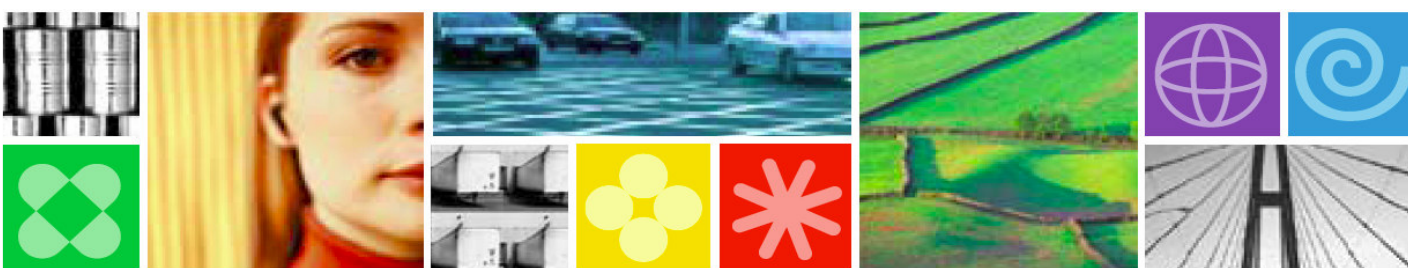




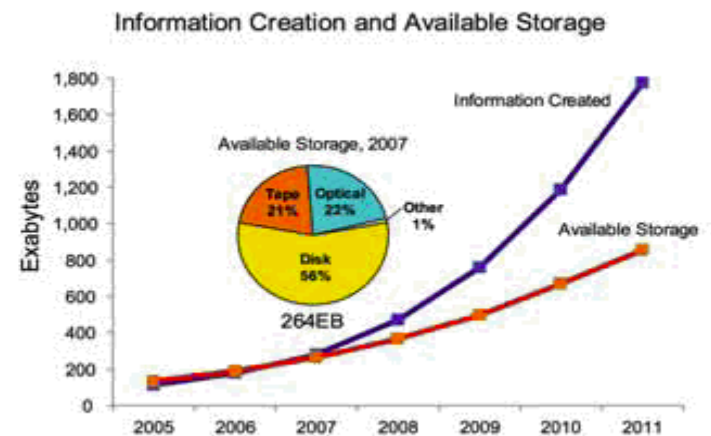
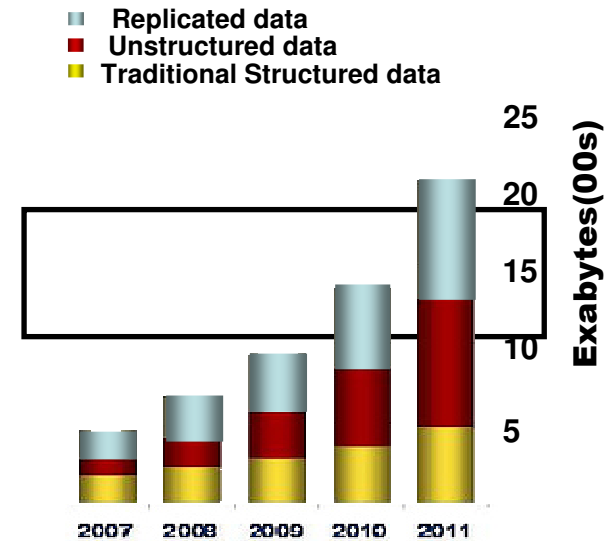
Exploiting All The Capabilities Of Your Mainframe

Handling The Information Explosion To Make Smarter Decisions



Data Volumes Are Exploding ...

- Information is doubling every 18 months
 - ▶ Structured data growing at 32%
 - ▶ Unstructured data growing at 63%
 - ▶ Replicated data growing at 49%
- IDC predicts by 2011, digital data will be ten times its size in 2006
- We now create more data than we can store
 - ▶ By 2011, half of the data created will not have a permanent home



Sources: IDC worldwide enterprise disk in Exabytes from "Changing Enterprise Data Profile", December 2007 and "The Diverse and Exploding Digital Universe", March 2008

Too Much Data!

Business growth and new regulations on data retention have our data growing out of control!



**Service Oriented Finance
CIO**

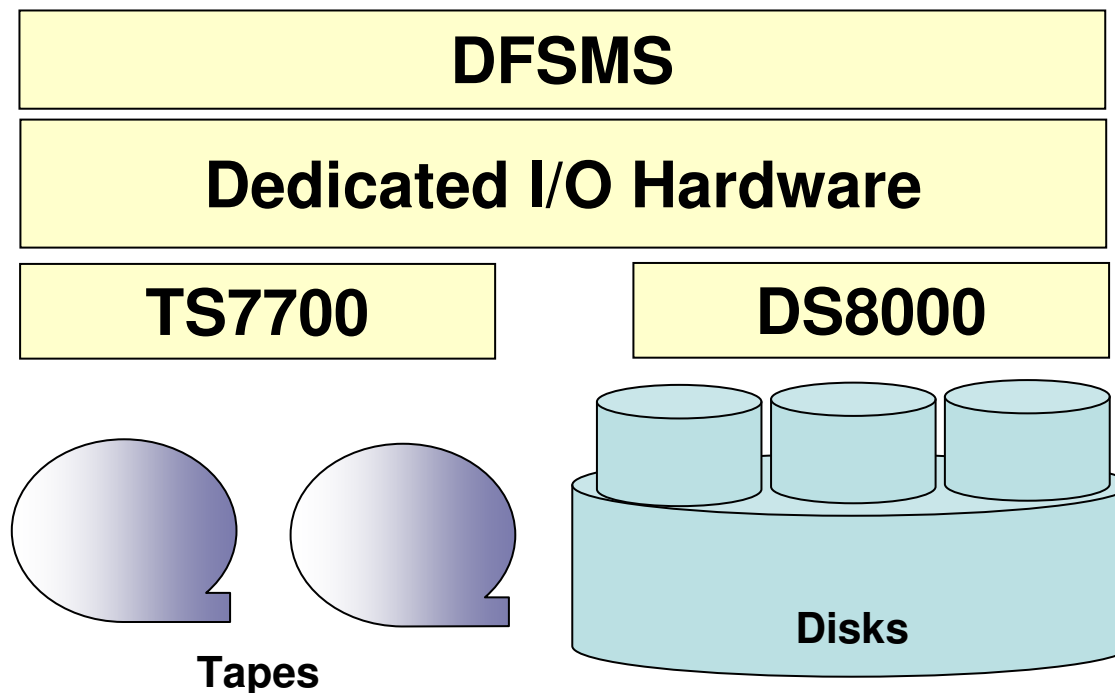
Building a scalable, cost effective storage environment is the first step



IBM

System z Storage Management Is Designed To Handle Massive Amounts Of Data

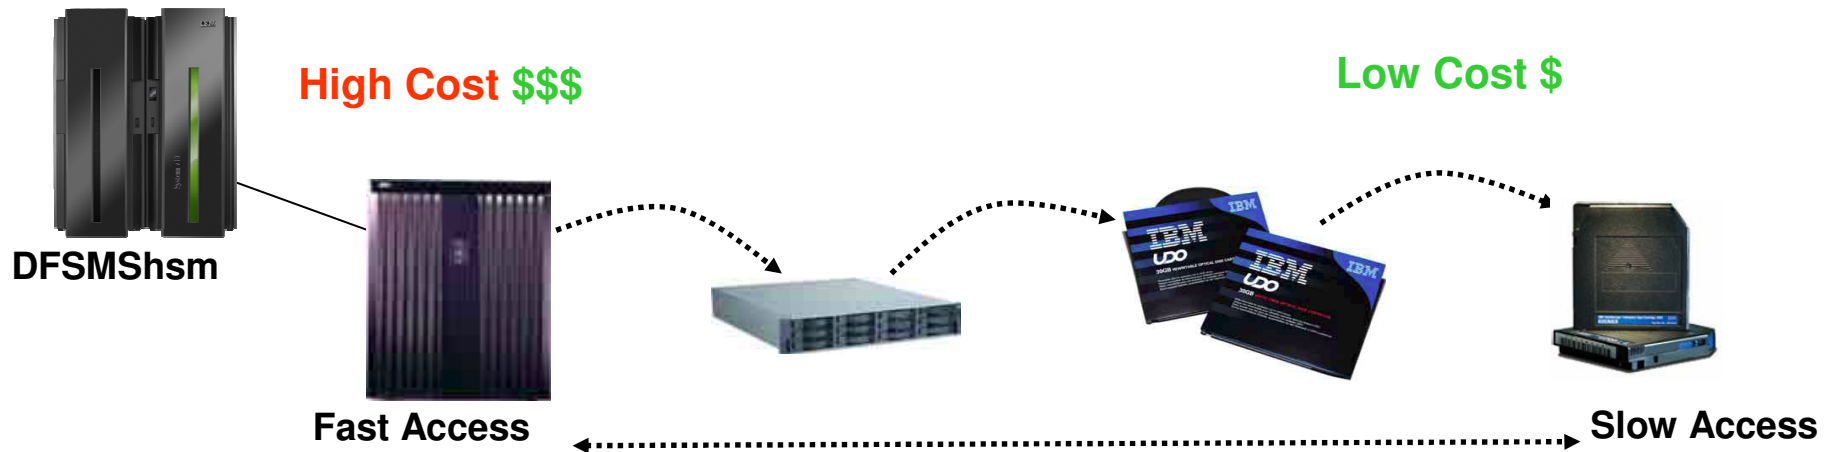
- System z Data Facility Storage Management Subsystem (DFSMS)
- System z dedicated I/O hardware offloads I/O processing cycles
- IBM System Storage DS8000 and IBM Virtualization Engine TS7700 virtualize storage and deliver massive capacity
- System z integrates these capabilities to deliver optimized storage



System z DFSMS Storage Management

- Provides System z file system and access methods
 - ▶ E.g. BSAM, QSAM, VSAM, z/OS Unix file system ...
 - ▶ Extendable while running
- Storage management features
 - ▶ Automate management of datasets, catalogs, objects, z/OS UNIX files and logical volumes
 - ▶ Move, copy, backup, recovery and automatic space management
 - ▶ Manage removable media
 - ▶ Manage movement of data in storage hierarchy
 - ▶ Concurrent access of VSAM data

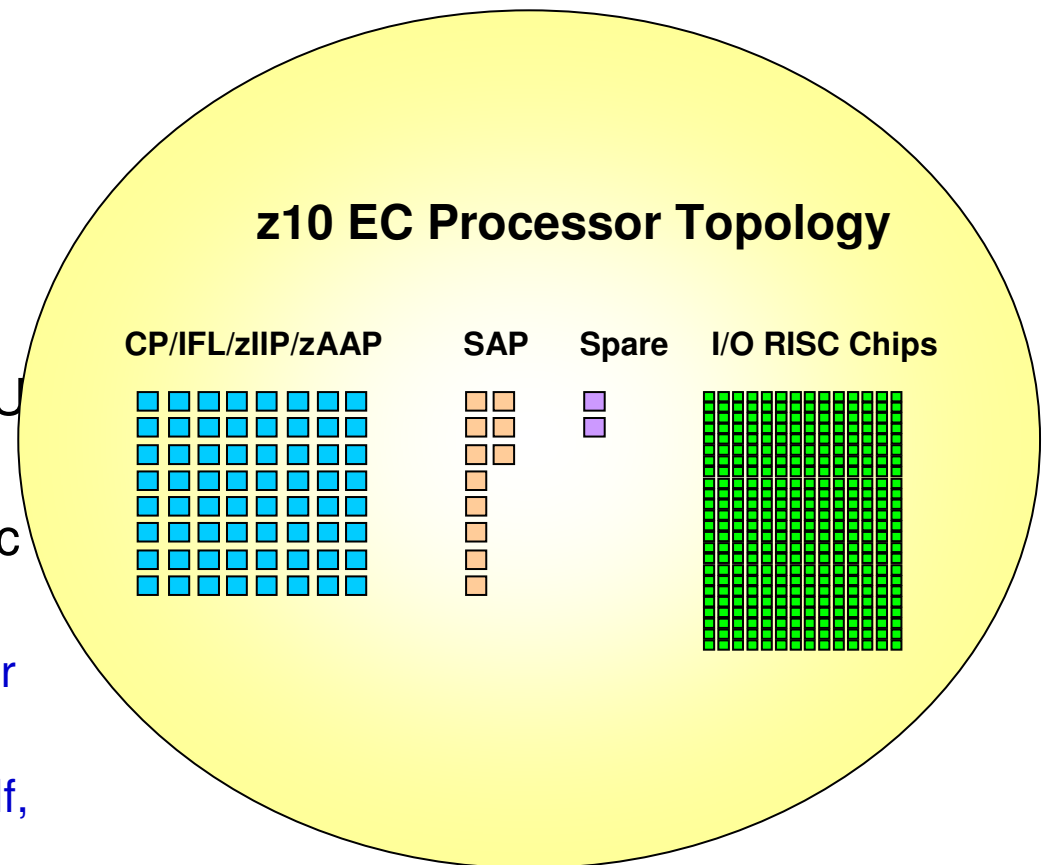
Hierarchical Storage Management (HSM) Autonomically Migrates Data For Archival



- References to data typically diminish over time
- Hierarchical storage management automatically moves older data to slower devices
 - ▶ Reference to migrated data initiates immediate retrieval to faster devices
- DFSMSHsm provides automated hierarchical storage management for System Z
 - ▶ Distributed servers require a separate product like Tivoli Storage Manager for Space Management

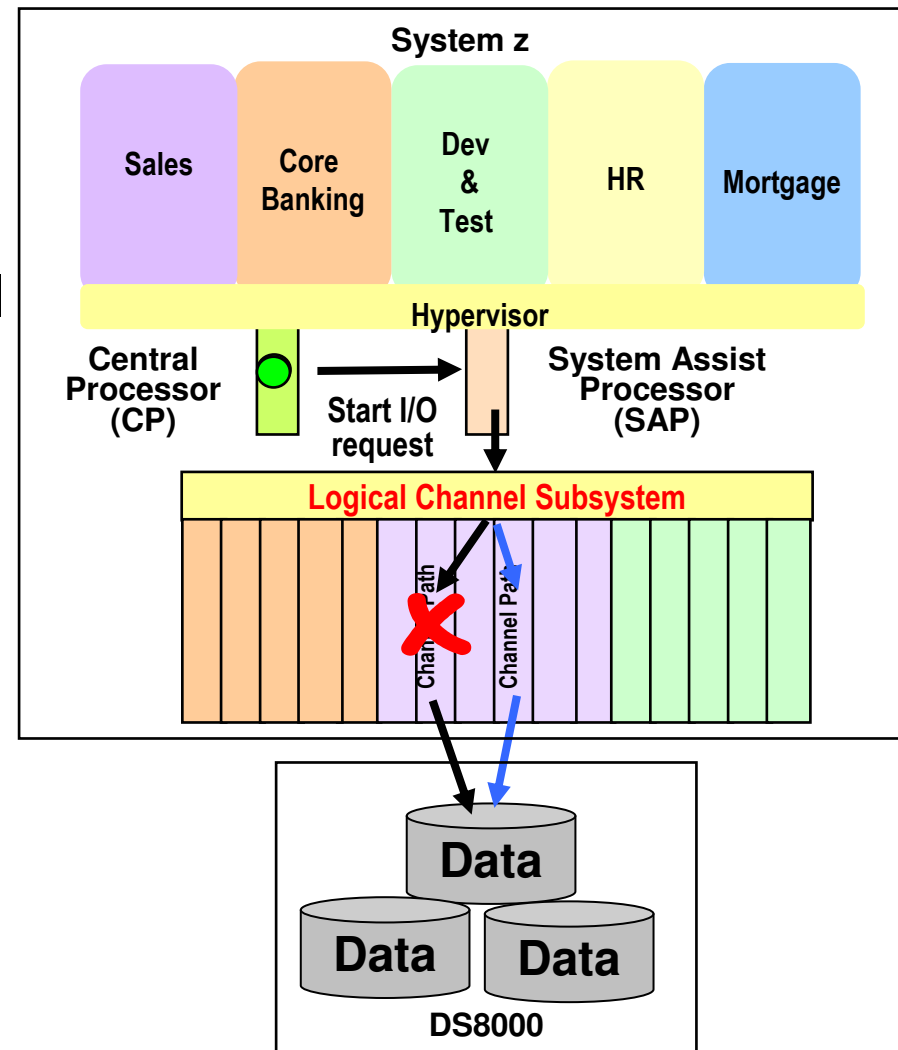
System z Also Has Dedicated I/O Hardware To Enhance Performance

- Offload I/O operations to dedicated hardware
- Up to 11 System Assist Processors (SAP) coordinate I/O requests
- Up to 336 RISC processors handle I/O operations
- I/O Offload saves general purpose CPU cycles
- Maximum I/O Bandwidth of 288 GB/sec without impact to workload capacity
- HP Superdome uses general processors for I/O – no dedicated processors
 - ▶ Sustained I/O bandwidth less than half, while impacting workload



Virtualization Of I/O Enables Redundant I/O Paths

- I/O Virtualization provided by Logical Channel Subsystem
 - ▶ Up to 1024 logical channel paths
- Virtualization enables optimal Physical I/O path to be used
 - ▶ Dynamic path selection
 - ▶ Load balances I/O traffic
- Transparent Failover
 - ▶ SAP recovers I/O operations in progress and switches to alternate path



Solid State Disk Drives Are Here To Revolutionize Storage

- Semiconductor storage delivered in DS8000 storage subsystems
 - ▶ Random access solid state storage – no moving parts
 - ▶ Electronically erasable medium
- Response times is around 0.8 milliseconds in contrast to 6 milliseconds for a typical hard disk drive
 - ▶ 5-10x improvement in throughput and queries
 - ▶ SSD drives can sustain I/O rates two orders of magnitude higher than traditional spinning disk
 - ▶ Reduce the “batch window”
 - ▶ DFSMS automatically controls allocation of new datasets to SSD drives
- Cost reductions
 - ▶ 75% reduction in space
 - ▶ 80+% reduction in power and cooling
 - ▶ Reduce RAM requirements



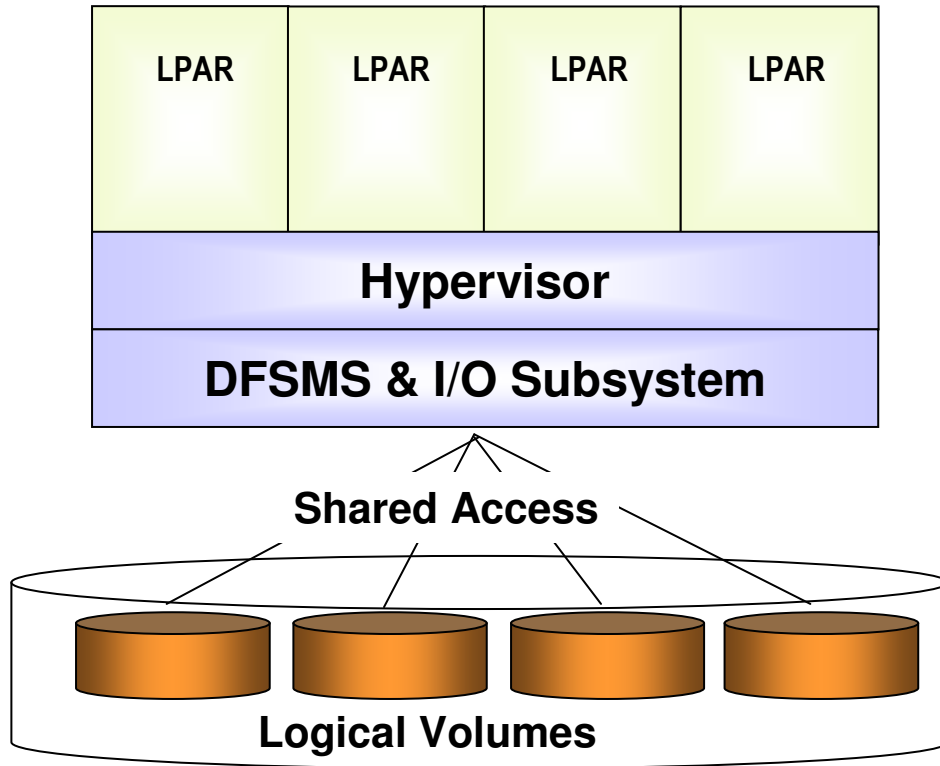
IBM DS8000 And TS7700 Provide High Capacity Storage For System z

- DS8000 supports a mix of disk drive types up to 461 TB
 - ▶ Maximum of 1024 disk drives
 - ▶ Solid State Disk drives (146GB)
 - ▶ 450 GB Fiber Channel Hard Disk Drives (450 GB)
- Up to 4.9 million I/O Operations per second
- Stripe data across multiple RAID arrays
 - ▶ Minimize disk “hot spots”
- Data mirroring for business resilience
 - ▶ Synchronous copies up to 300 km apart
 - ▶ Asynchronous copies over virtually unlimited distances
- Supports System z Extended Address Volume
 - ▶ Up to 223 GB per volume
- Supports Dynamic Volume Expansion
 - ▶ Increase volume size while running
- TS7700 provides virtual tape solution
 - ▶ Up to 70TB disk cache and 11PB capacity with TS3500 Tape Library



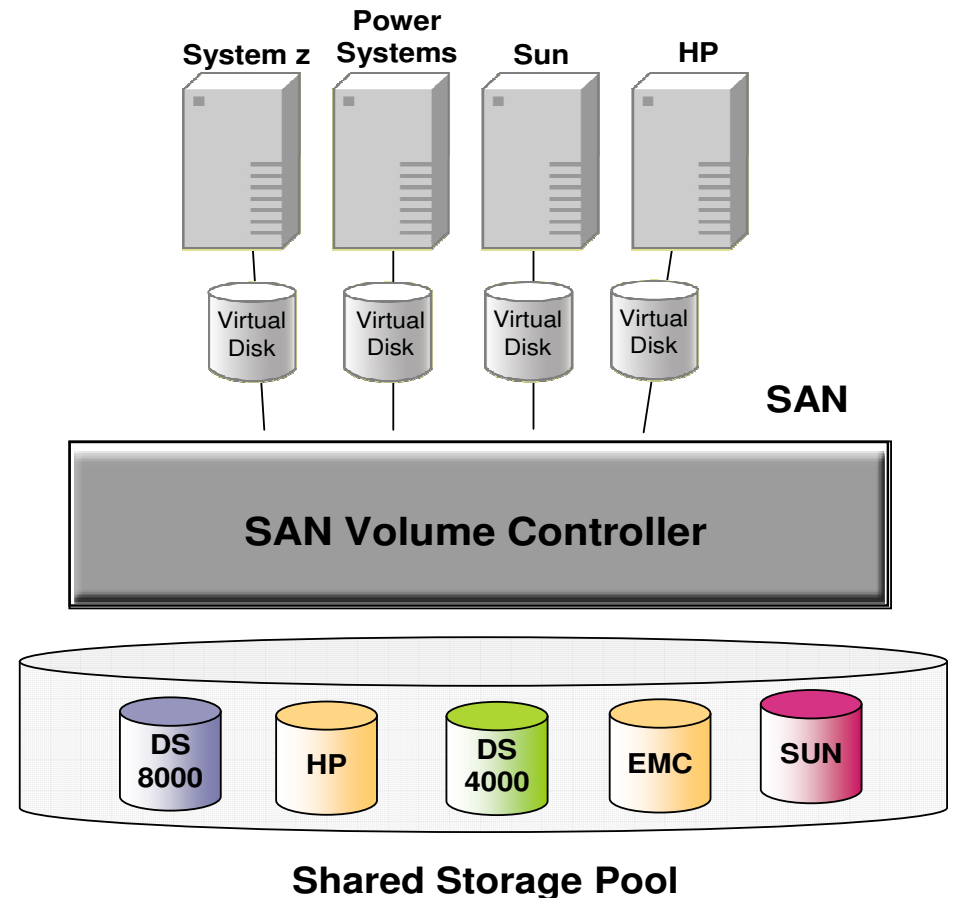
Storage Virtualization Is Built Into System z... Distributed Solutions Need Additional Products

System z



**Virtualized storage
(DS8000 & TS7700)**

Distributed



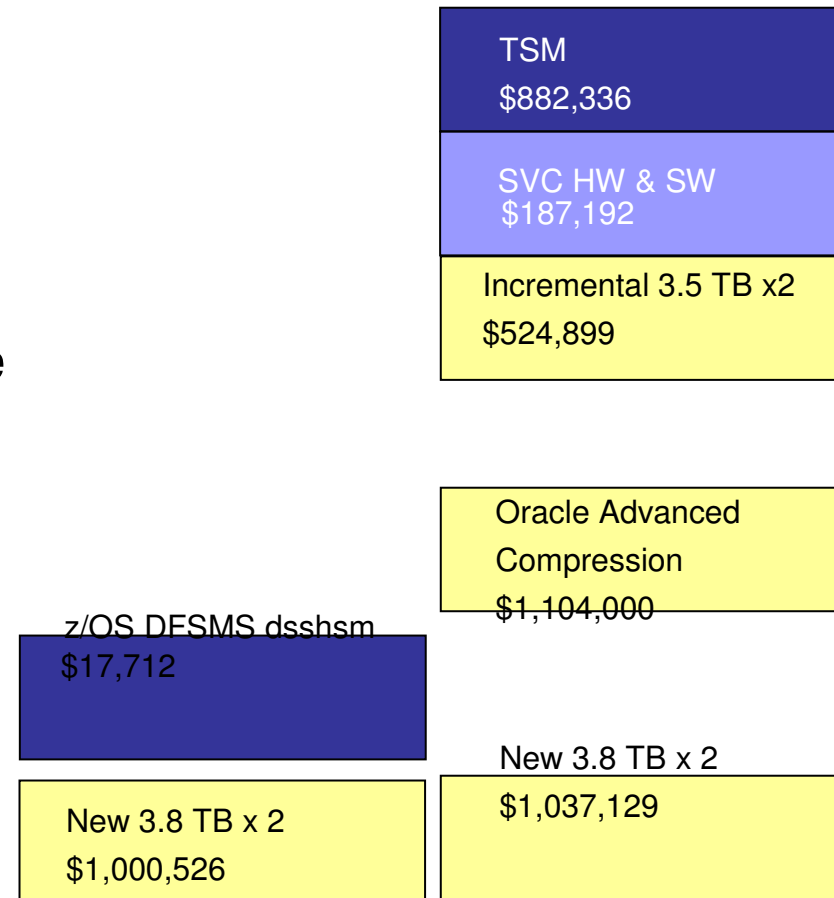
**PAV supports parallel access of logical volumes within the same system
and MA supports I/O parallelism across different systems**

DB2 Hardware Compression For System z Further Reduces Storage Costs

- Data Warehouses (TPC-H benchmark):
 - ▶ 62% (DB2) vs. 27% (Oracle)
- Save over **TWICE** as much on disk space over Oracle
- DB2's compression also saves on memory and I/O used
 - ▶ You'll need less buffer cache than with Oracle
 - ▶ You'll also do less I/O than with Oracle
 - ▶ You'll also need substantially less backup storage space
- Flexible DB2 compression algorithm applies to more database tables
 - ▶ Oracle algorithm limitations limits its effectiveness

System z And DB2 Reduce The Cost Of Storage By 73% For A New 10 TB Database

- For new storage capacity, 3.8TB x 2 (Primary+Secondary),
 - ▶ DS8100 for System z and HPXP2400 for Distributed
- Data Compression (10TB Storage)
 - ▶ System z – No incremental storage required, since DB2 uses built-in hardware compression, which supports up to 62%
 - ▶ Distributed – Incremental 3.5TB x 2 capacity since Oracle Advanced compression supports up to 27%
- Storage Management (HSM) and Virtualization (Data Sharing)
 - ▶ System z – DFSMS
 - ▶ Distributed – San Volume Controller (SVC) for Virtualization and Tivoli Storage Management (TSM) for HSM



System z \$1.02M

Distributed \$3.74M

Philippine Airlines Selects System z And IBM System Storage To Support Exponential Growth

- As PAL prepared to launch new routes to both domestic and international destinations, it realized it needed to upgrade its current information infrastructure
- PAL required better performance and superior throughput of the storage systems to run more efficiently. PAL also required an off-site fallback storage for business continuity and disaster recovery
- PAL replaced three different multi-vendor disk systems with an IBM storage solution consisting of the IBM System Storage Turbo DS8300

New Intelligence Needs To Be Extracted From The Exploding Information Silos

The mortgage line of business is doing well. Our late payment rates are very low.



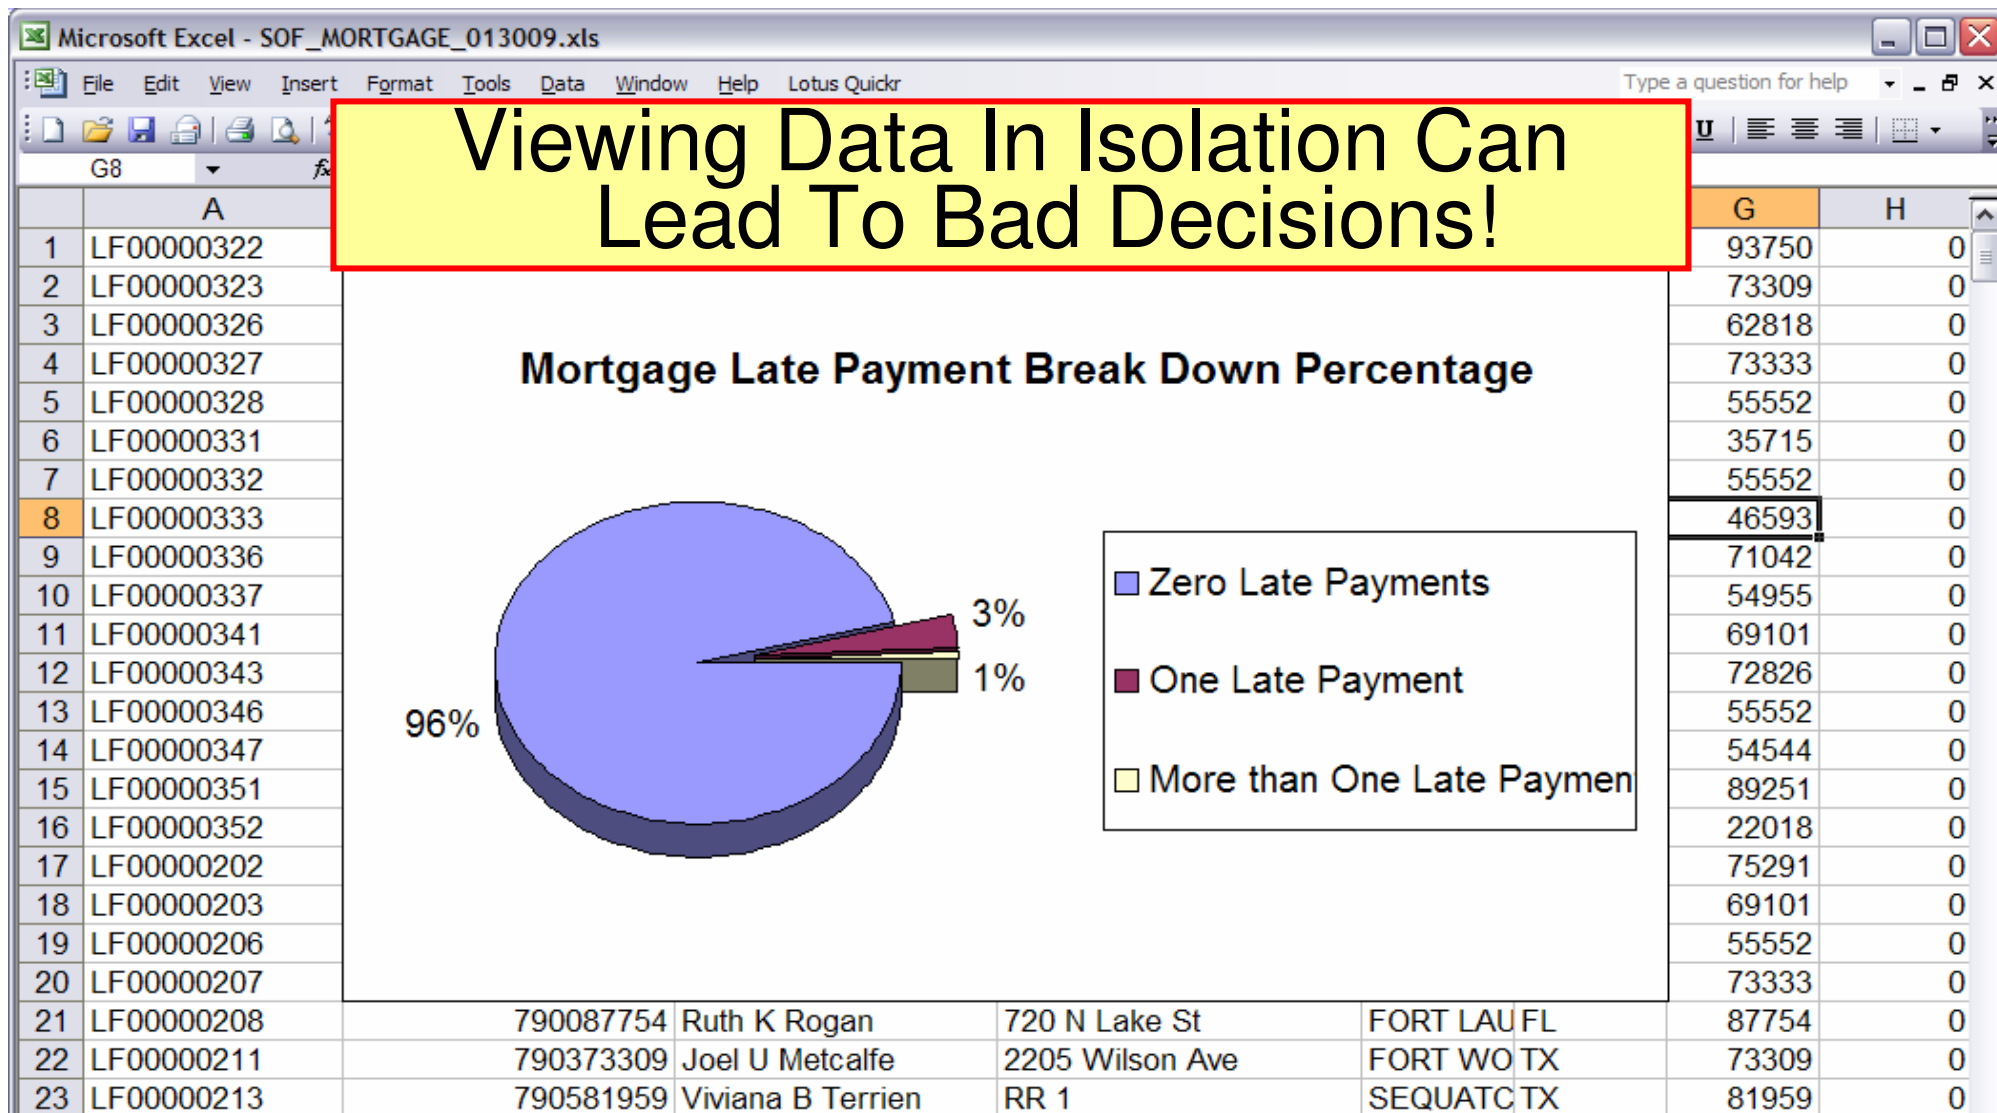
**Mortgage Line of Business
VP**

Then why are other areas of our business seeing problems...



**Service Oriented Finance
CEO**

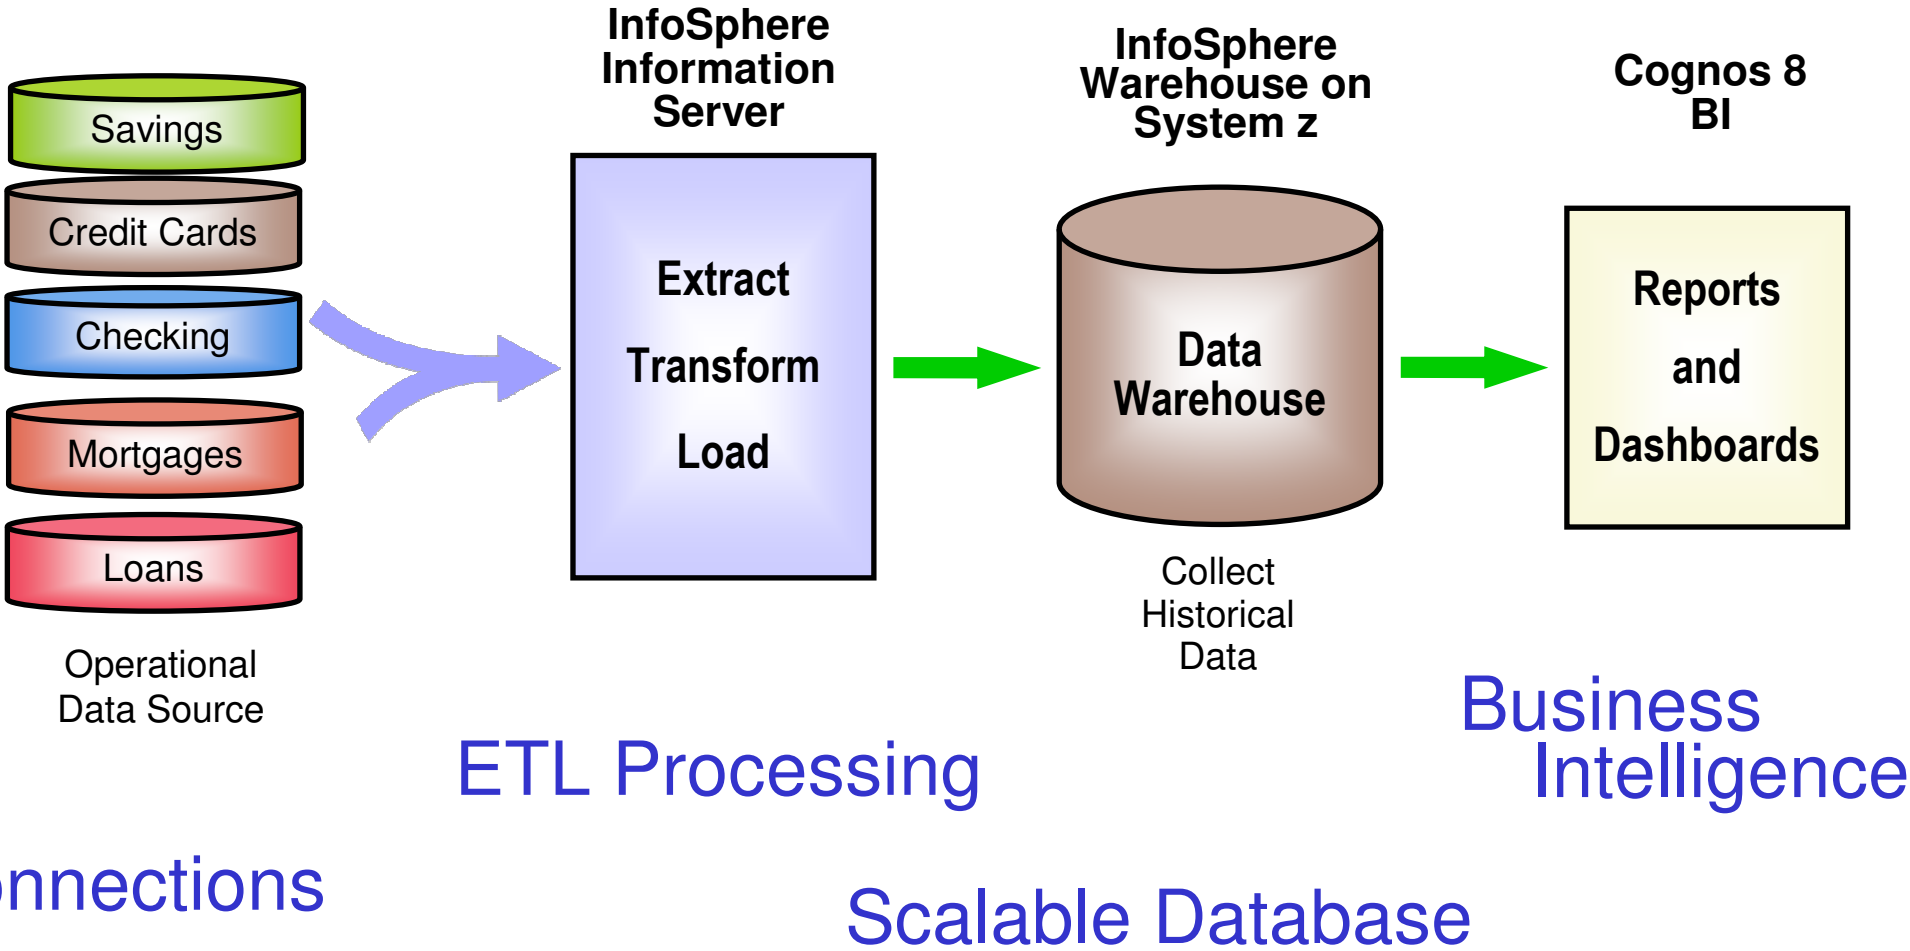
Isolated Customer Information Leads To An Incomplete View Of The Business



Service Oriented Finance Needs A Complete View Of Their Customers

- **Problem:** Segmented lines of business store their data in isolated silos
 - ▶ Banking, Credit Cards, Mortgage, Consumer Loans
- **Requirement:** Identify customer trends faster by viewing data from across all areas of business
- **Solution:** Create new intelligence by building an enterprise data warehouse containing a complete view of customer information

Create New Intelligence With IBM Information Management Software



Industry Data Models Help You Get Started

Where do we start?



**Service Oriented Finance
CIO**

IBM industry data models can help you get started quickly.



IBM

IBM Industry Data Models Accelerate Your Data Warehouse Solution

- Industry Data Models are:
 - ▶ Best practices from over 400 IBM clients
 - ▶ Built on InfoSphere Information Server and InfoSphere Data Architect
- Industry Data Models include:
 - ▶ Enterprise Data Warehouse (EDW) Model
 - ▶ Business Terminology Data Model
 - ▶ Business Solution Templates (BST)
- Industry Data Model Business Benefits
 - ▶ 83% report their Data Warehouse is better aligned with business needs
 - ▶ Over 50% report that businesses are now getting the information they want
- Industry Data Model Development Benefits
 - ▶ 15-20% cost savings to build the warehouse
 - ▶ 20-25% decrease in the time spent in design phase
 - ▶ 30-40% decrease in time spent in the modeling phase

Source: Hurwitz

InfoSphere Warehouse on System z Is An Excellent Base For Your Data Warehouse

- Based on DB2 for z/OS
- Superior scalability due to System z sysplex exploitation
- Parallel queries, Materialized Query Table, Star Join Enhancements optimize performance
- Near continuous on-line availability
- System z I/O bandwidth benefits warehouse performance
- Data compression beats Oracle
- Proven security
- zIIP exploitation achieves lowest cost
- Benefits from built-in storage virtualization

Rapid Data Integration With InfoSphere Information Server

Data integration has many complexities; Metadata, ETL, connectivity, performance, etc. How can we simplify our approach?



**Service Oriented Finance
CIO**

IBM has a consolidated platform that overcomes the difficulties of data integration. Let me tell you about InfoSphere Information Server



IBM

IBM InfoSphere Information Server

A consolidated platform for information integration

IBM InfoSphere Information Server for System z

Understand



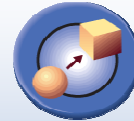
Discover, model, and govern information structure and content

Cleanse



Standardize, merge, and correct information

Transform



Combine and restructure information for new uses

Deliver



Replicate, virtualize and move information for in-line delivery

Platform Services

Parallel Processing Services



Connectivity Services



Metadata Services



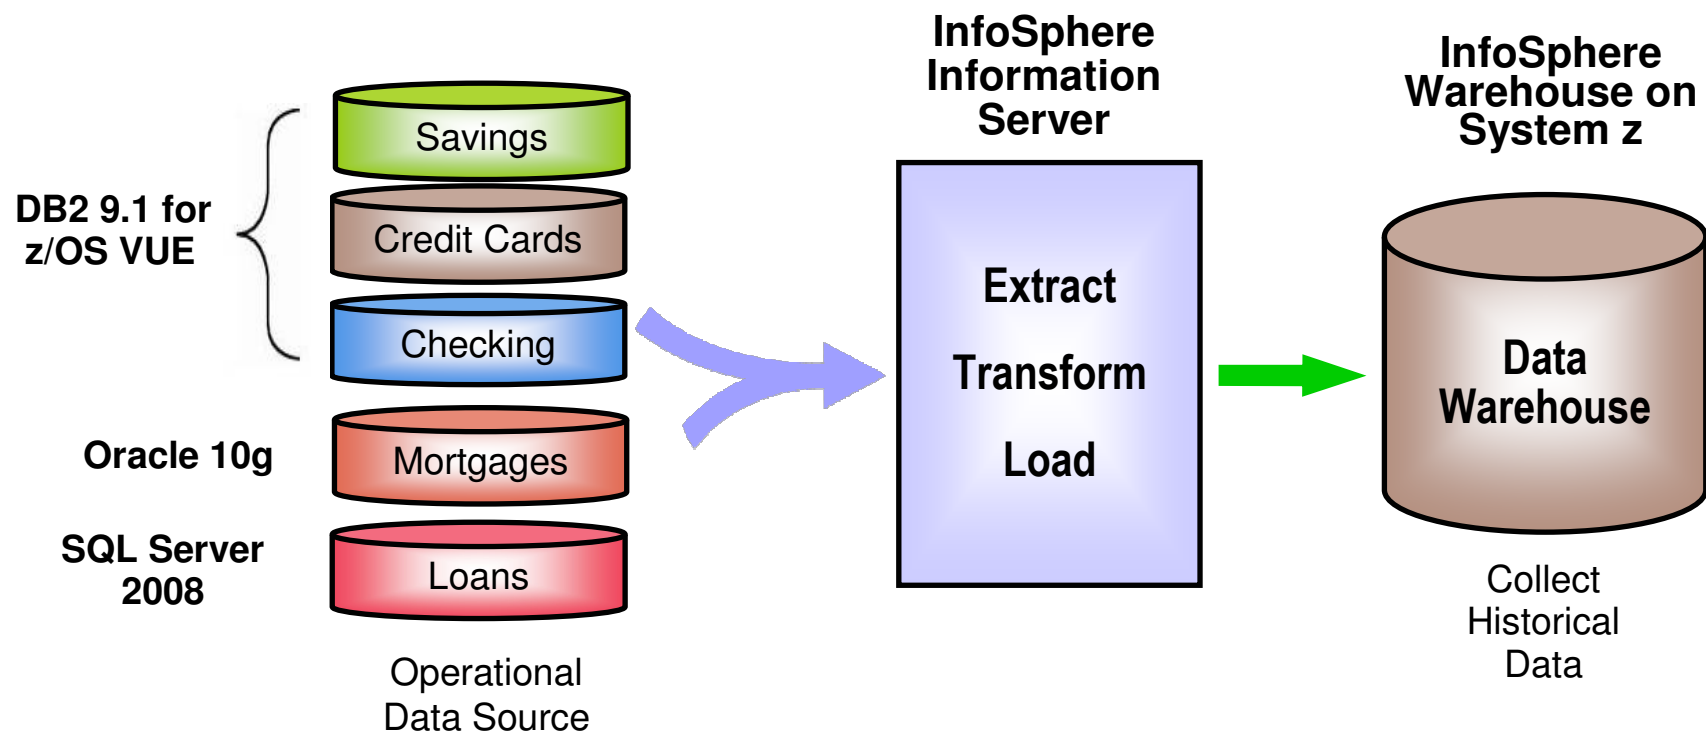
Administration Services



Deployment Services



InfoSphere Information Server Can Load Your Data Warehouse



Extract, Transform, And Load (ETL) Jobs Map Data From Sources To Targets

A few simple examples:

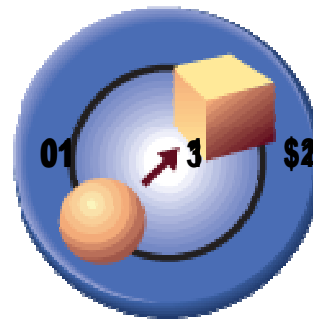
- Mapping source columns to targets
- Different column names and order
- Generating new column values
- Converting data types and formats



PROD ID	CUST ID	SOURCE ID	QTY	BAL	SALEDATE
000 101	100	01	1	\$10,000.00	2007-02-28
000 121	100	01	3	\$500.50	2007-02-28
000 102	101	01	1	\$ 20,000.00	2007-03-01

Target: Data Warehouse

000 121 100



01 3 \$20,000.00 2007-02-28

Transform

PRODUCT	QTY	CUSTNO	BALANCE	DATE
101	01	100	10000	02-28-2007
121	03	100	500.50	02-28-2007
102	01	100	20000	03-01-2007

Source: Operational Data

A successful data integration project requires a detailed specification for the business goals and technical requirements!

InfoSphere FastTrack Creates Data Maps And Specifications For Your Data Integration Projects

- Create simplified data maps and transformations using drag and drop
 - ▶ Automatically discover source and target columns
 - Uses database introspection and Web 2.0-style tagging
 - Use business terms to accurately match source to target
- Data analysts and developers share project specifications
 - ▶ Collaboration and reuse improve productivity
 - ▶ Use metadata common to all Information Server tools
 - ▶ Standard formats and centralized management for governance
 - Synchronize work across global teams
- Generate ETL code directly from job specifications
 - ▶ Reduces costs and errors in ETL job development

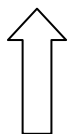
Oracle doesn't offer any of these capabilities

InfoSphere FastTrack Automatically Discovers Data Mappings Using Business Term Tags

Mortgage To Warehouse Mapping Specification

Source	Target	Tag
Times_Past_Due	Credit_Events	Failure_To_Pay
Current_Balance	Ending_Balance	Period_Balance
Account_Num	Account_ID	Arrangement
Account_Holder	Customer_ID	Party_ID

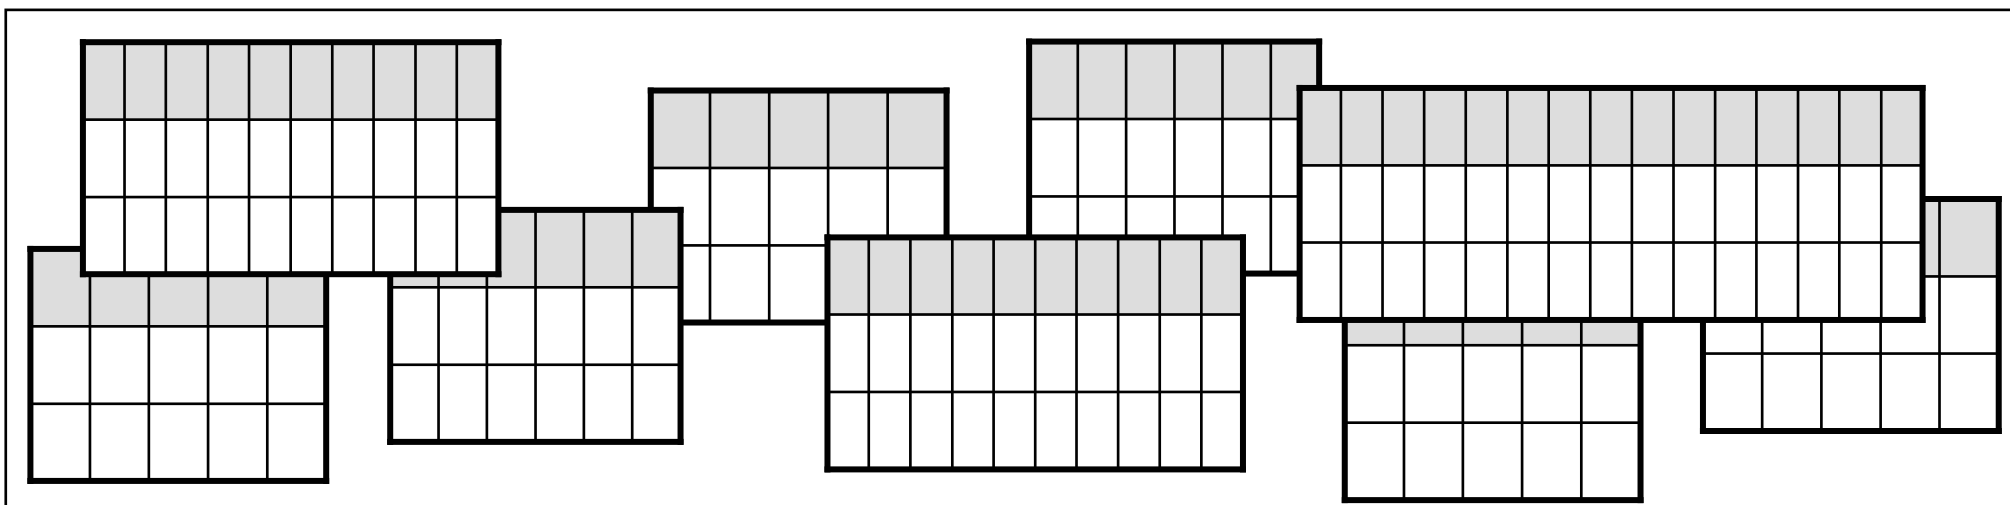
Source Discovery
Mortgage.Times_Past_Due
Checking.NSFCCount
Loans.MissedPayments



Metadata From Information Server

Tags come from:

- Industry data models
- Your corporate standards



DEMO: Use InfoSphere FastTrack To Create ETL Specification For Warehouse

- Use discovery feature to find source columns matching business term tags
- Generate ETL job for InfoSphere DataStage and execute it

The screenshot displays the IBM Information Server FastTrack interface. The main window is titled "LoadMortgageFact" and shows a mapping editor with source and target columns. The source columns include "Results available.", "SOF_MORTGAGE.CURRENT_BALANCE", "SOF_MORTGAGE.ACCOUNT_ID", "SOF_MORTGAGE.ACCOUNT_HOLDER_ID", and "SOF_MORTGAGE.DATE_ID". The target columns include "BALANCE_SNAPSHOT_FACT.CREDIT_EVENTS", "BALANCE_SNAPSHOT_FACT.ENDING_BALANCE", "BALANCE_SNAPSHOT_FACT.ACCOUNT_ID", "BALANCE_SNAPSHOT_FACT.CUSTOMER_ID", "BALANCE_SNAPSHOT_FACT.BRANCH_ID", and "BALANCE_SNAPSHOT_FACT.DATE_ID". The Business Term column lists terms like "FAILURE TO PAY [WBG]", "PERIOD BALANCE [WBG]", "ARRANGEMENT IDENTIFIER [WBG]", "INVOLVED PARTY IDENTIFIER [WBG]", "BUILDING IDENTIFIER [WBG]", and "STATUS DATE [WBG]".

The "Discover" tab is active, showing a table of discovered source columns:

Name	Score %	Data Type	Related
ZLNXPCK.SOFIOD.FISHING.BALANCE_SNAPSHOT_FACT.CREDIT_EVENTS	100	Integer	true
ZLNXPCK.Oracle Mortgages.INST1.SOF_MORTGAGE.TIMES_PAST_DUE	100	Double	true

A "Match Details" dialog box is open, showing the match between "FAILURE TO PAY" and "TIMES_PAST_DUE". The match is defined as a "Glossary business term" relationship.

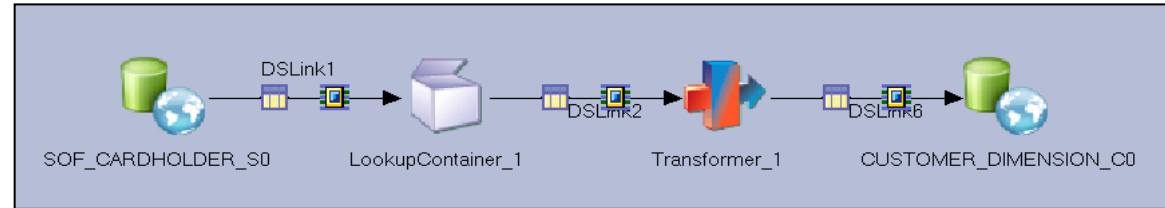
The "Database Metadata" pane on the right shows a list of tables and columns, including "CITY [Char]", "CURRENT_BALANCE [Decimal]", "INT_RATE [Decimal]", "JOINT_ACCOUNT_HOLDER [Char]", "MONTHLY_PAYMENT [Decimal]", "NAME [Char]", "ORIG_AMOUNT [Decimal]", "ORIG_DATE [Timestamp]", "SOURCE [Char]", "SS_NUM [Char]", "STATE [Char]", "TERM_YEARS [Double]", "TIMES_PAST_DUE [Double]", "ZIP [Double]", and "ZIP_FOUR [Double]".

The "Properties" pane on the right shows the properties for the selected column, including Name Alias, Data Type (Double), Position (18), Nullability (true), Length (8), Precision (8), Scale, and Cardinality Type (Not Constrained).

The bottom of the interface shows the "Mapping Specifications" pane with buttons for "Generate Job", "Save", and "Close".

InfoSphere DataStage Creates The Technical Implementation Of Data Integration Jobs

- Creates graphical data integration jobs using hundreds of pre-built transformation and data quality functions
 - ▶ Batch & real-time operations
- Stores and retrieves metadata from Information Server
 - ▶ Allows easy reuse of integration work between projects
- Advanced parallel processing capabilities
 - ▶ Dynamic partitioning and pipelining
 - ▶ Scale jobs across additional hardware without modification
- Easily deploy data integration jobs as services for SOA



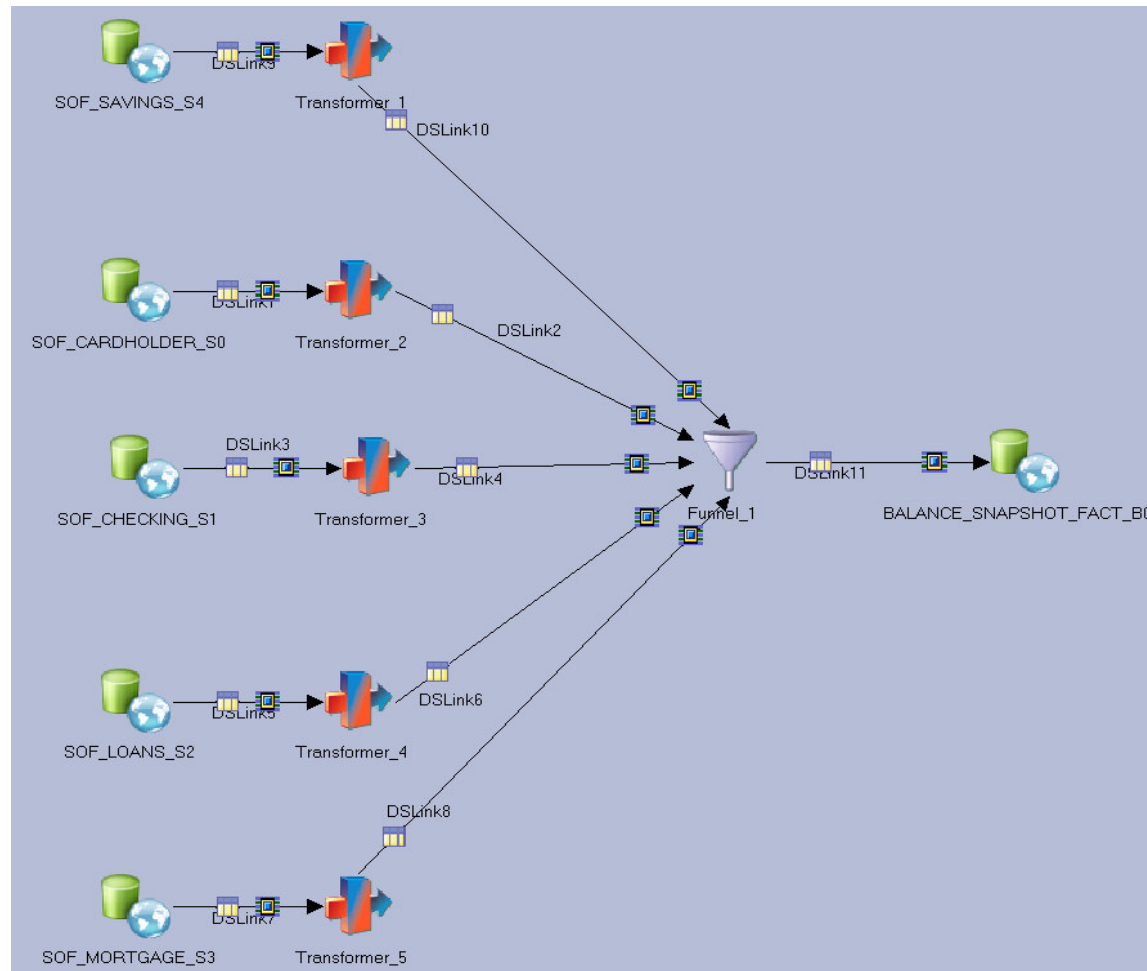
Database				
Classic Federation	DB2 UDB API	DB2/UDB Enterprise	DB2Z	Dynamic RDBMS
Informix CLI	iWay Enterprise	ODBC	Oracle Enterprise	Stored Procedure

Processing			
Aggregator	Change Apply	Change Capture	Compare
Compress	Copy	Decode	Difference
Encode	Expand	External Filter	Filter
FTP Enterprise	Funnel	Generic	Join
Lookup	Merge	Modify	Pivot
Remove Duplicates	Slowly Changing Dimension	Sort	Surrogate Key Generator
Switch	Transformer		

Real Time			
Java Client	Java Transformer	Web Services Client	Web Services Transformer
WebSphere MQ Connector	WISD Input	WISD Output	XML Input
XML Output	XML Transformer		

DEMO: Use InfoSphere DataStage To Load The Data Warehouse

1. Show the results of the ETL job that populated the data warehouse fact table



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
Netezza
Progress
RDB
RedBrick
SQL/DS
SQL Server
Sybase (ASE and IQ)
Teradata
Universe
UniData
NonStopSQL
And more.....



Offering more connectivity than Oracle

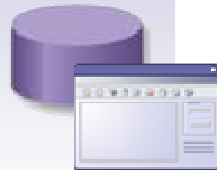
General Access

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
Turbolimage
Unify
And many more....



IBM Leads In Data Integration

- Only InfoSphere Information Server delivers unified metadata across all tools for collaboration and reuse
 - ▶ Oracle has no integration of metadata across products
 - ▶ Manual import/export required
- Model-driven design with FastTrack and DataStage speeds development
 - ▶ Oracle has no tools to help manage source to target mappings
- InfoSphere Information Server works in heterogeneous environments
 - ▶ InfoSphere gathers, processes, and cleanses more data from more sources than Oracle

"FastTrack enables our analysts to **capture more complete business requirements**. The ability to translate this information directly into DataStage jobs with up to 70 percent of the code completed will **significantly shorten our development lifecycle**."

- Roderich Hofmann, project manager, WAVE, IT-Solutions provider of Bank Austria and member of UniCredit Group

Using New Intelligence Creates New Business Opportunities

If we can identify our risky mortgage assets, we can work to remove them from our books



**Service Oriented Finance
CEO**

We can identify risky mortgage customers by watching their activities in other business areas

- ▶ Bounced Checks
- ▶ Missed Credit Card Payments
- ▶ Missed loan payments



**Mortgage Line of Business
VP**

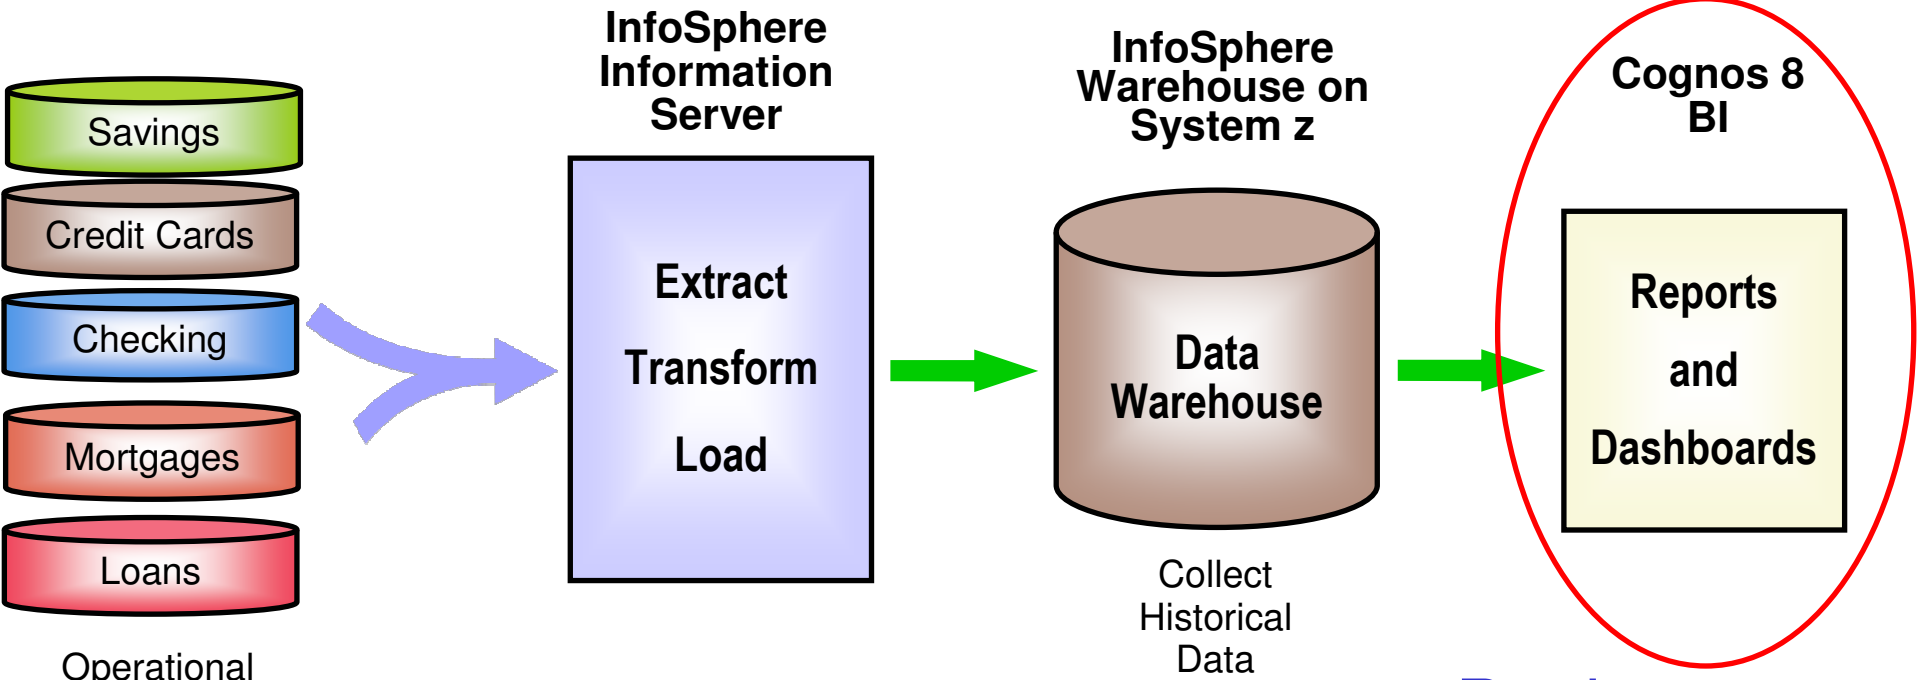
Create New Intelligence With IBM Cognos

Now that you have a data warehouse, you can create this new intelligence using IBM Cognos 8 BI



IBM

Use IBM Cognos 8 BI To Optimize Business Decisions



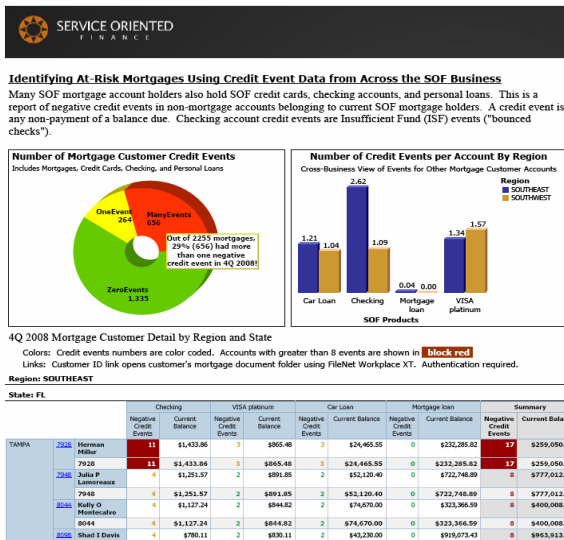
Connections

Scalable Database

Business Intelligence

DEMO: Identify Risky Mortgage Accounts Using Cognos 8 BI

1. Show report generated in Cognos Report Studio in PDF format
2. Report identifies high-risk mortgages by looking at negative credit events in customers other SOF accounts (CC, Checking, etc...)
3. Report uses both structured and unstructured data (link to mortgage data stored in FileNet)



That report is just what we need!

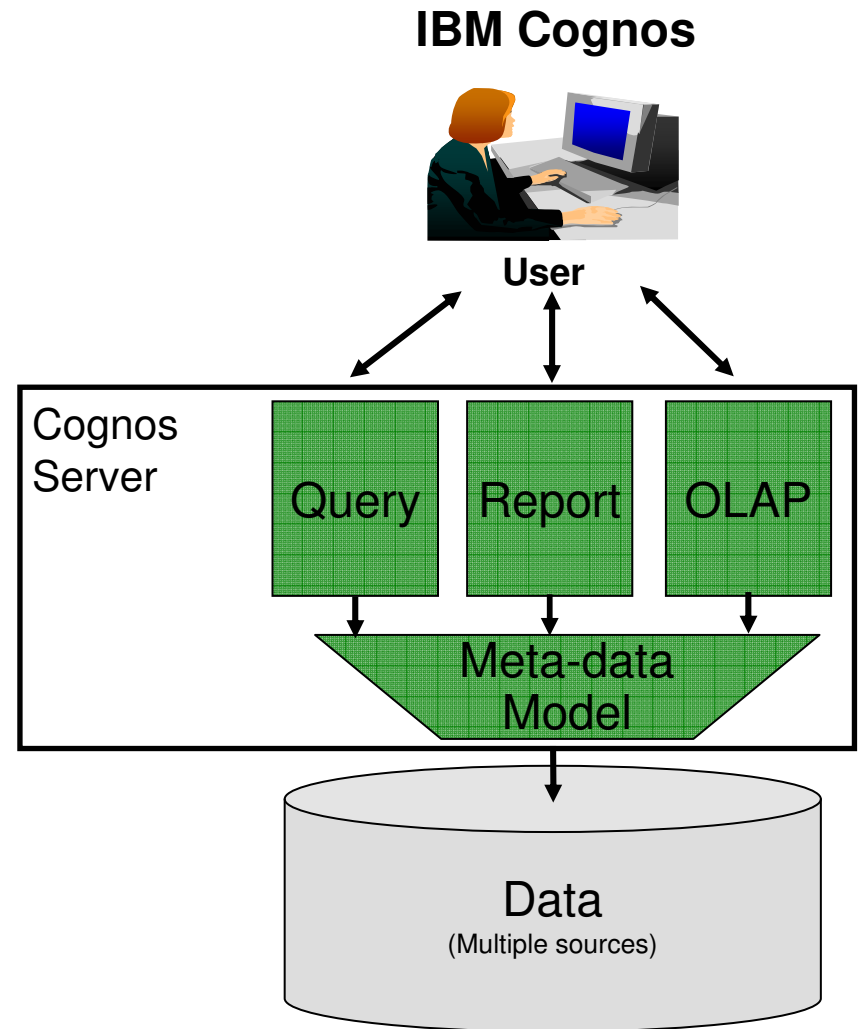


Mortgage Line of Business VP

- At risk customers are identified and contacted to refinance
- Risky mortgages can be sold

IBM Cognos Is An Integrated Platform Built On SOA

- Implemented in Java, runs on WebSphere
- 100% browser based access
 - ▶ Server side business intelligence
 - ▶ Users can access new intelligence from anywhere
- Easiest for IT to deploy and manage
 - ▶ Scales up and out across heterogeneous hardware and operating systems
 - ▶ Unified security
 - ▶ Unified administration
- Consistent user interface across tooling
 - ▶ Greater user satisfaction and increased business agility with lower IT costs
- Common meta data model
 - ▶ Author new intelligence assets once, consume anywhere
 - ▶ Common view enables open data strategy
 - ▶ Supports Unicode and multilingual features without recreating reports

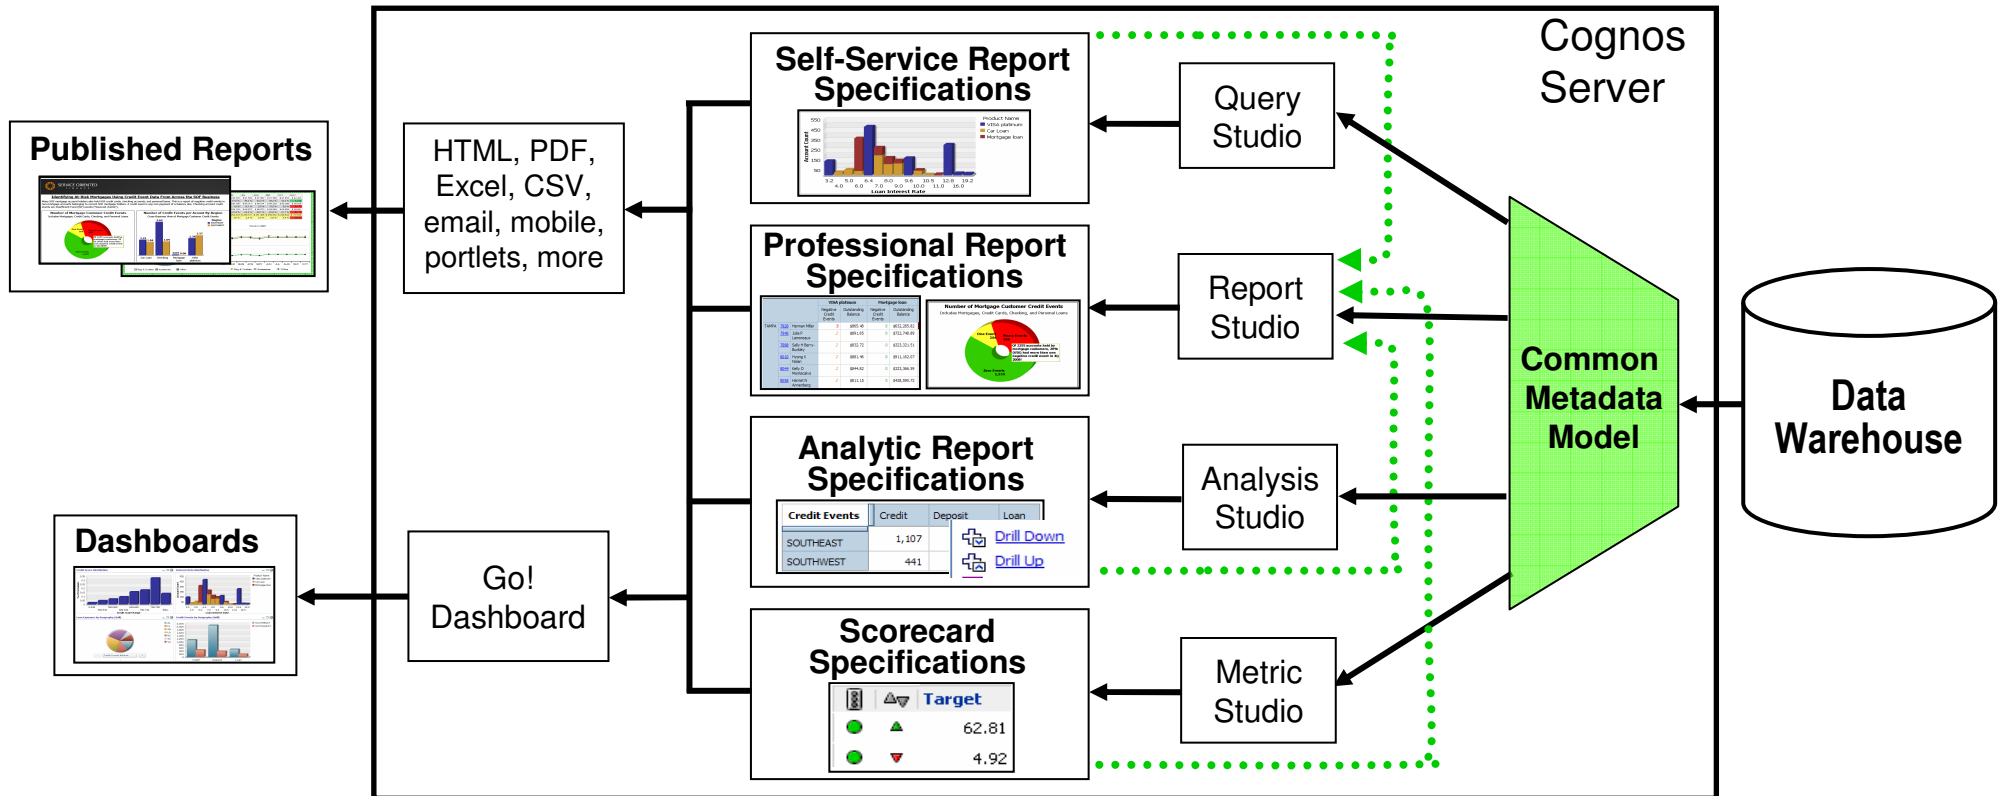


■ IBM Cognos

Users Can Create The Reports They Need Using Cognos 8 BI

- Query Studio is an easy to learn self-service reporting tool requiring minimal reporting knowledge
 - ▶ Helps alleviate report authoring backlog
 - ▶ Use existing self-service reports to create a new report
 - ▶ Modify the style and layout of self-service reports
- Report Studio is a professional reporting tool to create any style of report
 - ▶ Invoices, financial statements, inventory, payroll, etc
 - ▶ Provides “pixel-perfect” formatting with absolute control over visual layout
 - ▶ Library of lists, crosstabs, charts, maps, operators, constants, functions, filters, more

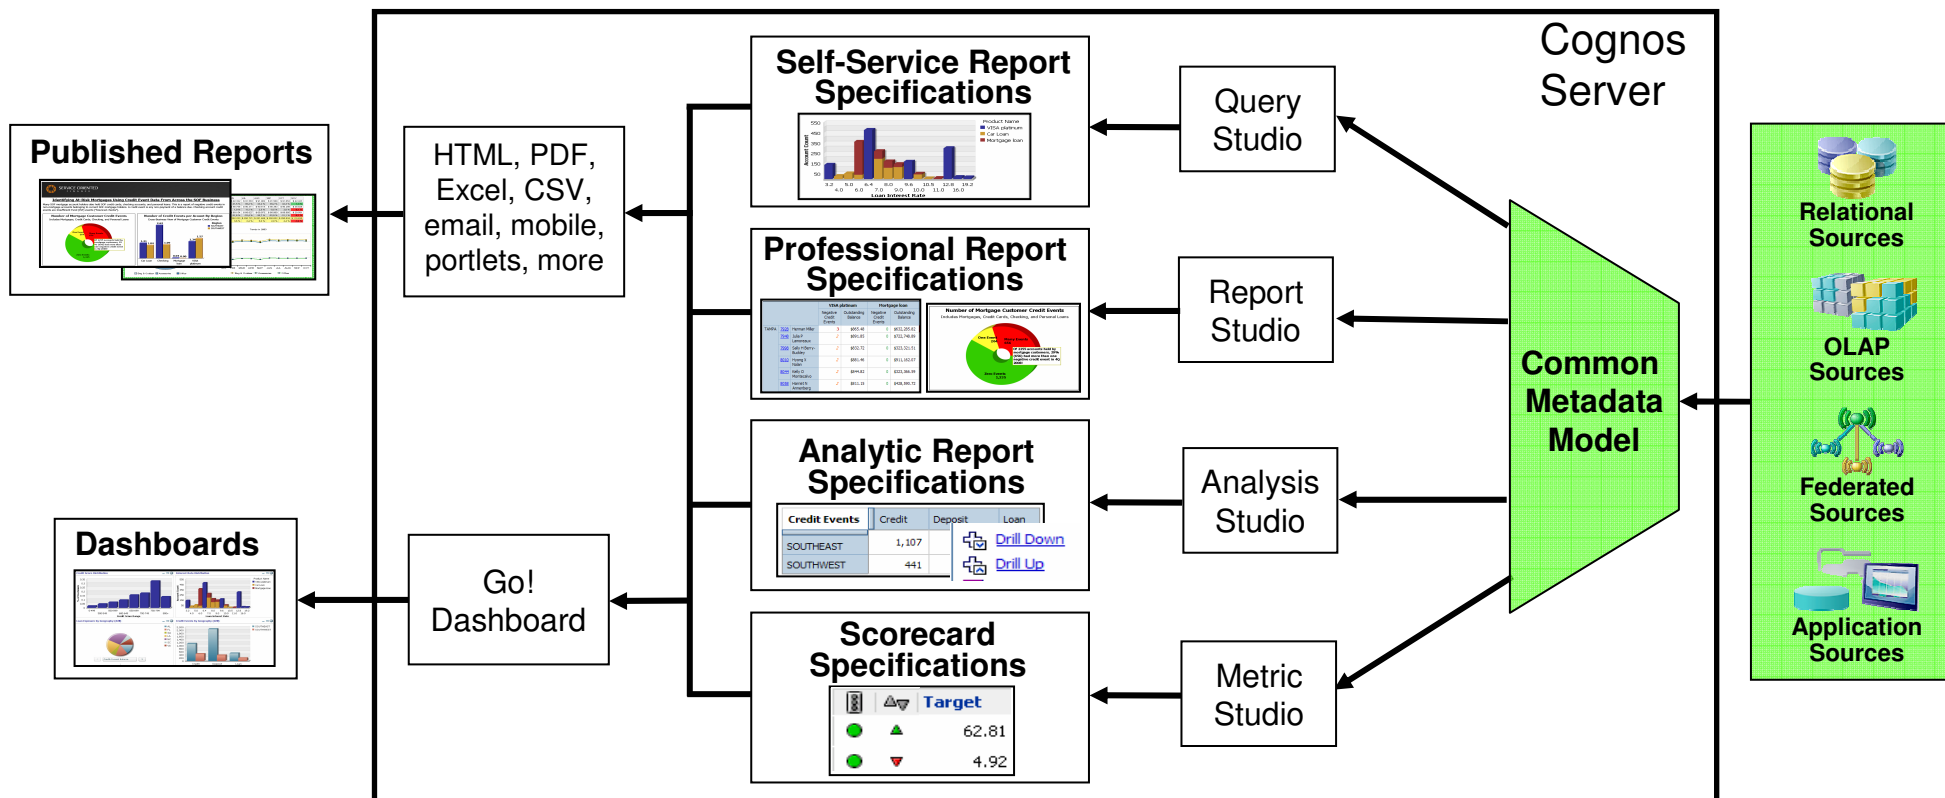
Reuse Trusted New Intelligence Assets Across the Cognos 8 Platform



- All new intelligence assets share a common metadata model and common report specification
- Author Once – Consume Anywhere
- Ensures consistent information and enables reuse across platform functions

- Oracle has multiple metadata models depending on source type
- Oracle has multiple different report formats
- Oracle cannot reuse assets between tools

Include Any Data Source In Your New Intelligence With the Cognos 8 Platform

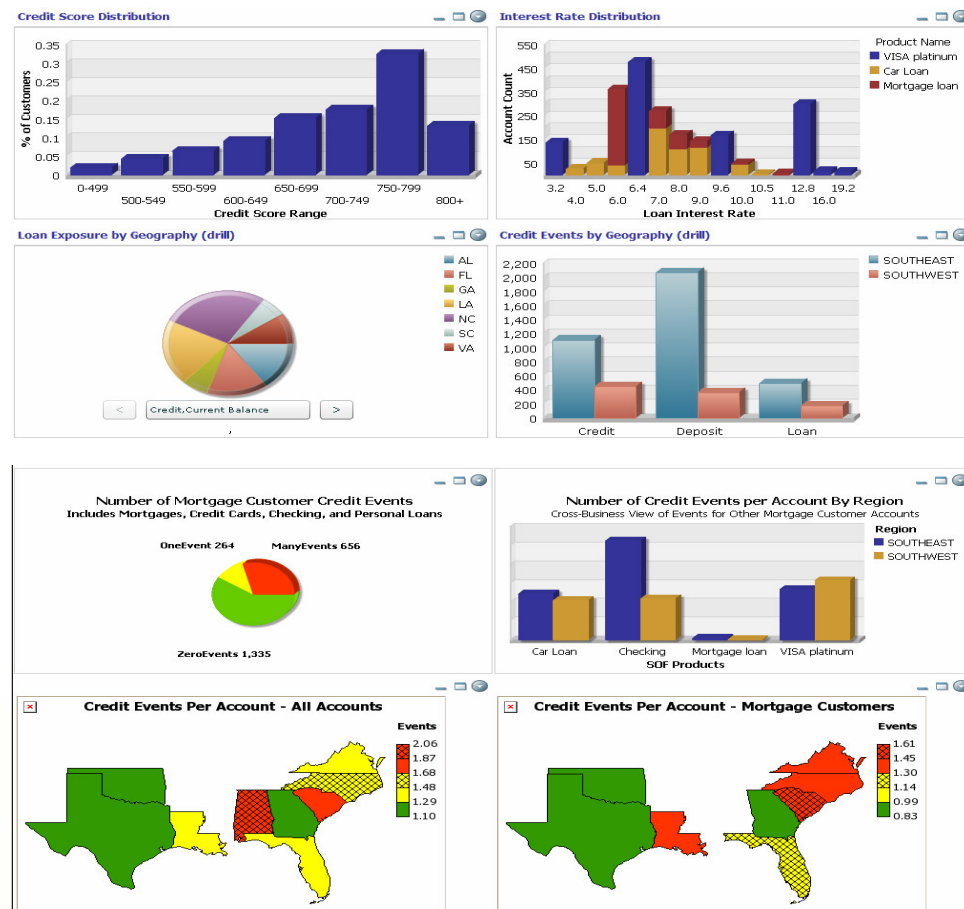


- Open data strategy enables a common view across a variety of data sources
- Support for application data sources such as SAP ERP
- Combine relational, OLAP, federated, and other data sources in any tool

- All capabilities access a trusted set of information defined in the common metadata model
- As sources change, metadata model can control and identify impacts to report specifications

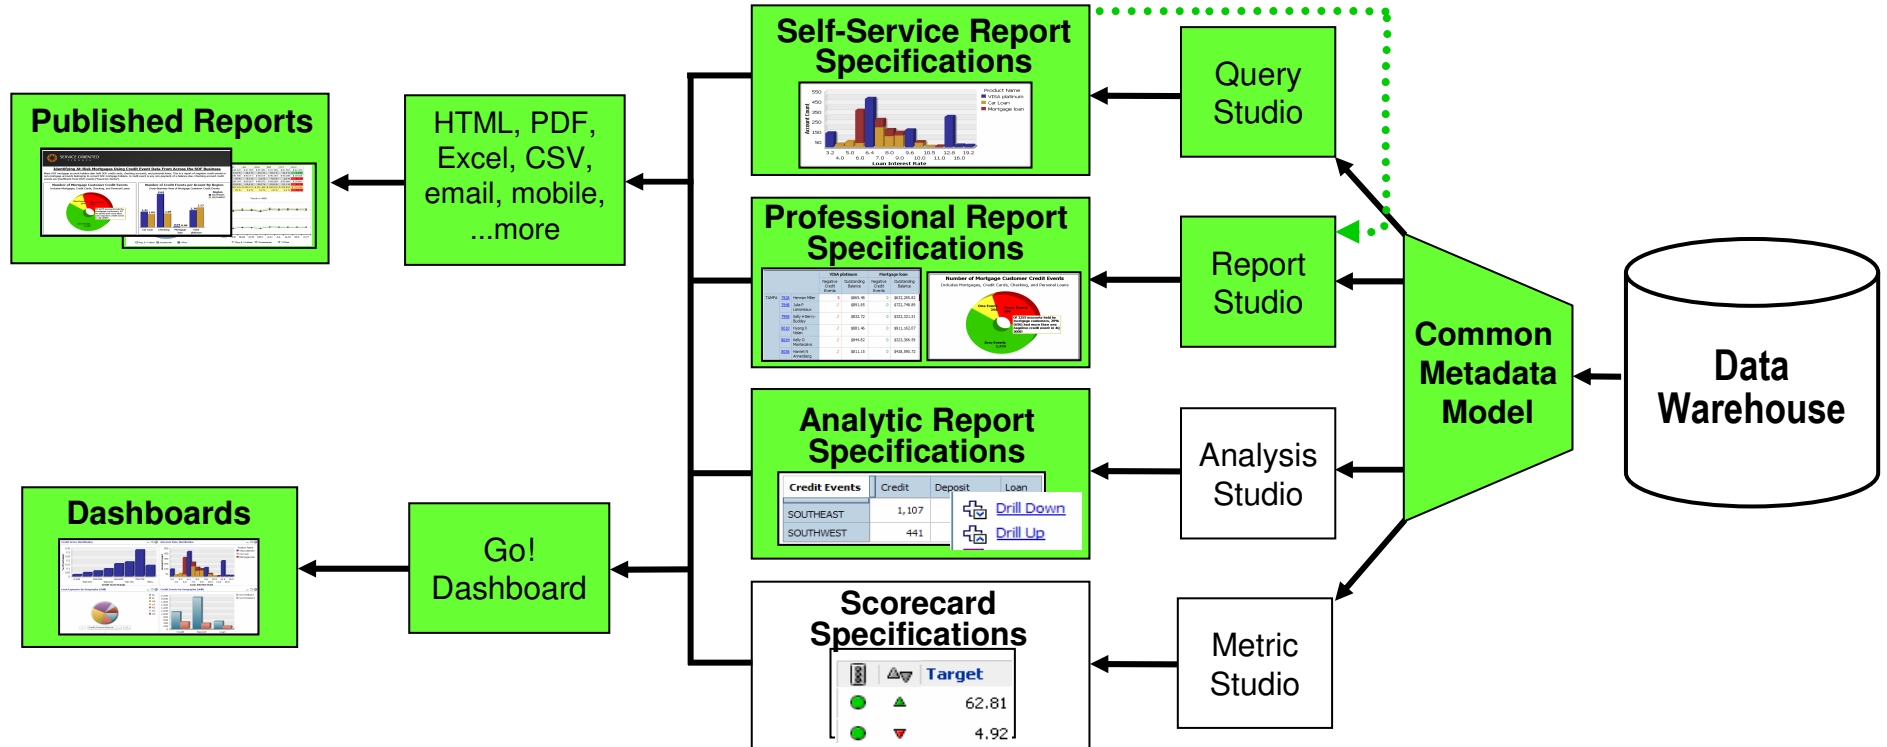
Cognos Go! Dashboard Enables You To Monitor Business Operations

- Cognos report specifications can be incorporated into dashboards using Cognos 8 Go! Dashboards
- What goes into a dashboard?
 - ▶ Self service reports
 - ▶ Professional reports
 - ▶ Analytical reports
 - ▶ Scorecards
 - ▶ RSS feeds, HTML, search, more
- Users can create their own dashboards from existing Cognos report assets
- Everything you need to monitor a particular aspect of the business
- Information from several different subjects areas presented at the same time
- Provides dynamic and visually appealing capabilities by using Adobe Flash



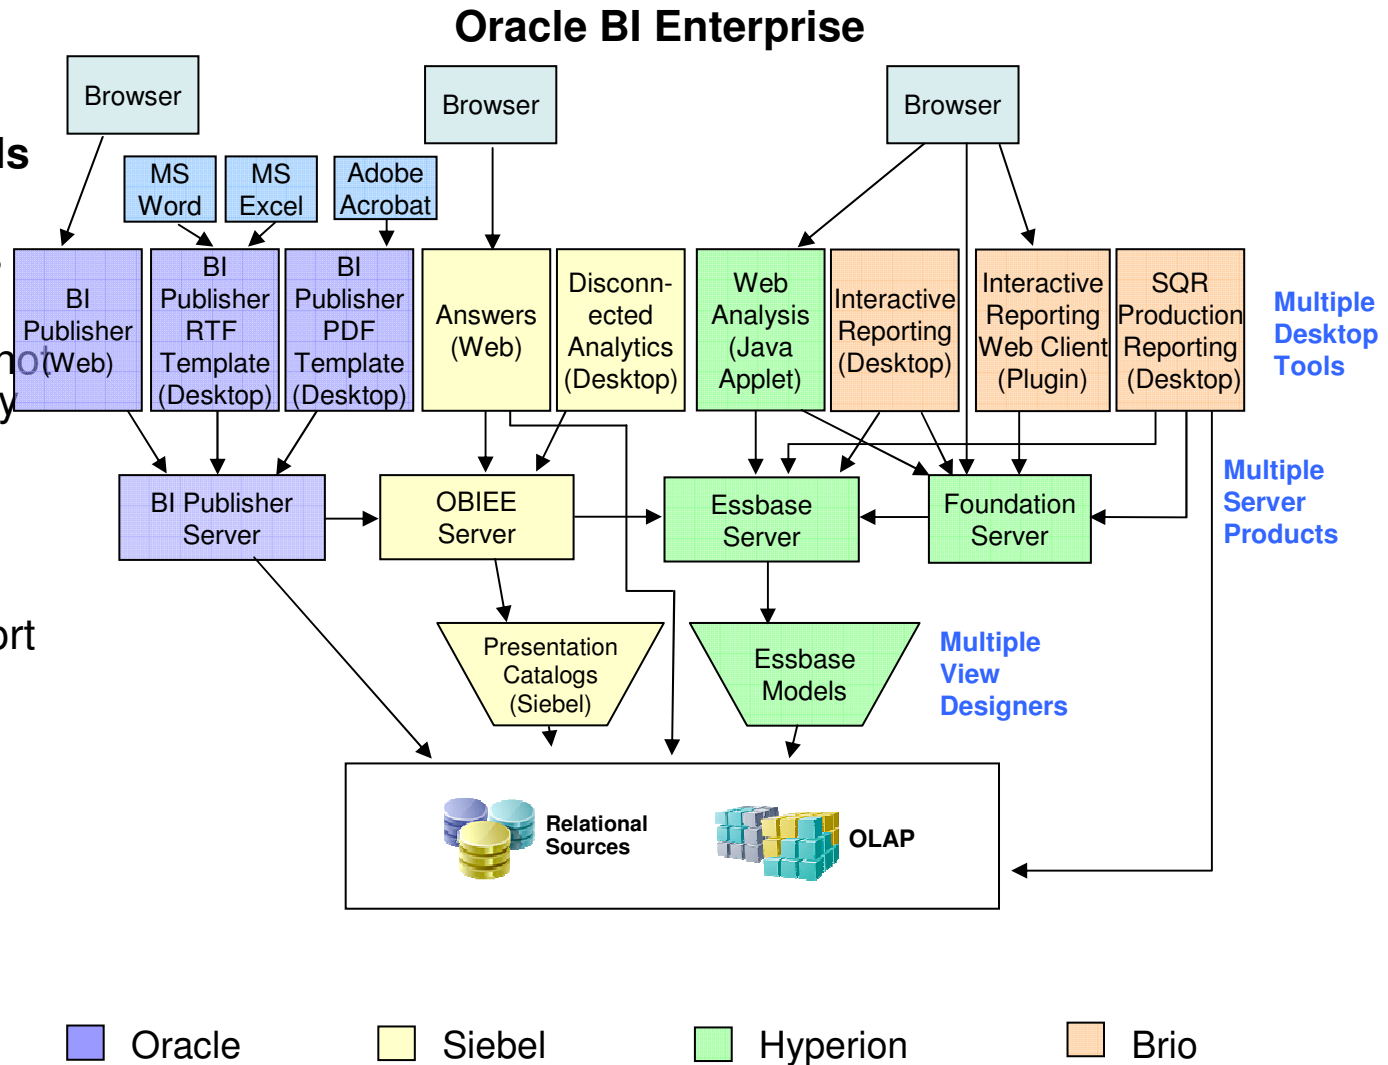
DEMO: Gain Business Insight Through IBM Cognos 8 BI

- Use Go! Dashboard to quickly monitor the business operations
- Use Cognos Query Studio to customize an existing report
- Open Cognos Report Studio and add a chart to the report



Oracle Business Intelligence Enterprise Edition Is A Complicated Bundle

- **Too Many Products!**
- **Multiple Desktop Client Tools**
 - ▶ Multiple report formats
 - ▶ Multiple metadata models
 - ▶ Creates More work
 - ▶ Report specifications cannot be shared or reused easily
 - No common meta data model
- **Creates IT Burden**
 - ▶ Install, Patch, User Support all happen at the desktop



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

Existing Mainframe



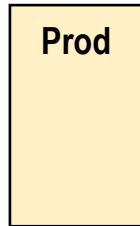
Existing z10:
2 GP 1,720 MIPS
DB2 and utilities
With 20Tb storage

Existing Disaster Recovery Site



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

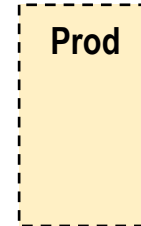
Add 1 LPAR for New Data Warehouse w 3.8 TB Storage



2,184 MIPS additional workload on z/OS and 1,840 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
\$5.71M

Capacity Backup:
2 GP
1 zIIP
2 IFLs

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

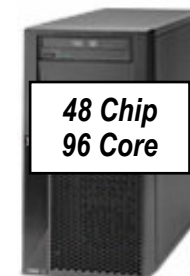
Prod



350,299*
Performance Units

And add Disaster Recovery W 7.3 TB storage

Prod

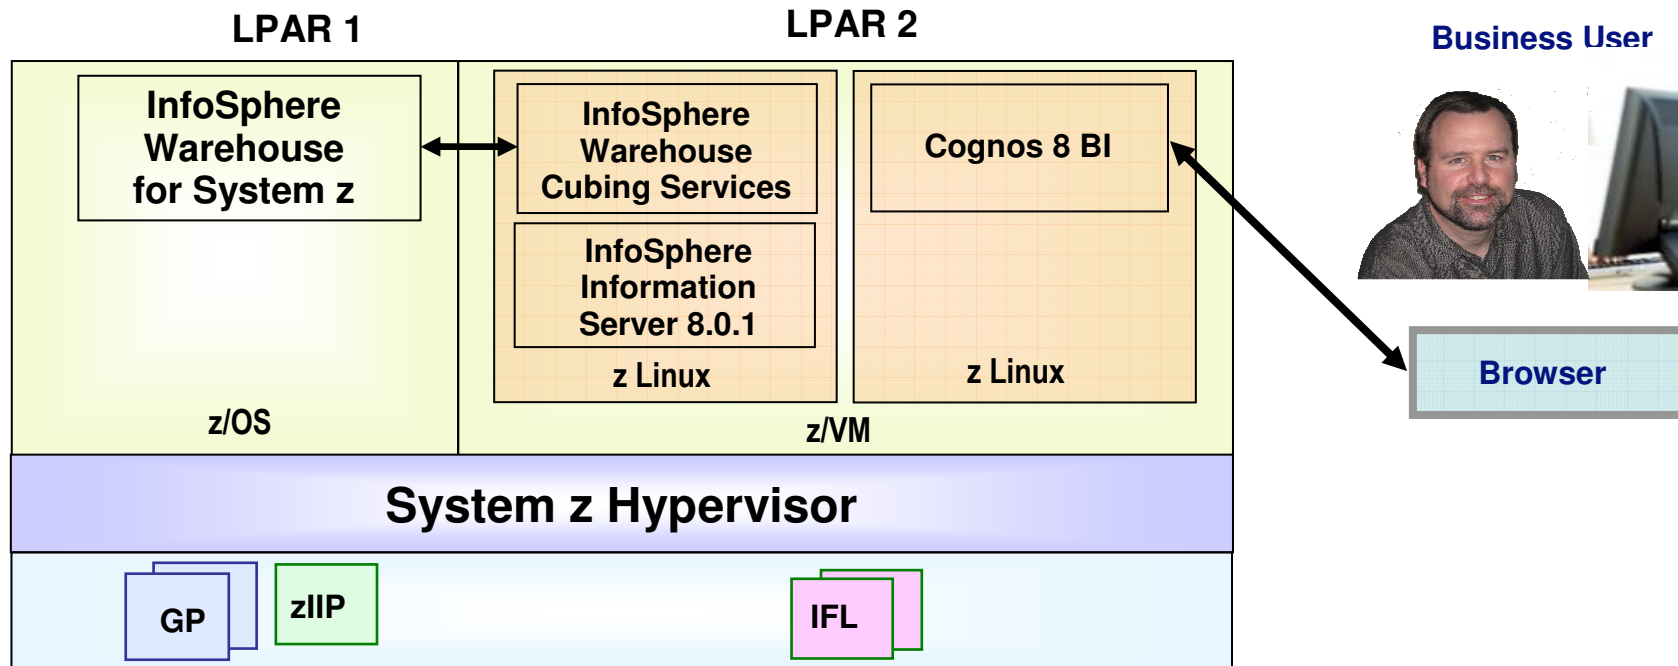


350,299*
Performance Units

3 year cost of acquisition
\$24.98M

*Production Performance Units required = (2,184+1840) MIPS x 87 = 350,088

System z Provides A Comprehensive BI Solution



System z Offerings for Enterprise Data Warehouse and BI:

- InfoSphere Warehouse for System z
- InfoSphere Information Server for System z
- IBM Cognos 8 BI for System z
- Only IBM can provide an end to end Platform DW and BI Solution