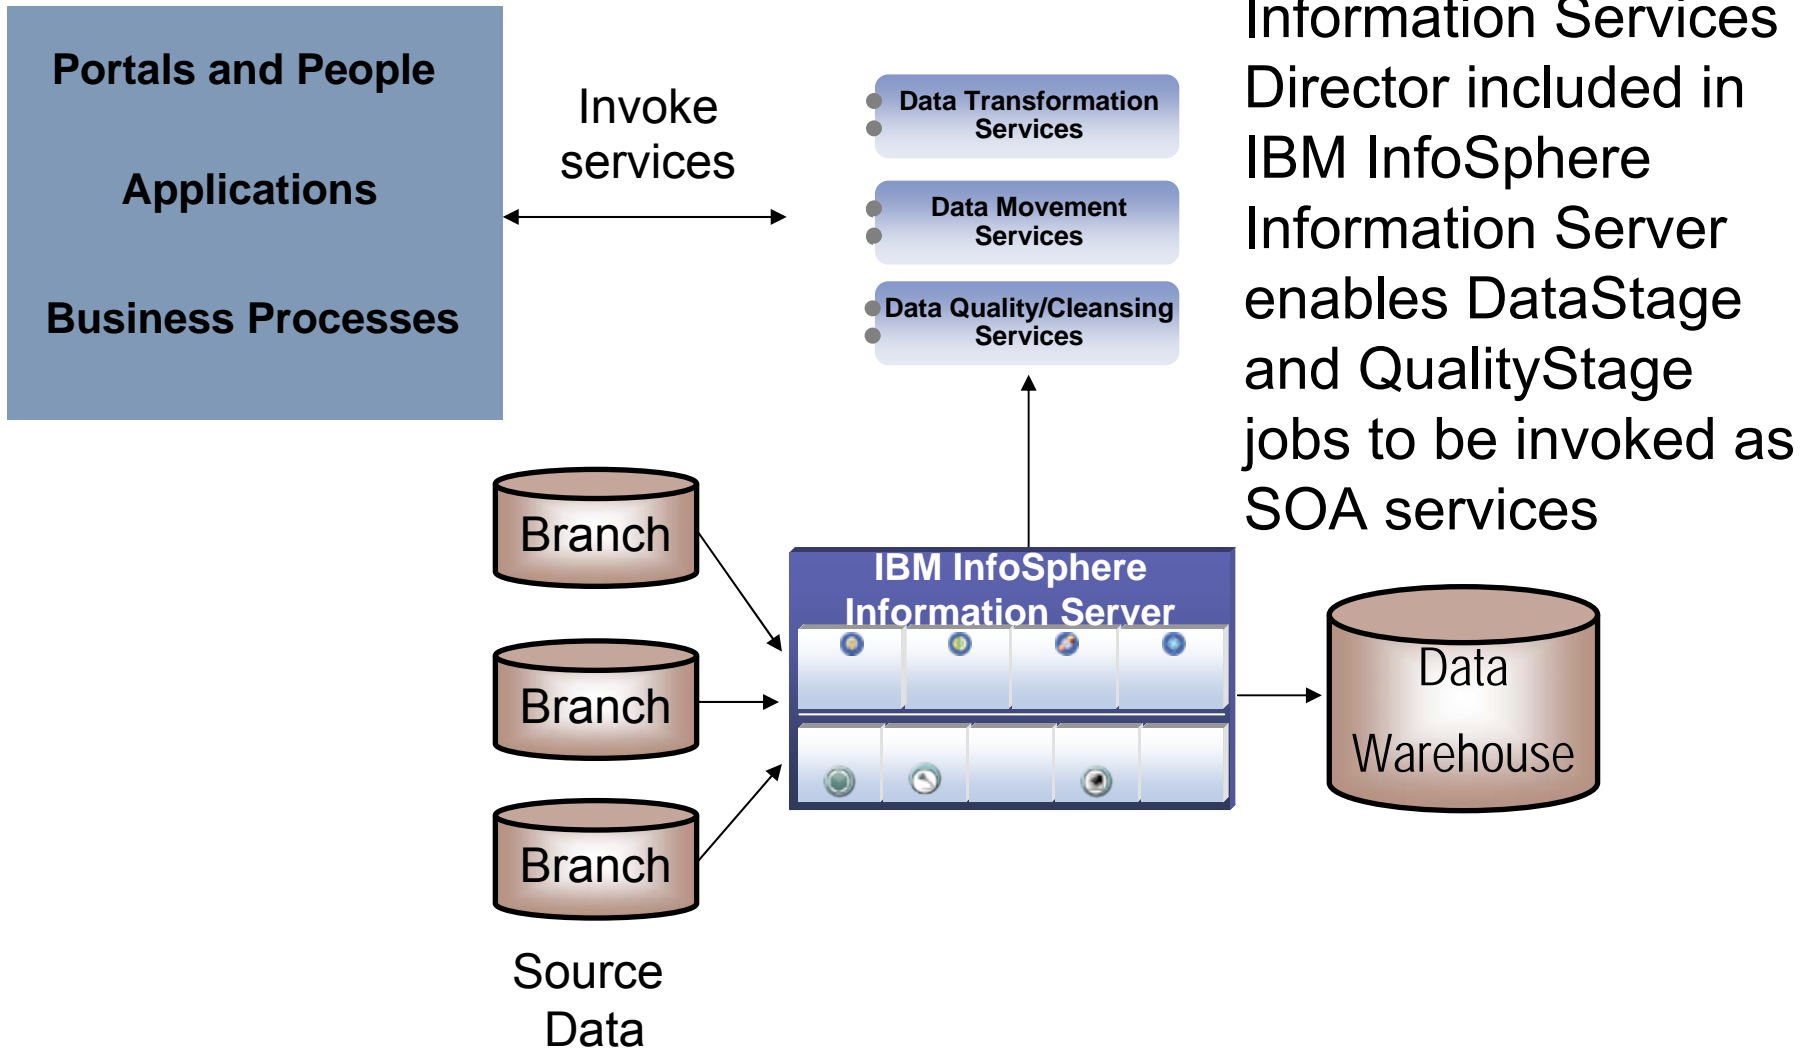
ETL Performance
Records/Sec (bigger is better)



Development Time
Minutes (smaller is better)




IBM InfoSphere Information Server Exposes Jobs As SOA Services



IBM Beats The Competition In Data Warehouse Solutions

IBM




- Integrated data services platform
- Extensive connectivity
- Enterprise scalability
- Easily expose jobs as services

Oracle



- Two separate un-integrated products
- Warehouse Builder can only load Oracle databases
- Oracle Data Integrator (used to be Sunopsis) has no data quality capabilities
- Coding required to expose jobs as services

Microsoft



- Data quality limited to "Fuzzy Search" and MS SQL only
- Limited connectivity and limited support for non-Microsoft
- Lacks enterprise scalability
- Coding required to expose jobs as services

Service Oriented Finance Wants A *Dynamic* Data Warehouse

Our data is updated frequently,

We need the data in our data warehouse to be more current



Change Data Capture together with Information Server can “trickle feed” data into your data warehouse as it changes



IBM

IBM InfoSphere Information Server For System z Can Provide Near Real-Time Data Movement

IBM InfoSphere Information Server for System z

Understand



Information Analyzer
on Linux for System z

Business Glossary
on Linux for System z

Rational Data Architect

Cleanse



QualityStage on
Linux for System z

Transform



DataStage on Linux
on for System z

DataStage for z/OS

DataStage MVS

Deliver



Classic Fed. Server
Classic Replication
Replication Server

**Change Data Capture
for System z**

Platform Services

Parallel Processing Services



Connectivity Services



Metadata Services

Metadata Server

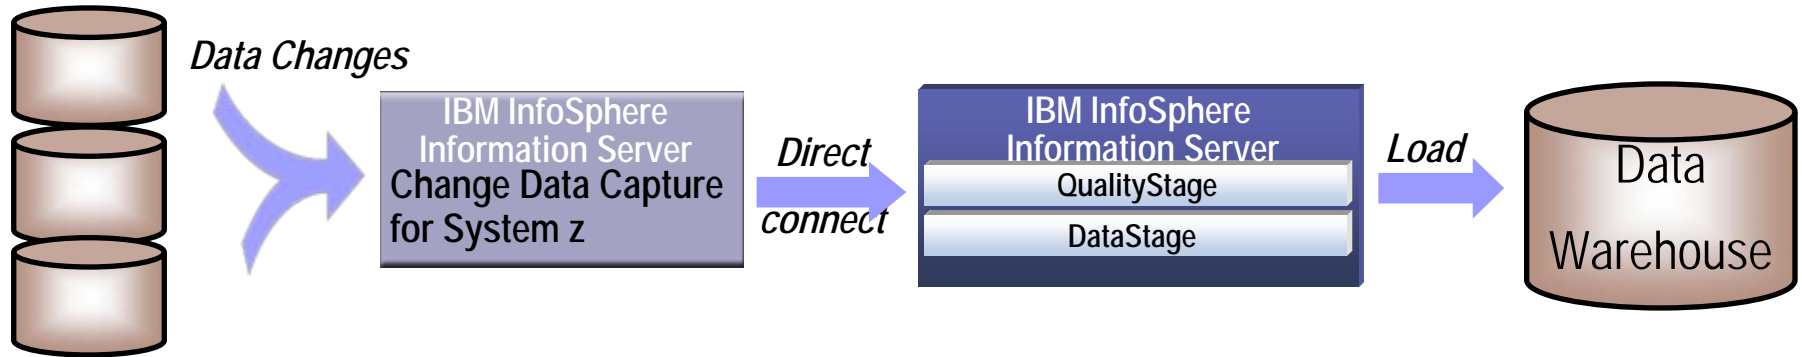
Administration Services



Deployment Services

Information
Services Director

Change Data Capture For System z



- Event is triggered automatically by a change in the data
 - ▶ Monitors DB2 log file for data changes
 - ▶ Sends message with data changes
- Message can initiate a business process
 - ▶ Example: After \$10,000 in sales, a service request is sent to WebSphere Information Services Director, invoking a DataStage job to load sales data for analysis
- Automates the process of loading the data warehouse
 - ▶ Improves the currency of data in the data warehouse
 - ▶ Alternative to statically scheduled updates - which can lead to stale data and lost opportunities

Log-Based Change Data Capture

- **Using existing native database recovery logs to capture table and row level activity in a relational database**
 - ▶ **Most databases have a native log format that can be accessed**
 - DB2 and DB2 z/OS = DB2 Log, DB2(i5)=OS/400 Journal
 - Oracle = Re-do Log, SQL Server = Transaction Log
- **Many advantages to log-based Change Data Capture**
 - ▶ No changes to existing applications or schemas required
 - ▶ Little performance impact to source application and system
 - 0.05% system resources required to process over 300 GB of data
 - Sends only the changes – efficient use of bandwidth
 - Scalable
 - ▶ Alternatives have many drawbacks
 - SQL Select, File Comparison, Database Triggers, Modifying Source Application
- **Wide range support**
 - ▶ Information Server CDC for - DB2 (all platforms), Oracle, SQL Server, Sybase, and more
 - ▶ Data Event Publisher for non relational sources

Business Optimization

We'd like to make it easy for managers to closely monitor how our bank branches are performing



**Service Oriented Finance
CIO**

This is an example of Performance Management and you can do this with IBM Cognos BI

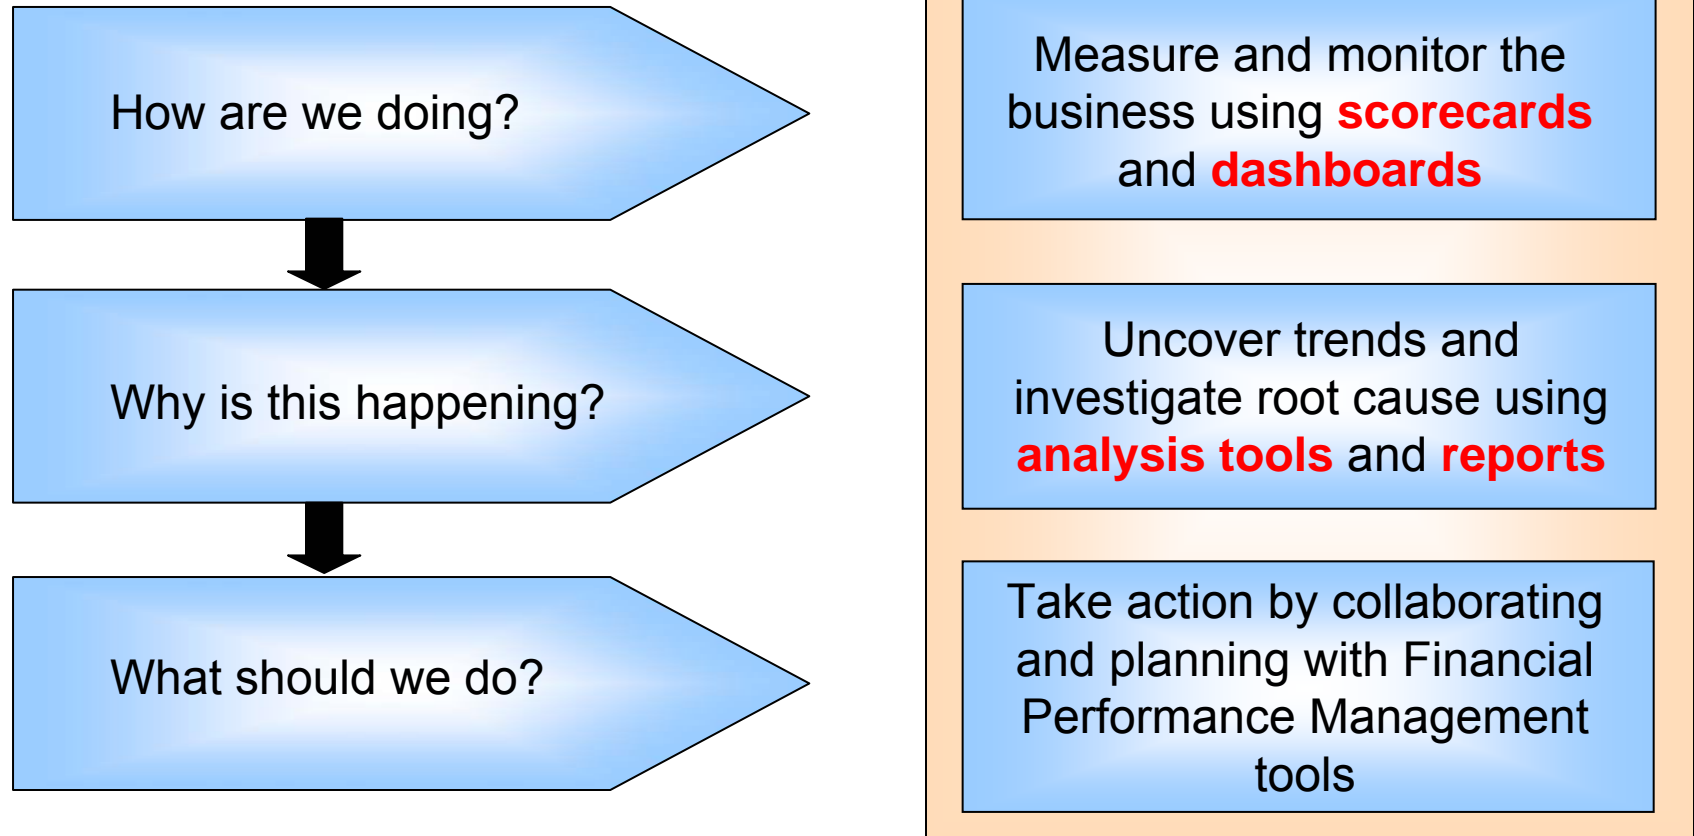


IBM

Achieve Business Optimization Using Performance Management

- Make effective use of information to improve decision making
- Enhance processes to improve business results
 - ▶ Define strategic goals, then measure and manage performance to ensure the goals are achieved
 - ▶ Link employee objectives with organizational objectives
- Examples of Performance Management
 - ▶ Monitoring and managing key performance indicators
 - ▶ Reporting and analyzing strategic information
 - ▶ Creating, monitoring, and managing financial plans

IBM Cognos Provides an Integrated Platform for Performance Management



Server based business intelligence – only requires a browser

Use IBM Cognos BI Scorecards And Dashboards To Answer The Question “How are we doing?”

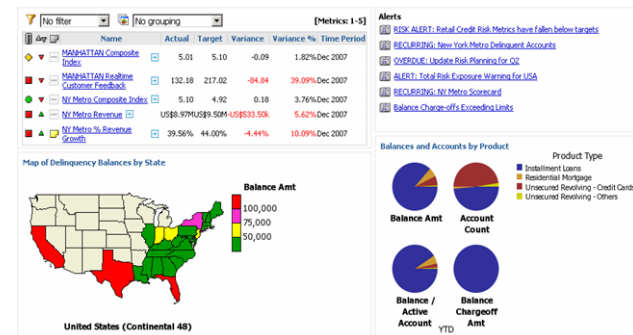
■ Scorecards track progress towards objectives

- ▶ Metrics display results relative to targets
 - Defined on any indicator of business performance
 - Actual results imported on a scheduled basis
- ▶ Scorecards organize metrics logically
- ▶ Diagrams organize metrics visually
- ▶ Take action using action items, projects, comments, and watch lists

	Name	Actual	Target
● ▲	NY Metro Profitability Index	65.82	62.81
● ▼	NY Metro Composite Index	5.10	4.92
● ▼	NY Metro Customer Retention/Attrition	2,085.87%	2,077.21%
◆ ▼	NY Metro Realtime Customer Feedback	198.00	200.00
■ ▲	NY Metro Revenue	US\$8.97M	US\$9.50M
■ ▲	NY Metro % Revenue Growth	39.56%	44.00%

■ Dashboards give quick access to information

- ▶ Present at-a-glance, high impact views of relevant information
- ▶ Use visual indicators such as gauges and dials
 - Drill-down into details
- ▶ No additional licensing requirements to create dashboards with Cognos 8 BI



Use IBM Cognos BI Analysis And Reporting To Answer The Question “Why is this happening”?

- **Use Online Analytical Processing (OLAP) analysis to gain insight**
 - ▶ **Cognos Analysis Studio** provides simple and accessible multidimensional analysis
 - ▶ Powerful comparative analysis features such as automatic time trending and mixed-grain detail
 - ▶ Works with many OLAP and “dimensional relational” data sources
- **Create, distribute, and manage any style of report**
 - ▶ **Cognos Query Studio** provides “self service” reports created as needed by novice users
 - ▶ **Cognos Report Studio** provides “pixel perfect” production reports for payroll, invoices, financials
 - ▶ Supports many report formats including HTML, PDF, Excel, CSV, XML
 - ▶ Creates reports from both relational and OLAP data
 - ▶ Supports Unicode for working with reports in multiple languages
- **Easily manage the full reporting lifecycle**
 - ▶ Schedule analysis and reports and “burst” to many users
 - ▶ Users configure alerts and watch rules to automatically receive updated content via E-Mail or portal “News Items”

IBM Cognos Allows an Open Data Strategy

- All capabilities access a consistent set of information defined in the Cognos metadata model
 - ▶ As backend data sources change, the Cognos metadata model can be updated to control these changes and can identify impacts to deployed reports
- Open Data Access allows *any combination* of data sourcing strategies (e.g. combining ETL with Federation)



REPORTS



ANALYSIS



SCORECARDS



ERICO International

International manufacturer builds enterprise data warehouse with IBM

Challenge

- Information systems weren't keeping pace with growing volume of information generated throughout company.
- Isolated databases did not provide the global views needed to analyze market position and make critical business decisions.
- Launched corporate initiative to build global data warehouse to centralize information and better analyze data.

Solution

- Selected InfoSphere Information Server to access, understand and cleanse data from its worldwide operations.
- Information is delivered into an enterprise data warehouse leveraged by Cognos business intelligence and performance management software for reporting.
- Replaced Microsoft SQL Server with DB2 Data Warehouse Edition.




Business Benefits

- **Simplify** company's e-commerce and electronic publishing initiatives by **reconciling, cleansing and updating** product information, part numbers, pricing data, renderings and drawings.
- **Deliver updated information** to different processes and systems that can use it to **improve sales, marketing and product development.**

IBM Delivers The Best Support For Business Optimization

- Consolidated information integration platform
- Unified metadata management
- Extensive connectivity
- Enterprise scalability
- Highly scalable enterprise data warehousing platform
- Cognos 8 Performance Management
- Supports many hardware platforms and operating systems



Data Warehouse Capabilities On System z

- Large Capacity Data Base
 - ▶ **DB2 for z/OS Value Unit Edition**
 - ▶ Parallel queries, Materialized Query Table, Star Join Enhancements
- Connect, Extract, Transform and Load
 - ▶ **IBM® InfoSphere™ Information Server For System z**
 - ▶ Runs on zLinux (DataStage also runs on z/OS)
- Analysis Tools
 - ▶ **Cognos**
 - Solution for reporting, analysis dashboards and scorecards
 - Other tools (QMF, Alphablox, Data Quant, Hyperion, Business Objects, SAS, IBI)
 - ▶ **InfoSphere Master Data Management Server for Linux on System z**
 - More effectively manage high-value operational data
- Performance Monitoring
 - ▶ **IBM Tivoli Omegamon XE for DB2 Performance Expert on z/OS**
- Security and Compliance
 - ▶ **DB2 Data Archive Expert, DB2 Test Database Generator, DB2 Audit Manager Expert, IBM Encryption for DB2 and IMS Databases**
- Dedicated Query Hardware
 - ▶ zIIPs for parallel queries & remote access
 - ▶ Superior I/O bandwidth, multiple I/O paths

IBM vs. Oracle and Microsoft – Effective Use of Data (Business Optimization)

	IBM	Microsoft	Oracle
Extensive Enterprise Connectivity	InfoSphere Information Server	Requires 3rd party	Oracle-focused
Enterprise Scalability	InfoSphere Information Server	Windows only	Waiting for Fusion
Integrated Data Quality	InfoSphere Information Server	NO	YES
Support for SOA	InfoSphere Information Server	Write code using ADO.Net	Write code to call PL/SQL as web service
Integrated, Comprehensive Performance Management Platform	IBM Cognos	NO Multiple products with different metadata models	NO Multiple products (Oracle, Hyperion, Siebel) with different metadata models
100% web-based zero footprint BI	IBM Cognos	NO Requires MS Office, BI Dev Studio, Report Builder (plug-in)	NO BI Publisher uses MS Word, Many BI functions require desktop apps and plug-ins

Where Should We Deploy This Solution?

Capability is important, but cost is a big concern for us



CIO

Deploying this solution on your System z will cost less than Oracle on distributed



IBM

Case Study: Deploy New 10TB Data Warehouse on z/OS With Disaster Recovery (Cognos Base Function)

Existing Mainframe




Existing z10:
2 GP 1,720 MIPS
DB2 and utilities
With 20Tb storage

Existing Disaster Recovery Site



Existing:
1 GP processor for hot
disaster switch-over
1 "dark" DR processor
With 20Tb storage

Add 1 LPAR for New Data Warehouse w 3.8 TB Storage




2,184 MIPS
additional
workload on z/OS
and 1840 MIPS on
zLinux

Incremental:

2 GP 1,310 MIPS (60%) DB2 & Utilities
1 zIIP 874 MIPS (40%) DB2
1 IFL 920 MIPS DataStage
1 IFL 920 MIPS Cognos
Add 10 GB memory

And add Disaster Recovery w 3.8 TB Storage




3 year
cost of
acquisition
\$5.93M

Capacity Backup:
2 GP
1 zIIP
1 IFL


Or add HP Integrity Superdome sx2k 9150N Server w 7.3 TB storage

Prod



And add Disaster Recovery W 7.3 TB storage

Prod



3 year
cost of
acquisition
\$15.16M

HP DR solution is used in
software and hardware

Storage Costs: DB2 Provides More Storage Savings Than Oracle

- DB2 for z/OS lowers TCO by reducing storage needed
 - ▶ TPC-H Benchmark: DB2 compression of 62% vs 27% for Oracle RAC
- Storage savings with DB2 vs. Oracle for a 10 TB data base

	Oracle	DB2 for z/OS*
Storage System	HP XP24000 Storage	IBM System Storage DS8100
Overall database compression ratio (using TPC-H benchmark results)	27%	62%
For 10 TB uncompressed data storage needed	7.3 TB of HP Storage	3.8 TB of IBM Storage
Cost of storage (3 year TCA)	$\$888,399 + \$37,560 \times 3$ = \$1,001,079	$\$192,205 + \$7,992 \times 2^{**}$ = \$208,189
With compression, storage for DB2 costs <u>79% less</u> than for Oracle		

*DB2 for z/OS achieves similar compression ratios to those of DB2 for LUW

**IBM storage maintenance fee for the first year is included in the warranty

Data Warehouse With Disaster Recovery Incremental Cost Breakdown

Mainframe Incremental Hardware

OTC		ANNUAL	
GP	\$2,293,200	Processor Maintenance *	
1 zIIP&2 IFL	\$375,000	(For year 2, 3)	\$175,177
DR Processors	\$60,000		
Memory (10 GB)	\$60,000		
IBM Storage (3.8TB x 2)	\$384,410	Storage Maintenance (For year 2, 3)	15,984
TOTAL	\$3,172,610	TOTAL	\$191,161 (year 2, 3)

Mainframe Incremental Software

OTC		ANNUAL	
Utilities	\$550,125	Data Stage S&S	\$51,600
Cognos	\$86,500	Cognos S&S (Yr 2,3 all above)	\$17,300
DataStage	\$258,000	Utilities S&S	\$79,041
		DB2 MLC x12	\$163,392
		z/OS MLC x12	\$87,672
		zLinux	\$36,000
WAS	\$154,050	WAS Maint	\$30,810
TOTAL	\$1,048,675	TOTAL	\$396,915 (year 1) \$465,815 (year 2, 3)

Distributed Incremental Hardware

OTC		ANNUAL	
HP Processors	\$2,652,717	Processor Maintenance	\$233,653
DR Hardware	\$1,591,630	(prepaid in year 1 for 3 years)	
HP storage (7.3TB x2)	\$1,776,798	Storage Maintenance	\$75,120
TOTAL	\$6,021,145	TOTAL	\$776,079 (year 1) \$75,120 (year 2,3)

Distributed Incremental Software

OTC		ANNUAL	
Oracle EE&Utilities	\$2,982,750	Oracle S&S	\$656,205
Unix	\$227,520	Unix S&S (prepaid in year 1 for 3 years)	\$85,504
Oracle Warehouse Builder ETL	\$557,750	Oracle Warehouse Builder ETL S&S	122,705
Oracle BI SE	\$1,115,500	Oracle BI SE S&S	\$245,410
TOTAL	\$4,883,520	TOTAL	\$1,280,832 (year 1) \$1,024,320 (year 2, 3)

Case Study: Deploy New 10TB Data Warehouse on z/OS With Disaster Recovery (Cognos Full Function)

Existing Mainframe




Existing z10:
2 GP 1,720 MIPS
DB2 and utilities
With 20TB storage

Existing Disaster Recovery Site



Existing:
1 GP processor for hot
disaster switch-over
1 "dark" DR processor
With 20TB storage

Add 1 LPAR for New Data Warehouse w 3.8 TB Storage




2,184 MIPS
additional
workload on z/OS
and 1840 MIPS on
zLinux

Incremental:

2 GP 1,310 MIPS (60%) DB2 & Utilities
1 zIIP 874 MIPS (40%) DB2
1 IFL 920 MIPS DataStage and QualityStage
1 IFL 920 MIPS COGNOS
Add 10 GB memory
CDC on Source servers running SQL Server

And add Disaster Recovery w 3.8 TB Storage




3 year
cost of
acquisition
\$8.85M

Capacity Backup:
2 GP
1 zIIP
1 IFL

Or add HP Integrity Superdome sx2k 9150N Server w 7.3 TB storage


Prod



48 Chip
96 Core

And add Disaster Recovery W 7.3 TB storage

Prod



48 Chip
96 Core

3 year
cost of
acquisition
\$20.83M

HP DR solution is used in
software and hardware

US Retailer Improves Response Time By Co-locating Data Warehouse And Operational Data

- A major US retailer moved their 5.5TB data warehouse from distributed servers to System z
 - ▶ Operational data bases were already located on System z
- On average they reduced query processing times by 80% due to better query parallelism in DB2 for z/OS
 - ▶ (17 minutes to 3 minutes)
- They saved CPU cycles to load the data warehouse
 - ▶ Avoided network processing

Deliver Business Insight With A Data Warehouse

We learned the best promotion to maximize our business profits

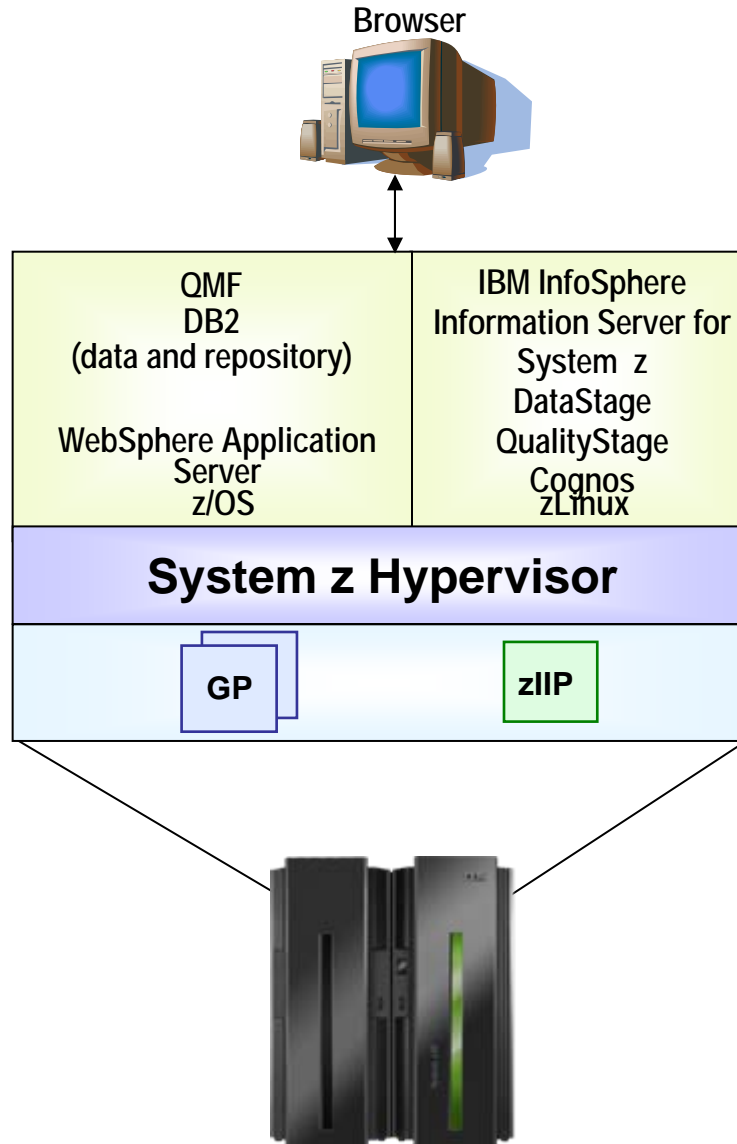


Marketing VP

And I saved money by deploying our data warehouse solutions on System z



CIO



Best

- ▶ Quality of service
- ▶ Co-location
- ▶ Lower cost of acquisition



IBM

Accelerate Development With Pre-built Data Models For Industry Data Warehouses

- Quick start your data warehouse design with pre-built **IBM Industry Data Models**
- Use **Enterprise Model Extender** to customize the data model
 - ▶ Eclipse plug-in for Rational Data Architect
- Models capture the best practices of over 400 IBM customers
 - ▶ Banking, finance, health, insurance, retail and Telco



