



Software Group – Information Management

Making data warehousing and business intelligence drive better decisions

Drew Moore, DW/BI technology specialist, LRS IT Solutions, a Cognos technology and IBM partner

Claudia Imhoff, President of Intelligent Solution

Mike Biere, Worldwide Specialist for Business Intelligence Analytics Tools, IBM Corporation



Claudia Imhoff

President and Founder Intelligent Solutions, Inc.

A thought leader, visionary, and practitioner in the rapidly growing fields of business intelligence and customer focused-strategy – Claudia Imhoff, Ph.D., is an internationally recognized expert on analytical CRM, business intelligence, and the infrastructure to support these initiatives – the Corporate Information Factory (CIF). Dr. Imhoff has co-authored five highly-regarded and popular books on these subjects and writes monthly columns (totaling more than 100) for technical and business magazines.



Email: cimhoff@intelsols.com
Phone: 303-444-6650





Presentation Topics

- ➔ ■ Today's BI Environments
 - Distributed versus Consolidated Architectures
 - Conclusions





Trends in Today's BI

- Significant increase in number of users
- Significant growth in volume of data
- Significant pressure for faster performance
- Significant availability requirements





Today's BI Environment – Users

- Numbers of users increase significantly
 - Traditional BI rarely supported a few hundred, maybe a thousand or so users until operational BI entered
 - Opening BI up to operational personnel means ramping up into tens of thousands of users
 - These users have very different interface requirements
 - Means BI implementers must rethink how BI is delivered to business users
 - Means tighter and faster connectivity of enterprise decision support environment to rest of the company
 - These users can be more naïve, less technically savvy, less analytical in their thinking...





Today's BI Environment – Volumes of Data

- Volumes of data increase substantially
 - Detailed intraday snapshots of data are trickle-fed into data warehouses or operational data stores
 - Tens of terabytes to hundreds of terabytes are not unusual storage requirements for operational BI
 - Scalability now a mandatory requirement in any BI technology
 - Whether in processing and integration of data, storage of massive volumes, or retrieval of query responses
 - Granularity is significantly different
 - Operational BI requires much more detailed data to support specific operational functions





Today's BI Environment– Performance

- Faster performance
 - Operational BI query performance must mimic or emulate response times in operational systems
 - Sub-second to just a few seconds to return data from a query.
 - Ability to prioritize queries not only according to their importance but also their response requirements is mandatory success criterion
 - This has stumped many BI implementers and BI vendors
 - Must have ability to handle mixed work load gracefully and simultaneously – short operational BI queries & larger, more intense strategic ones





Today's BI Environment– Availability

- Today's data warehouse directly supports transactional business processes (operational BI)
 - Must adhere to stringent service levels
 - 24 X 7
 - 365
- An outage of BI will have a direct impact on enterprise's ability to do business!





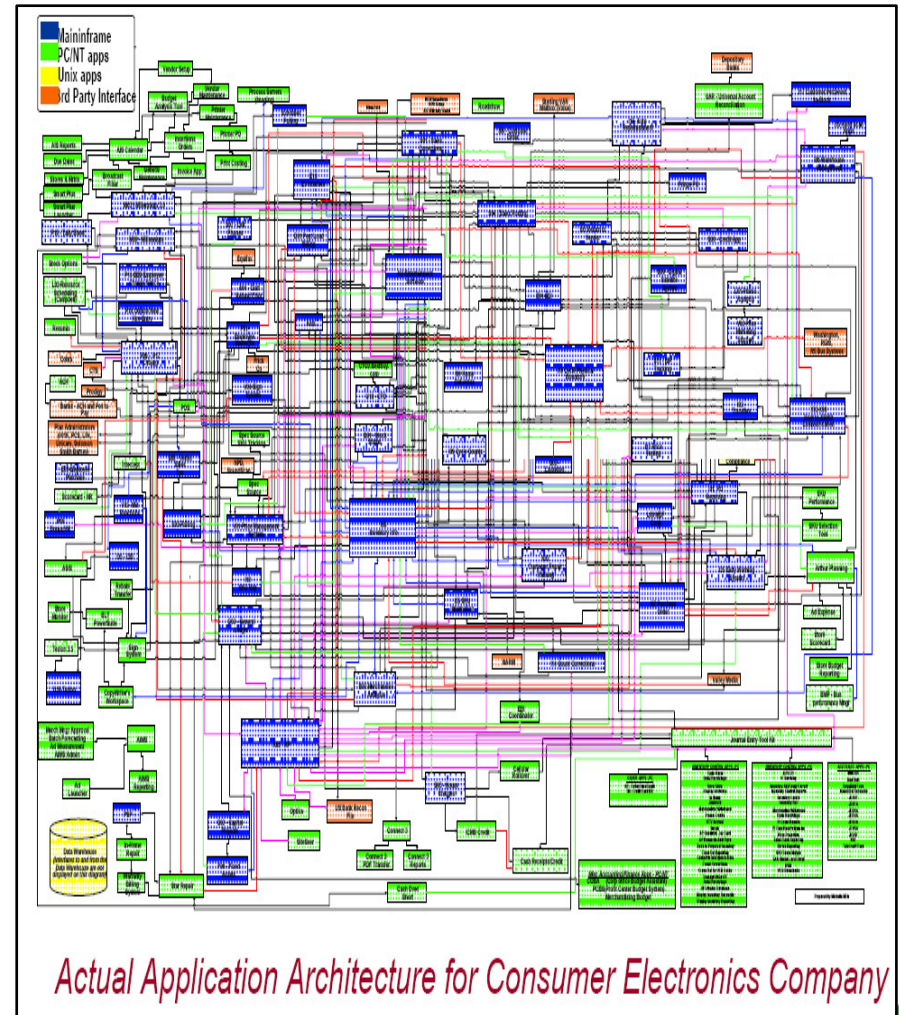
It's Not Getting Any Easier

- Number and disparity of information stores
- Information quality, accuracy, and consistency
- Timeliness of information
- Complexity of current decision-making systems
- Pressure to reduce costs of BI



Disparity of Information

- Operational business data
 - Maintained by legacy applications, application packages, web systems
 - Stored in transaction, event, master data stores
- Historical business information
 - Maintained by data integration software and BI applications
 - Stored in a data warehouse, data marts, data cubes
- Business content
 - Maintained by collaborative, content, operational, web systems
 - Stored in databases, text and rich media files, web pages, etc.



Source: IBM






Complexity of Current Decision-Making Systems

- “Pockets” of BI are isolated throughout the enterprise
 - Everyone has their own BI server...
- Distributed technology causing significant problems
- BI deployment still requires significant IT involvement



Biggest TCO Expense These Days Are People!*



* Based on IBM Scorpion customer analyses





But Are We Getting Better Decisions?*

6 in 10

Managers miss information that may be of value to them because they cannot find it

1 in 2

Executives 55% said their decisions rely on qualitative and subjective factors

3 in 4

Business leaders say more predictive information would drive better decisions

1 in 3

Managers frequently make critical decisions without the information they need

19+ Hours

Spent by knowledge workers each week searching for and understanding information

3 in 4

Executives are working to increase business analytics usage



* Various industry analyst surveys in 2008-2009

Copyright © 2010, Intelligent Solutions, Inc., All Rights Reserved






Presentation Topics

- Today's BI Environments
- ➔ ■ Distributed versus Consolidated Architectures
- Conclusions




Multiple Server Distribution vs Consolidated BI

Distributed Over Multiple Servers




Versus

Consolidated (Mainframe)




= Virtual Data Mart (views, etc.)



= Physical Data Mart




Comparison



- Distributed over multiple servers
- An easy sell but ...
- How consistent are extracts, transformations, integrations and summarizations?
- Redundancy of data magnified
- Security concerns over physically moving data
- Analytic results between groups are inconsistent
- Generally must create physical data marts

- Consolidated
- Is a sound foundation upon which to build future DSS needs
- Data acquisition and delivery is consistent and documented
- Maintains historical integrity
- Provides full data lineage meta data
- Stable, consistent data values and calculations
- Can create physical or (less costly) virtual data marts





Distributed Costs

- Cost of running additional workload on distributed servers goes up nearly linearly
 - Labor is highest cost element in distributed environments
 - Administrative staff costs increase in proportion to number of servers
 - New workload requires additional servers
 - Cost of additional servers / software licenses are nearly linear
 - Electrical and air conditioning costs also increasing
- Need to rethink scale out strategies in terms of cost per unit





Mainframe Costs

- Cost of running additional workload on mainframe generally goes down as total workload grows
 - Labor costs remain roughly same as workload grows
 - Mainframe pricing policies designed to favor additional workload
 - Lower software costs per transaction as workload grows
 - Lower electrical and air conditioning consumption than server farms
- Mainframes running high workloads are very cost efficient platform





Costs of Data Storage and Management

- Total Cost of Storage is about three times more in distributed environments
 - Application specific data silos tend to over-allocate
 - Storage utilization of 25-30% or less is typical in distributed environments
 - Mainframe fine grained allocation and data sharing yield typical storage utilizations of 80% +
- Environment management heartburn
 - Disaster recovery of separated data silos – not easy
 - Synchronization, and transfer requirements – also not easy and entail security risks





Presentation Topics

- Today's BI Environments
- Distributed versus Consolidated Architectures
- ➔ ■ Conclusions





Conclusion

- Reducing cost and complexity in BI infrastructure is mandatory today
 - Consolidation efforts to System z expand its usage from transactional processing to DW/BI
- Leverage your investment in System z:
 - Mainframe still THE standard for mission-critical apps – and BI now is mission-critical
 - New technology and price points make System z a cost-effective option (Cognos and SPSS for BI)
- Data integration (ETL, quality processing, etc.) easier if BI co-locates on same platform as most operational data
 - Reduce data latency, redundancy and complexity





Conclusion

- Maximize utilization of mainframe
 - They can be used at 90+% utilization, 24 hours by 7 days
 - Reverse trend and consolidate BI workload onto System z – think about virtual data marts
- Minimize other costs
 - Minimize software tool costs
 - Minimize power costs
 - Minimize outages and security breaches
 - Minimize labor costs



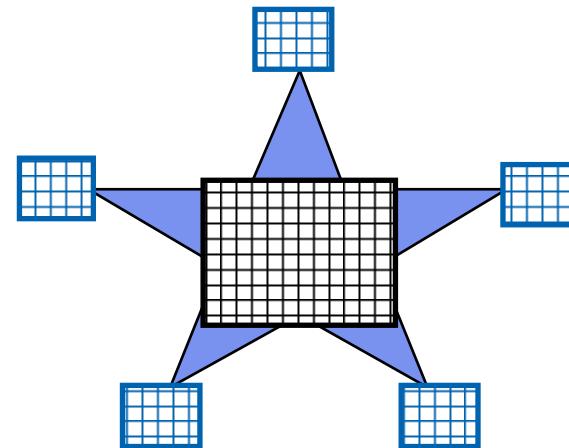


Software Group – Information Management

Data Warehousing and BI for System z

Mike Biere

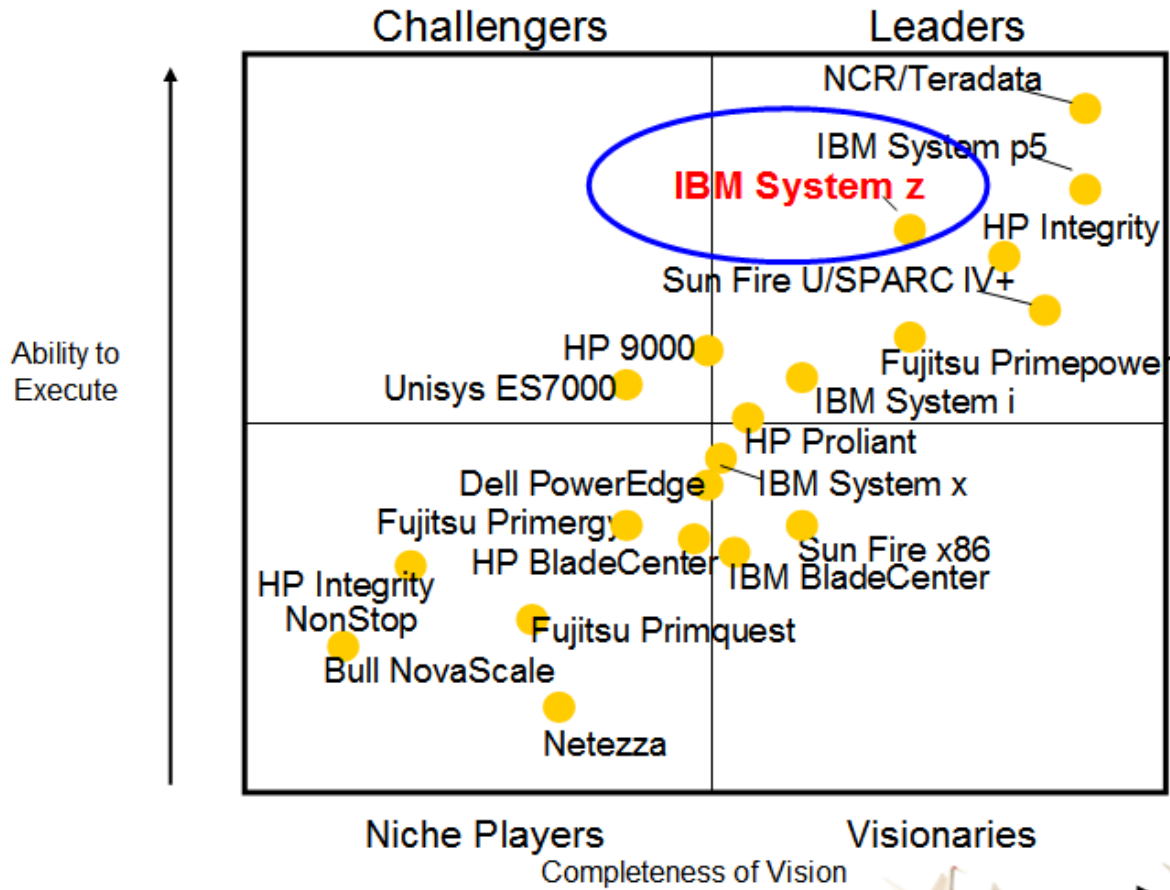
mbiere@us.ibm.com





Gartner's Data Warehouse DBMS Server Magic Quadrant

August 2006

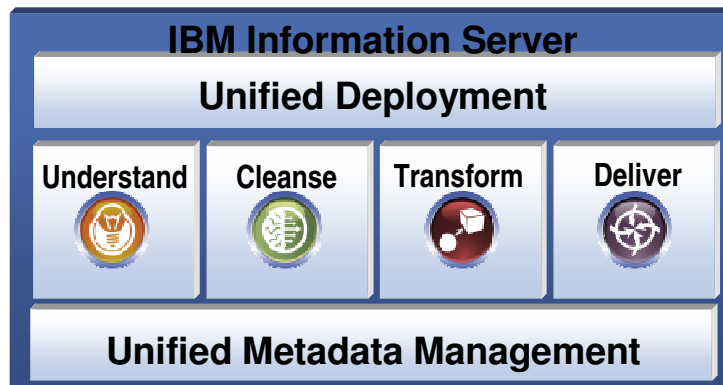


Mission-Critical Data Warehousing

InfoSphere Information Server for System z

Accelerating the delivery of trusted information


**Profile, cleanse, and transform information
from heterogeneous data sources
to drive greater business insight**




- Significant cost savings on System z
- Scalable to any volume and processing requirements
- Fully integrated, auditable data quality
- Metadata-driven integration for increased productivity

Physical Metadata: IBM WebSphere Information Analyzer


- Data-centric analysis of application, database and file-based sources
- Secure, detailed profiling of fields, across fields, and across sources
- Creation of metadata from profiling results
- Results instantly promotable across IBM Information Server



Subject Matter Experts



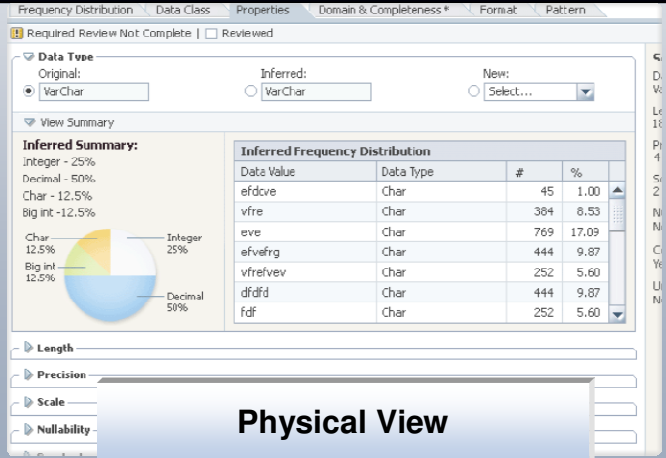
Data Analysts



Understand

IBM WebSphere Information Analyzer

Analyze source data structures, and monitor adherence to integration and quality rules



Physical View

Business Metadata: IBM WebSphere Business Glossary

- Web-based authoring, managing & sharing of business metadata
- Aligns the efforts of IT with the goals of the business
- Provides business context to information technology assets
- Establishes responsibility and accountability

Database = DB2

Schema = NAACCT

Table = DLYTRANS

Column = ACCT_NO

data type = char(11)



Technical



Business

GL Account Number

The ten digit account number. Sometimes referred to as the account ID. This value is of the form L-FIIIIVVVV.

Subject Matter Experts

Business Users

Understand

IBM WebSphere Business Glossary

Create and manage business vocabulary and relationships, while linking to physical sources

Business View

Logical Metadata: Rational Data Architect

- Data modeling for data structures and federations
- Federated data discovery
- Metadata relationship discovery & mapping
- Impact analysis, and synchronization across models
- SQL & XML generation capabilities



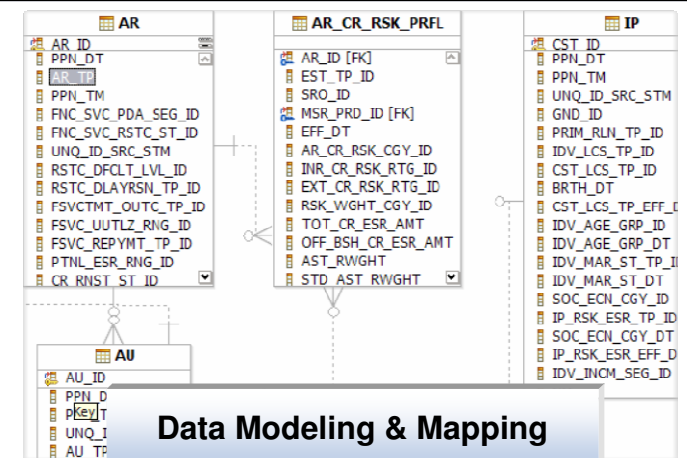
Subject Matter Experts



Architects


Rational Data Architect

Create and manage business vocabulary and relationships, while linking to physical sources




Data Cleansing: IBM WebSphere QualityStage

- Specialized data quality functions seamlessly integrated with DataStage
- Visual tools for defining complex matching and survivorship logic
- Ensures clean, standardized, de-duplicated information
- Enables a single version of the truth




Subject Matter Experts



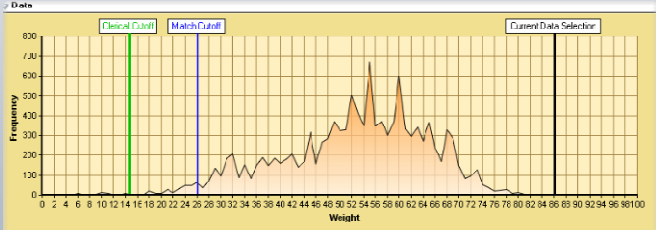
Data Analysts

Cleanse



IBM WebSphere QualityStage™

Standardize and correct source data fields, and match records together across sources to create a single view



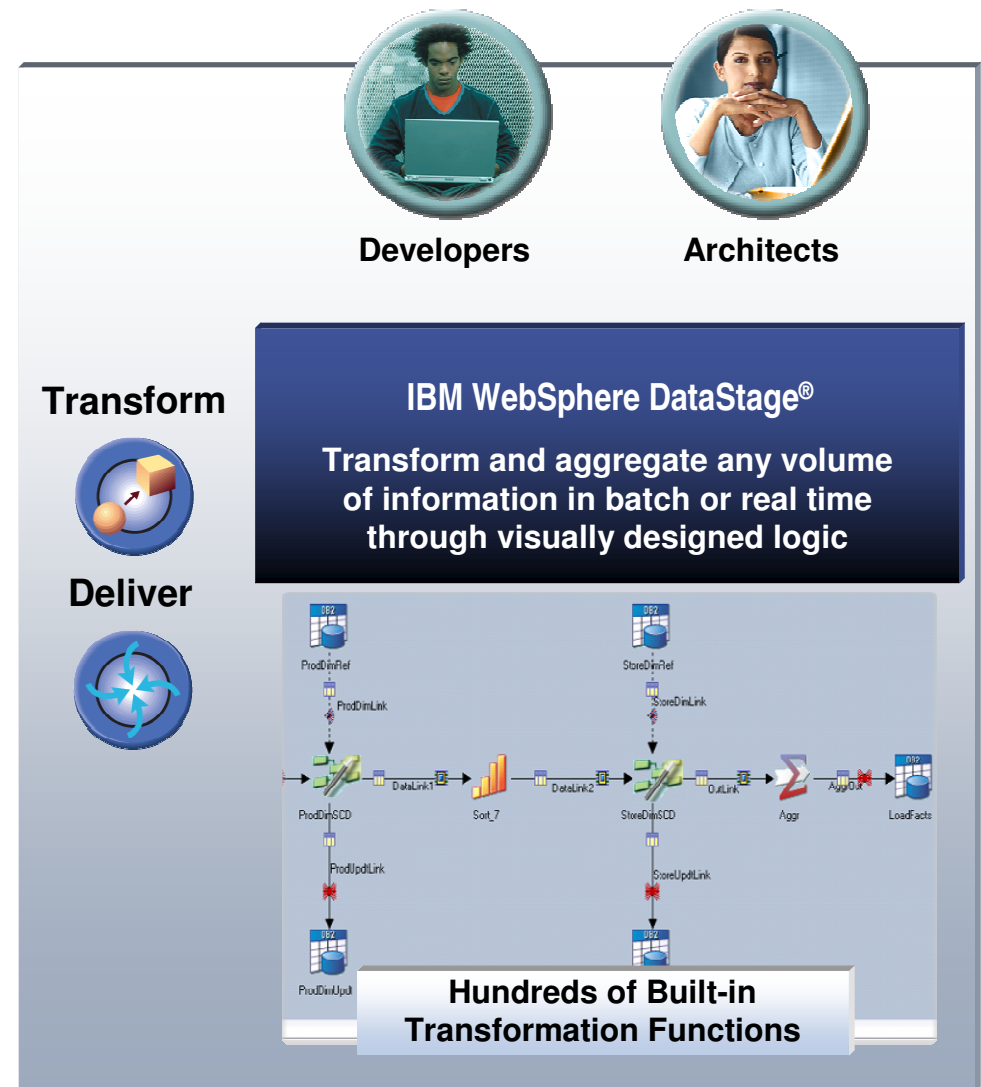
idID	RecordType	PassNumber	Weight	GenderCode	FirstName	MiddleName
4111	XA	1	85.57	IM	OTIS	GARLAND
4111	DA	1	85.57	IM	OTIS	GARLAND
5311	XA	1	84.08	IM	NICHOLAS	T
5311	DA	1	84.08	IM	NICHOLAS	T
3420	XA	1	81.68	IM	CLAUD	LAJALE
3420	DA	1	81.68	IM	CLAUD	LAJALE
3420	XA	1	81.68	IM	CLAUD	LAJALE
3420	DA	1	81.68	IM	CLAUD	LAJALE
4826	XA		80.20	IM	CLIFF	H
4826	DA		80.20	IM	CLIFF	H
7300						CHIN
7328						CHIN
7292						GILLIS
2282						GILLIS
6970						MCHELLE
6970						MCHELLE

Visual Match Rule Design



Data Transformation & Movement: IBM WebSphere DataStage

- Codeless visual design of data flows with hundreds of built-in transformation functions
- Optimized reuse of data integration objects
- Leverages parallel processing without requiring design changes
- Capable of supporting batch and real-time operations



Data Federation: IBM WebSphere Federation Server Server and Classic Federation

- Access diverse & distributed information as if it were in one system
- Industry leading query optimization with single sign-on, unified views, and function compensation
- Transactional write capabilities across heterogeneous sources
- Visual tools for federated data discovery & data modeling

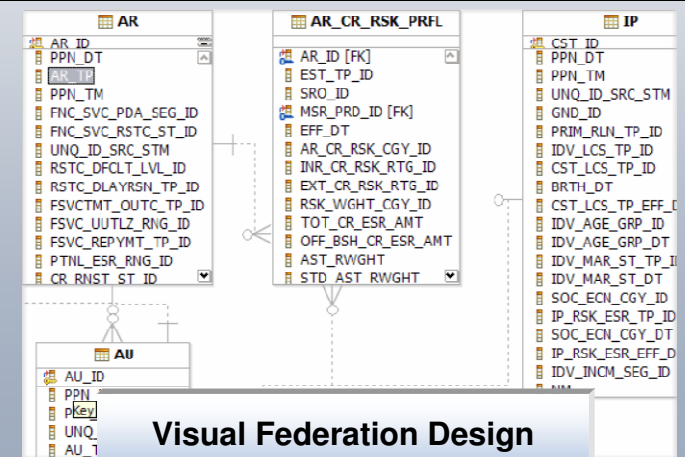
Deliver



IBM WebSphere Federation Server

Access and integrate heterogeneous information across multiple sources as if they were a single source

Extend value of existing analytical applications by providing real-time access to integrated information



Visual Federation Design



The IBM Information Server Advantage

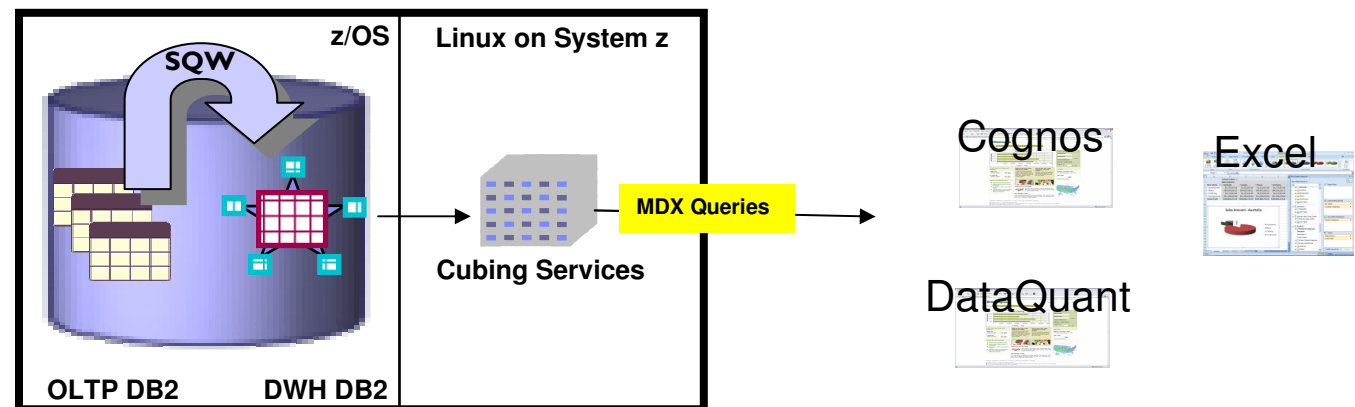
A Complete Information Infrastructure

- A ***comprehensive, unified foundation*** for enterprise information architectures, scalable to any volume and processing requirement
- ***Auditable data quality*** as a foundation for trusted information across the enterprise
- ***Metadata-driven integration***, providing breakthrough productivity and flexibility for integrating and enriching information
- ***Consistent, reusable information services***—along with application services and process services, an enterprise essential
- Accelerated time to value with ***proven, industry-aligned solutions*** and expertise
- ***Broadest and deepest connectivity*** to information across diverse sources: structured, unstructured, mainframe, and applications

InfoSphere Warehouse on System z

Adds core data warehouse and analytics capability to DB2 for z/OS

- Advanced physical database modeling and design
- In-database data movement and manipulation capabilities of SQL Warehouse Tool (SQW)
- Optimize multidimensional reporting and analysis of data with Cubing Services



[System z Environment Enhanced with InfoSphere Warehouse](#)

IBM Smart Analytics Optimizer

Technology Preview for System z

What is it?

- ✓ A high performance extension that easily integrates with IBM data systems, delivering predictable, order-of-magnitude faster, analytic query response times, while lowering operating costs



How is it different

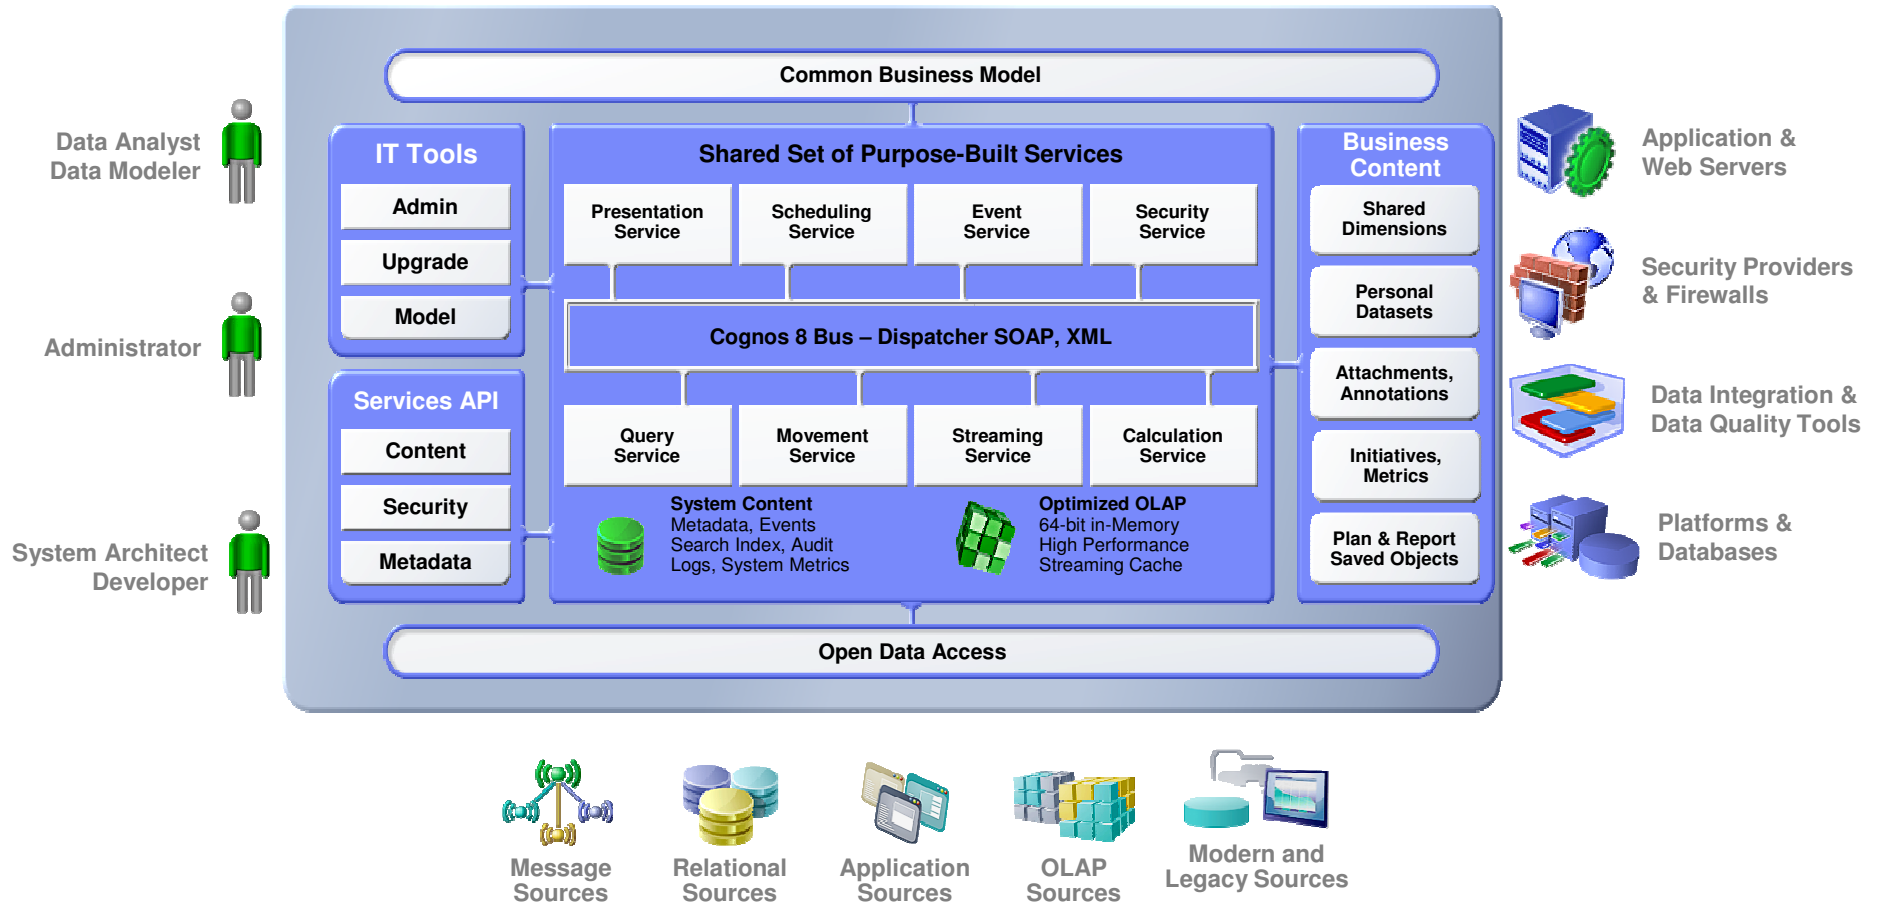
- ✓ **Deep integration with IBM data management systems**
- ✓ **High performance query software, based on advanced data in-memory technologies**
- ✓ **Leveraging existing data system investment and values without any changes to applications**
- ✓ **For System z, extends gold-standard manageability, security, and availability to high-performance analytic applications**

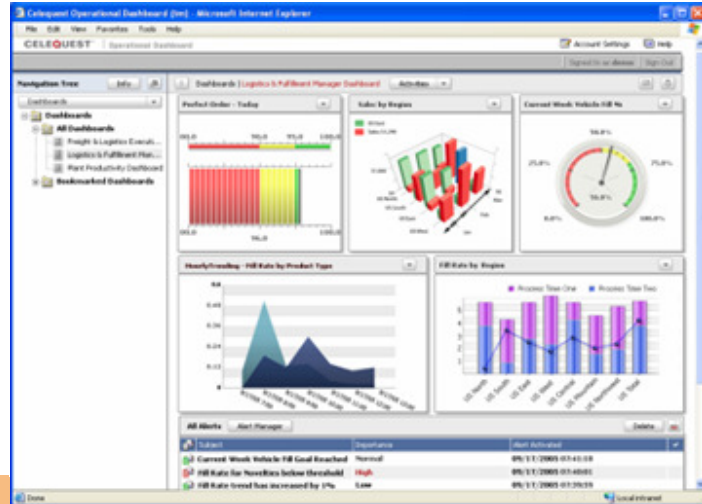


Next ... deliver a world-class BI platform

Built on an SOA to deploy successfully within changing environments

Cognos.
software





■ Dashboards

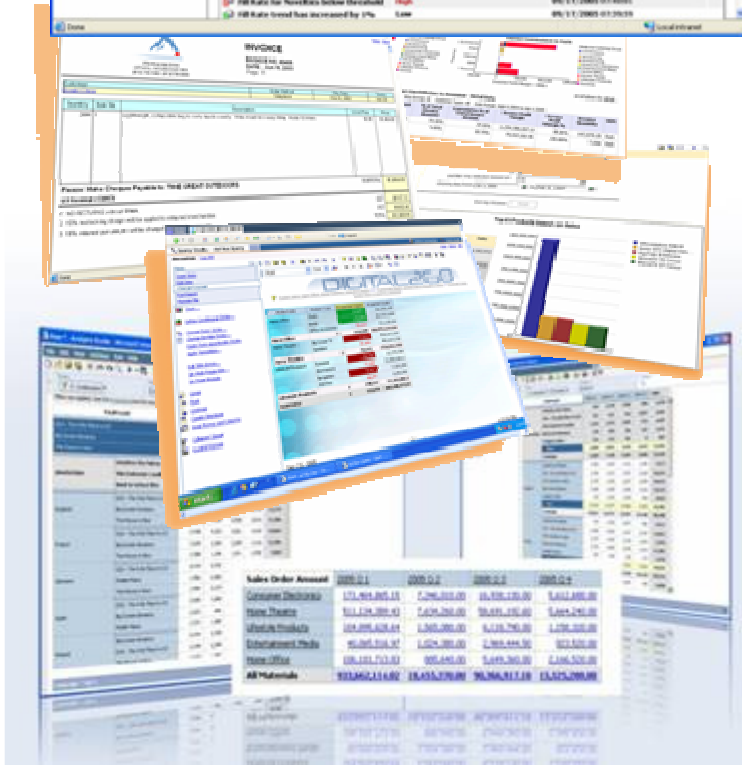
- Translate complex information into high-impact presentations
- Allow you to spot changes
- Are highly intuitive
- Align decision makers

■ Enterprise Reporting

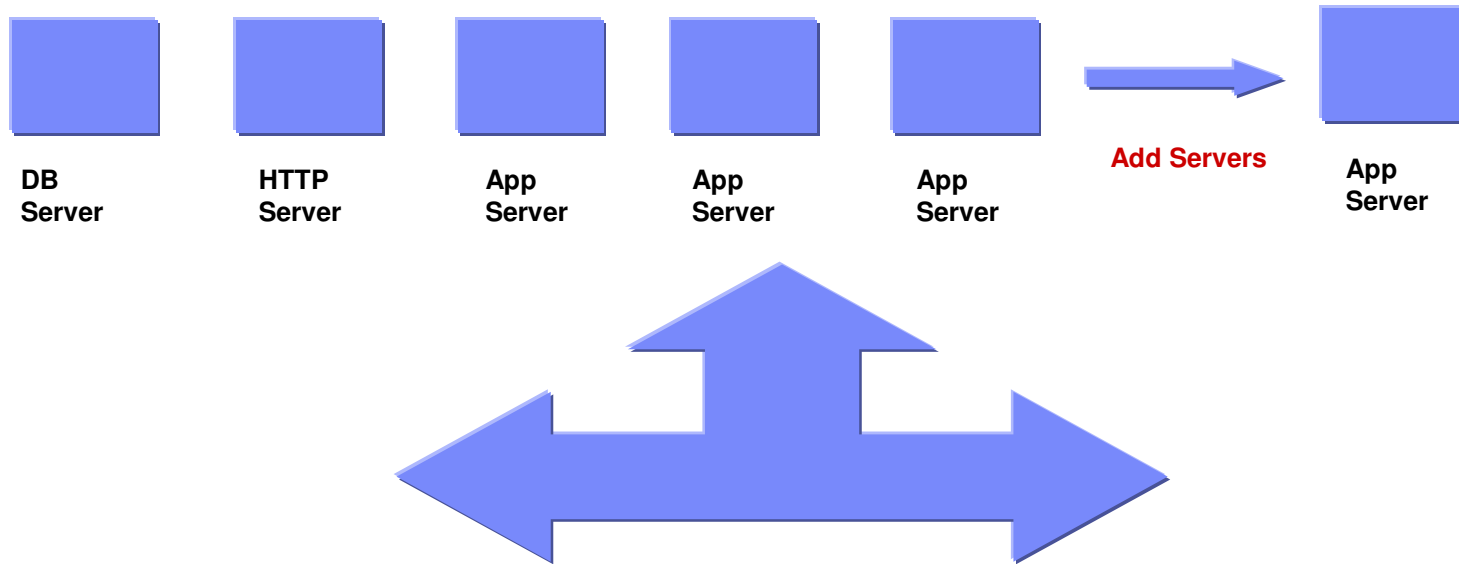
- Supports multiple report types: Production, Managed, Ad-hoc, Financial, etc
- Is adaptable to any data source
- Operates from a single metadata layer
- Can be personalized and targeted
- Can be distributed via email, portal, MS-Office, search application and mobile device

■ Analysis


- Enables the guided exploration of information that pertains to all dimensions of your business
- Performs complex analysis and scenario modeling easily and quickly
- Gets to the “why” behind an event or action to improve business performance.
- Moves from summary level to detail levels of information effortlessly



BI Server Consolidations




In a distributed environment growth is achieved by growing the server farm up or out or both





Processors and % utilization typically 20% (for actual work)

Number of virtual machines = 1 to Infinity!*




- Each virtual machine has RAM, Disk, NICs, processors, and HBAs (FCs) assigned to it
- Each copy of Linux sees an entire system z Server with the virtual machine's resources
- Intercommunication (TCP/IP) is facilitated by the hypervisor

LPAR = logical partition of the system hardware



z10 machines can have up to 60 LPARs (depending upon the processor class)

* Customers regularly run z10s at or near 100% utilization (that's efficiency!!!)



IBM FMS – currently running 40,000 users of Cognos 8 on System z Proof of Success with User Requirements

Replaced previous system in 5 months

- Exorbitant ISV charges erased
- On-demand reporting model
- WW deployment with substantial cost savings

Simplified User Experience

- Single, unified web portal for all their FMS reporting needs supporting multiple browsers
- Reduced number of reports (from 14 to 4) providing the same level of information
- Data populated on existing reports dramatically decreased due to drill down capabilities
- Significant improvement in reporting performance and response time
- Users now quickly and easily define what information they view and how they access it

Increased User Adoption

- Accommodated a larger user population as a result of System z strengths and capabilities
- Ran approx. 350,000 reports in the 1st 5 months, validating fast and broad user adoption

Delivered Increased analysis value to the Business

- New information for Territory Analysis - assist managers in analyzing a seller's territory coverage before achievement and commission payments are available

*Yes we are now
drinking our own
Kool-Aid!!*





Numius Case Study

Cognos.
software

- **Numius tested an existing customer's distributed Cognos environment on System z**
 - The application was successfully and without loss of functionality ported to the System z platform. This required no redevelopment.
 - The client's application would not require a redesign to accommodate its growth in data volumes or in terms of users.
 - Reports that are not practically useable at client's site now become relevant again. Reports that did not run at client's site now are runnable.
 - Client would be able to serve many multiples of current number of users with the very simple architecture from this PoC.
 - Client could scale out to more complex architecture without increased hardware complexity.
 - ***Throughput (not clock speed) 400x that of distributed***
 - Much of the improvement was a result of the processing synergy between Cognos 8 BI on System z and DB2 for zOS



50TB Summary – Operational BI validation


- System z and Cognos BI can respond to operational BI requirements
 - Successfully ran **400 active users** simulating call center agents accessing a prompted operational BI report
 - Average **1.75 seconds** response time for query and report creation per user over a 15 min run (steady state), at **56% Linux CPU** utilization
 - DB2 for z/OS provides **very efficient access** to operational BI data
- Cognos configuration options for Linux on System z
 - Multiple 31Bit WebSphere Application Servers on a single system
 - Varied resources assigned to Linux on System z and Cognos
- Load testing techniques using Rational Performance Tester
 - Strategic IBM tool for performance/load tests also recommended for customer tests
- Collateral
 - Best practices and results in Redbook
 - Collected detailed performance measurement data

Cognos.
software

What: Smart Analytics Cloud

A private cloud optimized for analytic services in large enterprises

	Defined as ...	To create...	That delivers ...
Smart Analytics Cloud	IBM Smart Business - services with industry leading hardware & software	A private cloud computing solution for business intelligence (BI) & analytics	A services solution for delivering business intelligence to the entire organization



IBM Services

- Create awareness of BI and understand the needs for a BI strategy across the organization
- Complete a readiness assessment to define the scope and priorities for the solution
- Deploy Cognos 8 BI for Linux on System z as a private cloud
- Provide the skills for the on going management & expansion of their BI private cloud deployment





IBM Cognos Now! – Real Time Monitoring

For critical, intra-day monitoring of operational KPIs and metrics

- Aggregated across multiple transactional systems and data sources
- No BPM system required

Closed loop business optimization

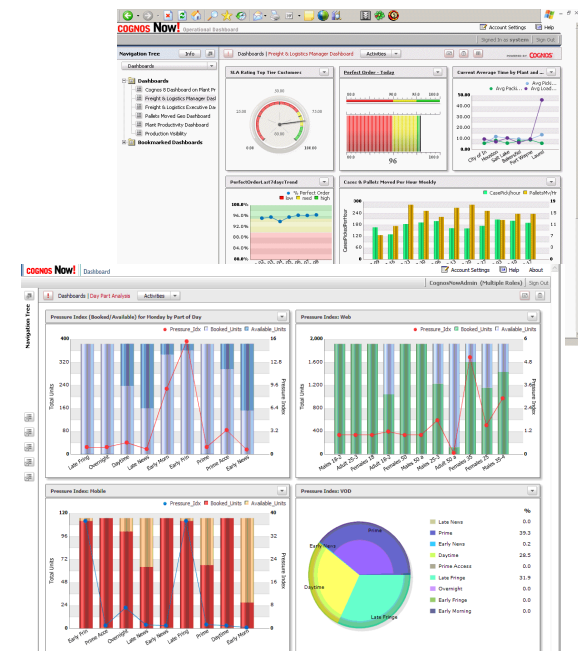
- Complete loop from monitor to alerting to corrective action
- Identify, customize operational KPIs and metrics
- Understand and perform root cause analysis
- Drive rapid, effective decision-making and action

Autonomy for line of business user

- Self service model
- User-defined thresholds, alerts
- Graphical watch points
- Customization by end users

Cost effective, low risk and rapid deployment

- No roles based pricing, unlimited user pricing in Americas
- Prepackaged hardware, software or VM appliance






Summary

- We have enhanced the System z portfolio significantly to accommodate requests for Data Warehousing and Business Intelligence
- IBM's major offerings on System z are
 - Information Server
 - InfoSphere Warehouse
 - Smart Analytics Optimizer
 - Cognos 8 BI for Linux on System z
- Data Warehousing requirements are changing to a more server-centric, near real-time mode of operation
- IBM now provides the full range of DW and BI solution elements on System z in a cost-effective and open architecture

LRS[®]  solutions
Experience at your Service



Corporate Overview

- ◆ **Incorporated**
in 1979; privately owned
- ◆ **Focused**
on software, IT Solutions and value-added services
- ◆ **Consistently Ranked**
among the top 200 software companies in the world
- ◆ **Highest “Financial Strength”**
rating from Dun & Bradstreet
- ◆ **Debt-Free**

Key Business Units



**Enterprise Output
Management**

LRS IT Solutions

LRS Consulting Services

Pension Gold

LRS Sports

LRS Education Services

LRS Web Solutions

Proven Track Record



- ◆ **LRS Software**
running on more than 5000 systems worldwide.
- ◆ **Market Success**
our customers:
 - 72% of the Fortune 100
 - 52% of the Fortune 500
 - 66% of the Fortune Global 100
 - 45% of the Fortune Global 500
- ◆ **Trusted by many**
of the largest corporations in the world.
- ◆ **IBM Premier Business Partner**
for the 12th straight year
- ◆ **Dynamic Infrastructure**
deep skills on all IBM platforms
- ◆ **System z Authorized**
one of only 12

Why LRS? – Qualifications

IBM Premier Business Partner – 12 straight years

IBM's highest level, certifications, customer satisfaction, joint planning

IBM Business Continuity Specialty

IBM Business Partners who **demonstrate proficiency** in delivering business continuity solutions.

IBM Dynamic Infrastructure Specialty

Recognizes Business Partners who invest in and sell the full IBM portfolio to address clients' IT transformation needs.

LRS was one of the first 10 D.I. Specialty Partners World Wide



Why LRS? – Qualifications

we have the skills and experience

- ◆ **National Reach**

LRS has customers from Coast to Coast

- ◆ **Industry Experience**

Manufacturing, Distribution, Finance, Insurance, Utilities, Government, Healthcare

- ◆ **Enterprise Customers**

Since 1996, our largest client is one of IBM's largest clients worldwide

- ◆ **Highly Skilled Team**

O/S Experience – zOS, zVM, zVSE, zLinux, AIX, OS/400, Linux, Windows

- ◆ **Defined Focus**

LRS has consistently focused on providing value around infrastructure solutions.

- ◆ **VALUE Partner**

LRS is not a VOLUME Partner. We are focused on the relationship, not the transaction.

Technology Solutions with a human touch

LRS  solutions

LRS – Information Management

Is this how you want to make major decisions?



◆ LRS BI Solutions

- Data Analysis
- Concept & Design
- Training
- Implementation & Project Mgt.
- Ongoing Support & Direction
- Cognos on System z – Linux IFL

Technology Solutions...

with a human touch

LRS IT Solutions Offerings

- ◆ **Server Hardware**
System x, BladeCenter, POWER Systems, System z
- ◆ **Storage Solutions**
XIV, Enterprise, SAN, NAS, Virtualization, Tape, VTS
- ◆ **Software Solutions**
Tivoli Storage Manager / Tivoli Productivity Center
Information Management / Business Intelligence
- ◆ **Financing Alternatives**
Flexible options from IBM Global Finance & LRS





Thank you.