Databases:
Yesterday, Today, and Tomorrow

Dr. Pat Selinger
IBM Fellow, retired

Yesterday, Before Databases:

- **Every Application did <u>Everything</u>**

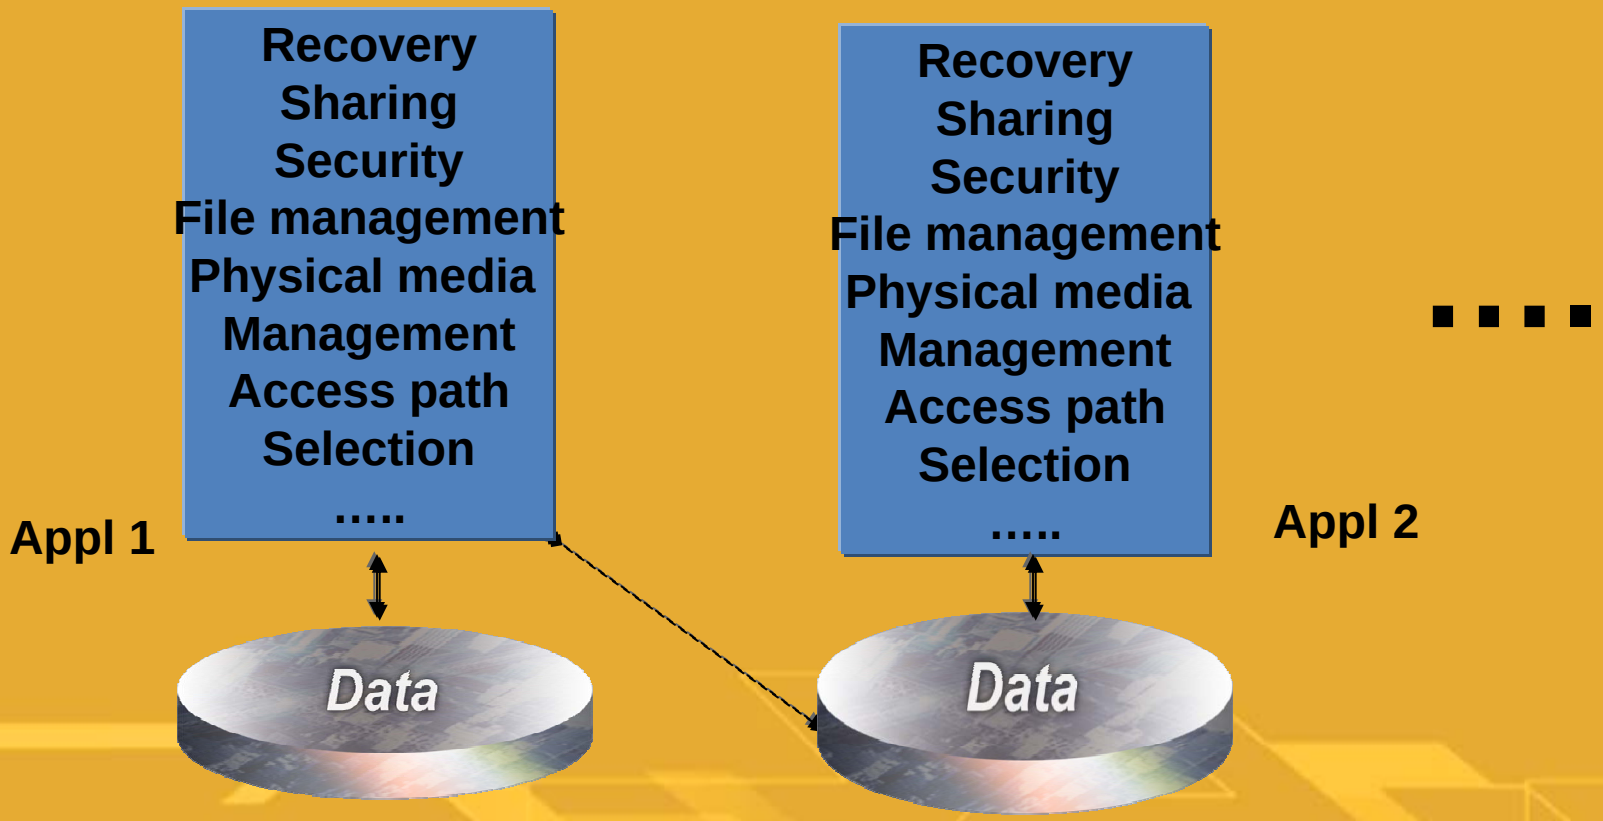

- **<u>Many</u> Copies of Data**

Dr. Ted Codd invented the Relational Model

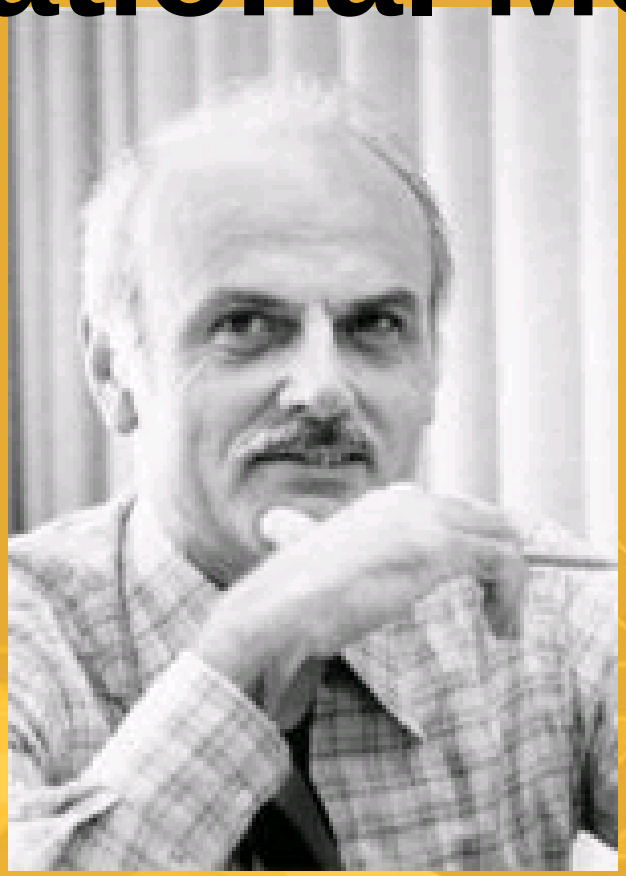

Dr. Don Chamberlin invented SQL

Query Optimization
Dr. Pat Selinger

1975

2005

Dr. Jim Gray

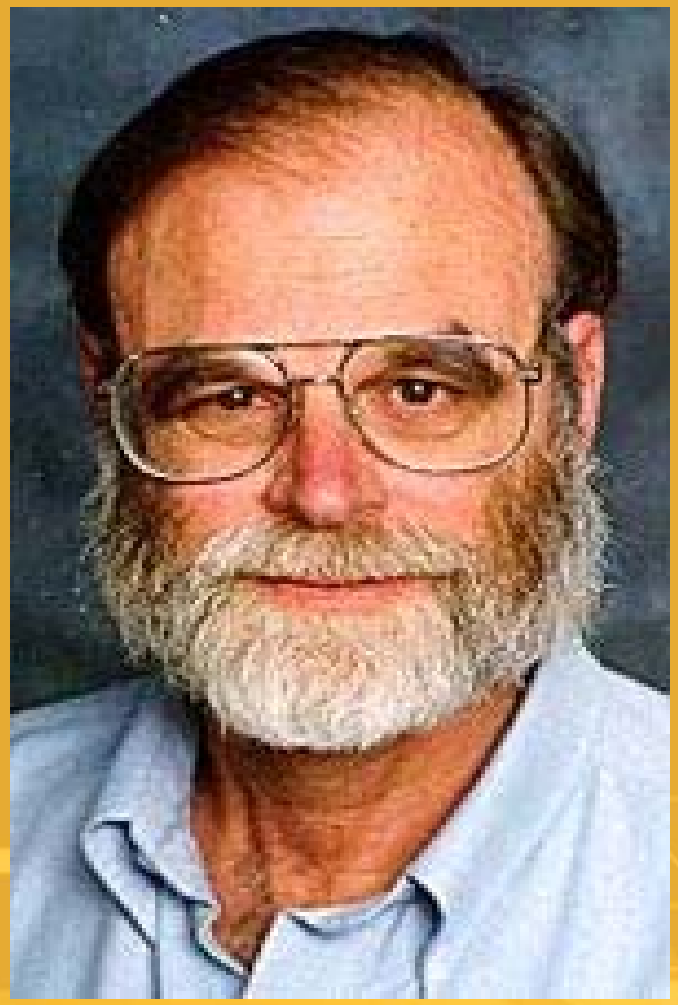

- Leader in System R and DB2 locking and transaction processing
- Missing at sea Jan. 28, 2007

Relational Databases

- Consolidated the Work

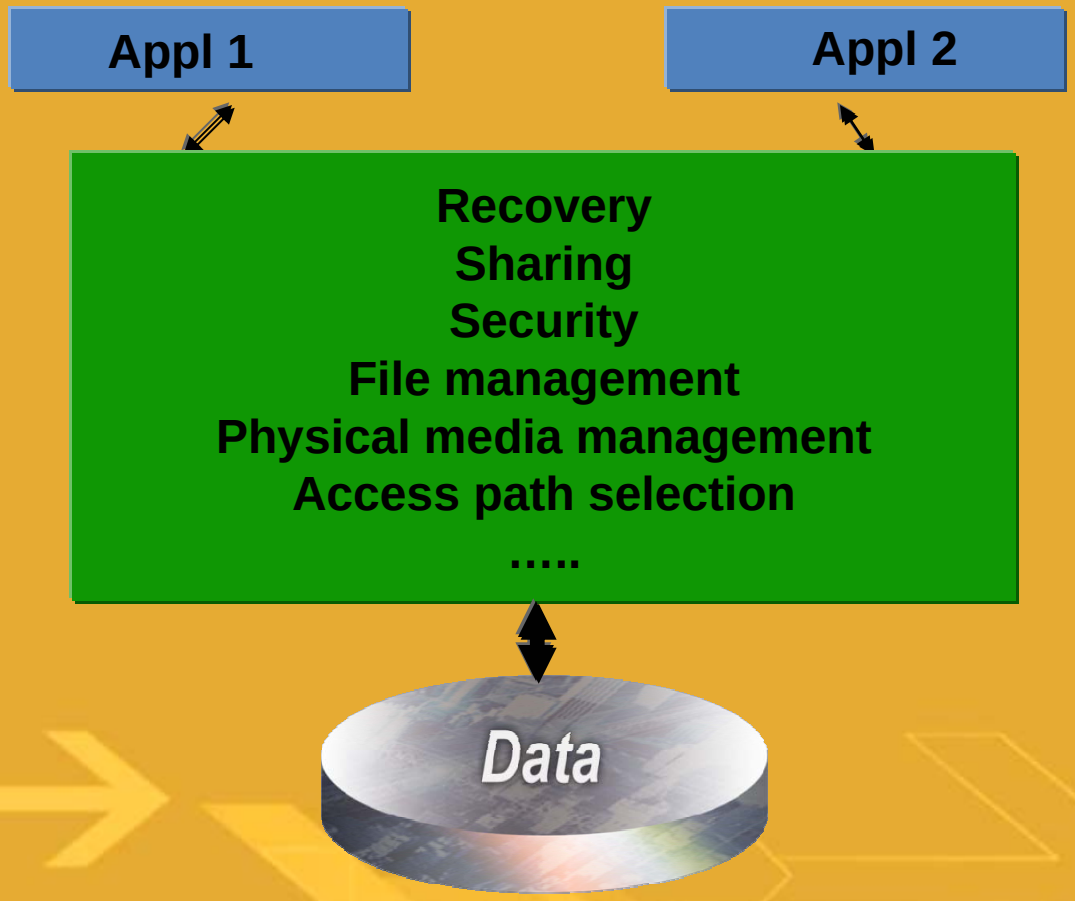

- Shared Single Copy of Data

We Shipped DB2

Don Haderle

The “Mother” of DB2

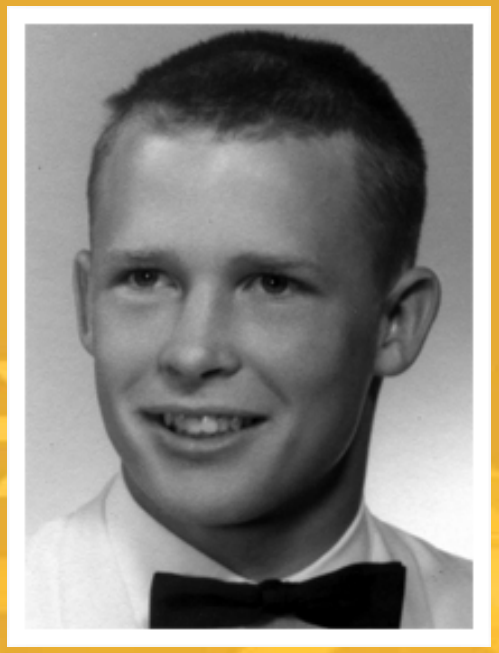

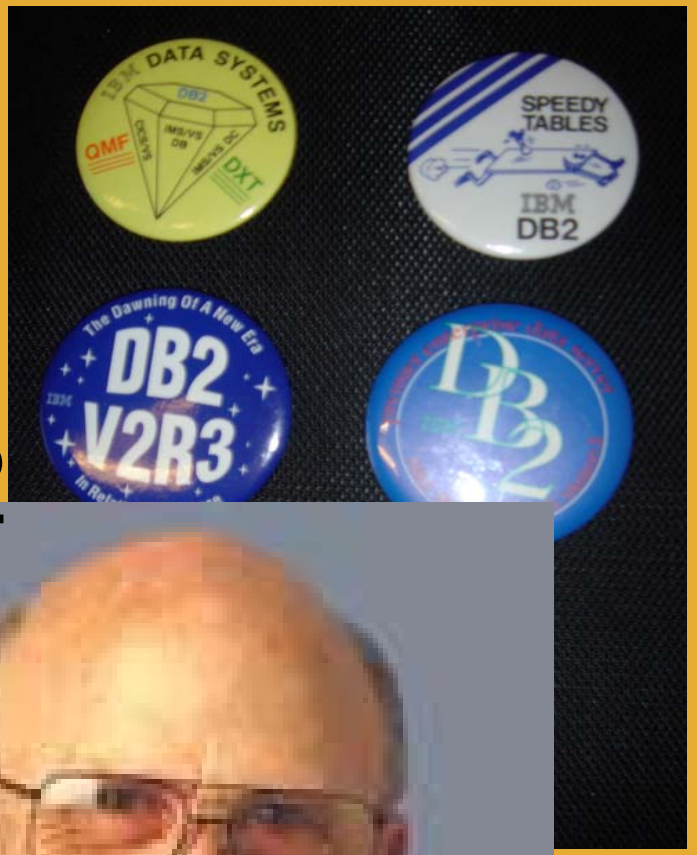

Roger Miller became “Mr. V-next”

Today

24 Years later…

DB2 Means More Than Ever Before…

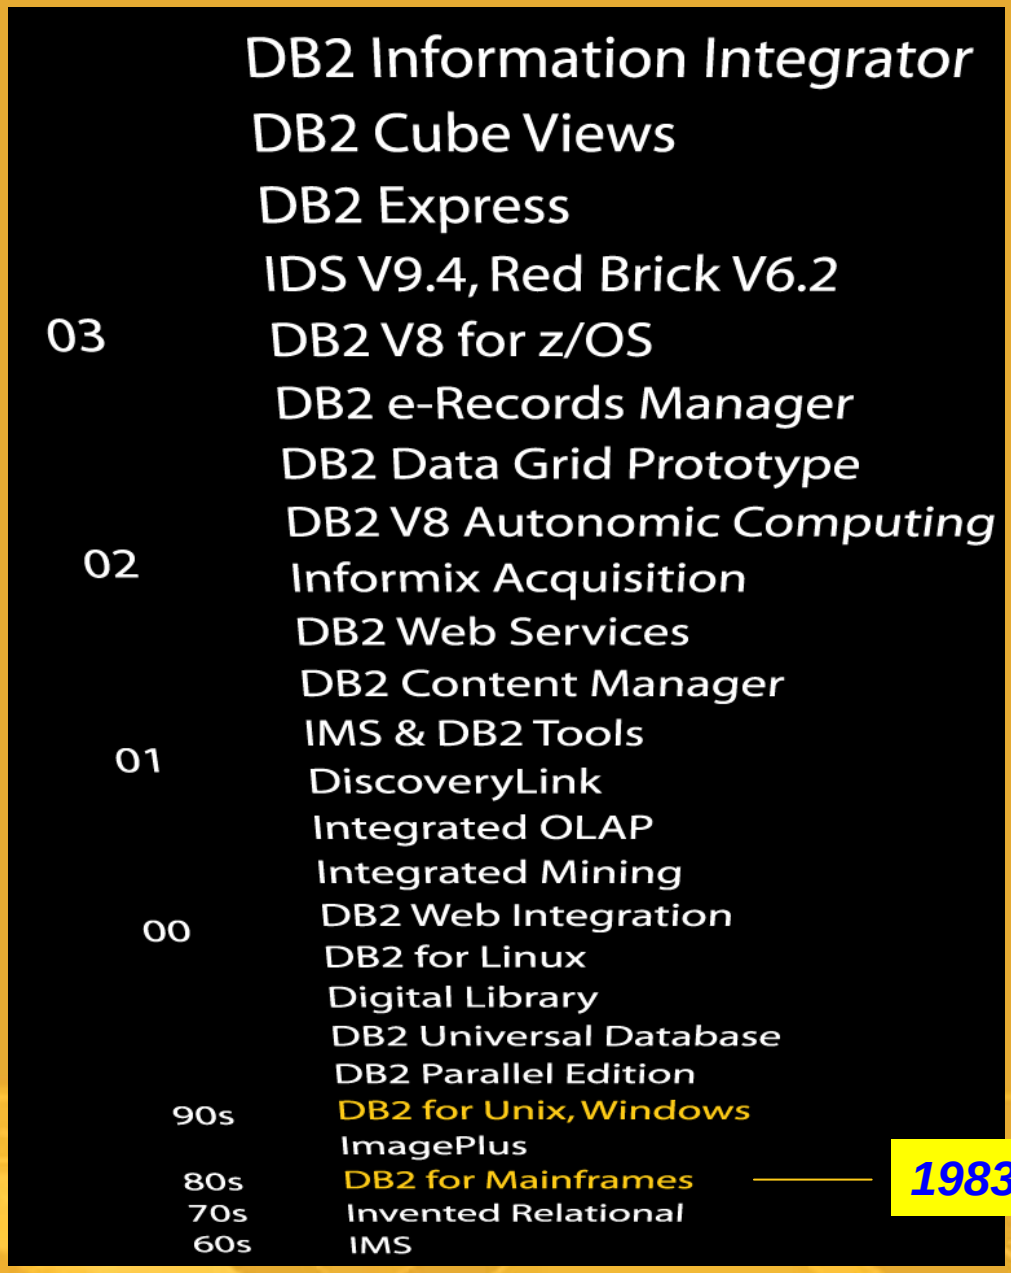

Our Business is Growing

IBM Information Management Software

#1 Market Share

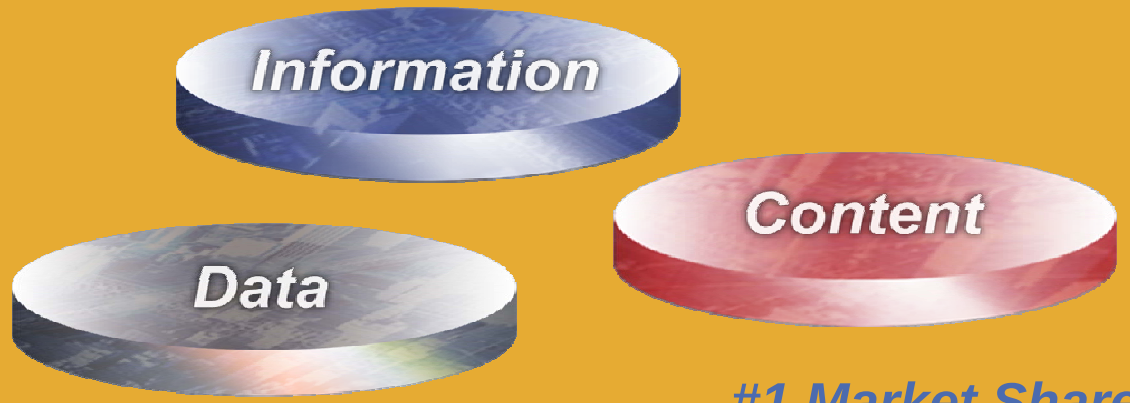

#1 Market Share

#2 Market Share
Growing Faster Than Market

Over 4,500 New Clients
in the Last 12 Months

Databases: Yesterday, Today, and Tomorrow

Data Servers @ IBM: Cool new Stuff

IBM DB2 System z

Technical Priorities

- ***Unmatched availability***
- ***Unmatched scalability***
- ***Unmatched reliability***
- ***Rock solid security***
- ***Support for all applications***
  - ***Legacy and new***
- ***Efficient data warehousing***
- ***Low TCO***

TD Bank Financial Group

10 YEARS

Continuous Availability

Land Registry

23TB

World's Largest Known OLTP Database

IBM DB2 System z

DB2 9 for System z Available: NOW

Availability, Scale & Resiliency

More Online Schema Changes

Volume Level Backup & Recovery

Business Flexibility

Faster, Cheaper, Granular Recovery

Bank of China

IBM DB2 System z

FNS BANCS Application on IBM Mainframe Redefines Transaction Processing Performance in Banking Industry

SYDNEY, AUSTRALIA and ARMONK, NY--(MARKET WIRE)--Feb 8, 2007

- ***380 million*** *client accounts*
- ***34 million*** *transactions per hour*
- ***9,445*** *transactions per second*
  - *Original goal:* ***4,100 TPS***

IBM DB2 System z

Specialty Engines Optimize Cost

- ***Significant value, available now***
  - Improved response times
  - Improved TCO
- ***Multiple workloads benefit***
  - SAP & other DRDA workloads
  - SAP BW & other warehouses
  - DB2 Utilities
  - DB2 9 remote native SQL stored procedures
- ***Strong customer adoption***

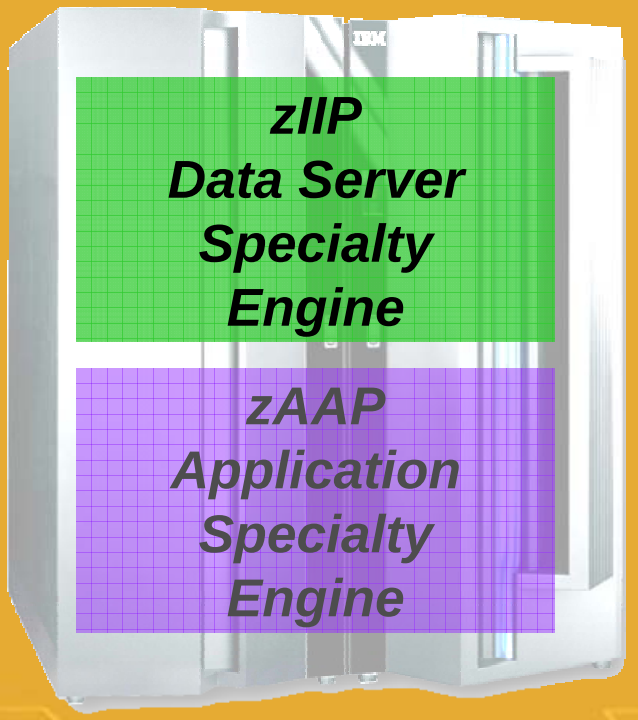

IBM DB2 Linux, Unix, Windows

Directions

- ***High Performance & Scalability***
- ***Reduced TCO & Accelerated Time-to-value***
- ***Resiliency through High Availability & Security***

IBM DB2 Linux, Unix, Windows

What's Coming in Performance Leadership?

- ***Go Faster…***
  - Faster engine
  - Faster utilities
    - Relational and XML
  - Consistent query performance

- ***Scale Further…***
  - Faster redistribute
  - Manage workloads
  - Expanded options for IBM Balanced Warehouse

"With DB2 9, we've seen a five-to-10 times performance improvement over a non-database environment."

CEO, Skytide

IBM DB2 Linux, Unix, Windows

What's Coming in Hands-off Administration?

- ***Prior Releases: Management-by-exception***
  - Notifies on unhealthy conditions
  - Advisors
- ***V9: Up-and-running***
  - Install, Provide e-mail / pager, "db2start"
  - Small to mid-size businesses
- ***Future: Continue to automate…***
  - Simplify upgrades
  - More automated storage
  - Integration with flash copy
  - Backup automation
  - "Just in Time" statistics
  - Auto-create compression dictionary

Databases: Yesterday, Today, and Tomorrow

what do you need to know for the future???

Pat's Observation #1:
Nature of "Interesting" Data is Expanding

Classic Data

Employee

Name	Dept Number	Employee ID	Ma
Jane	2	1000	-
John	3	1001	100

Department

Dept Name	Dept ID
Corporate	1
Manufacturing	2

Product Inventory
Sales Data
Bank Accounts
Warehouses
...

Unstructured Data

Multi-Modal Interactions,
e.g. speech, sensors

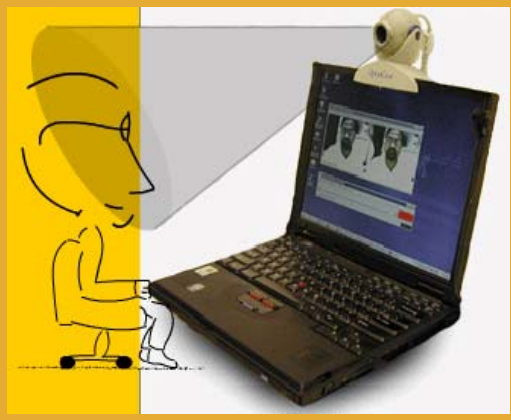

85% unstructured and not in DBMS

Increasing Need to Manage and Analyze New Types of Data

Changing Characteristics of Data

Transactions and structured data

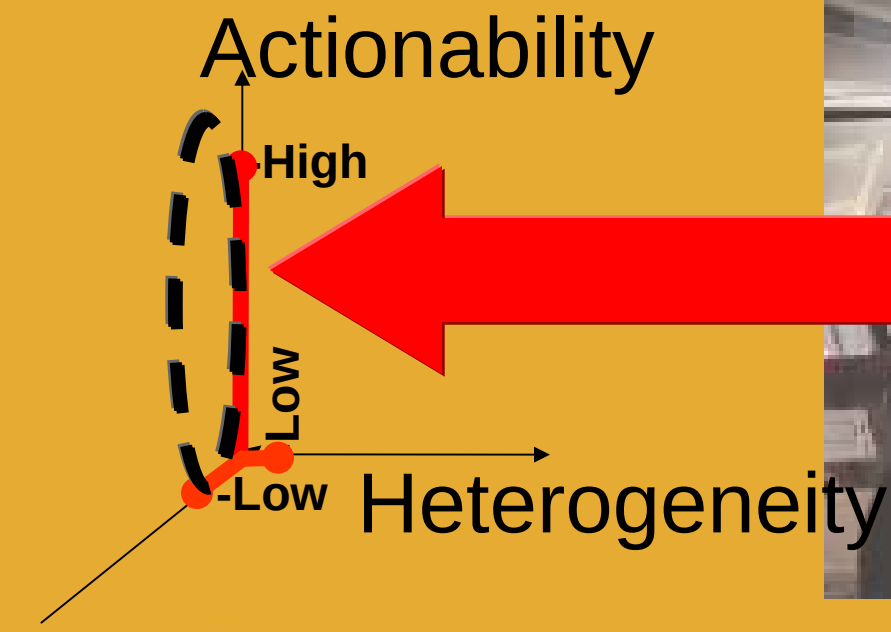

We are here!

Seat on an airplane:
easy to find,
structured data

Changing Characteristics of Data

Text and other human data

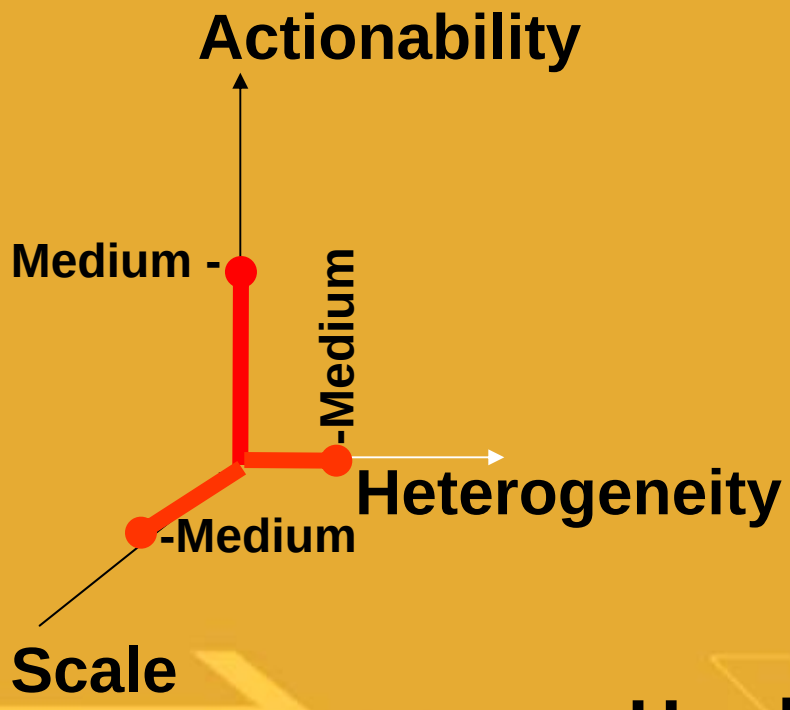

Harder to find but you know what you're looking for – Web Search

Changing Characteristics of Data

Machine-generated data

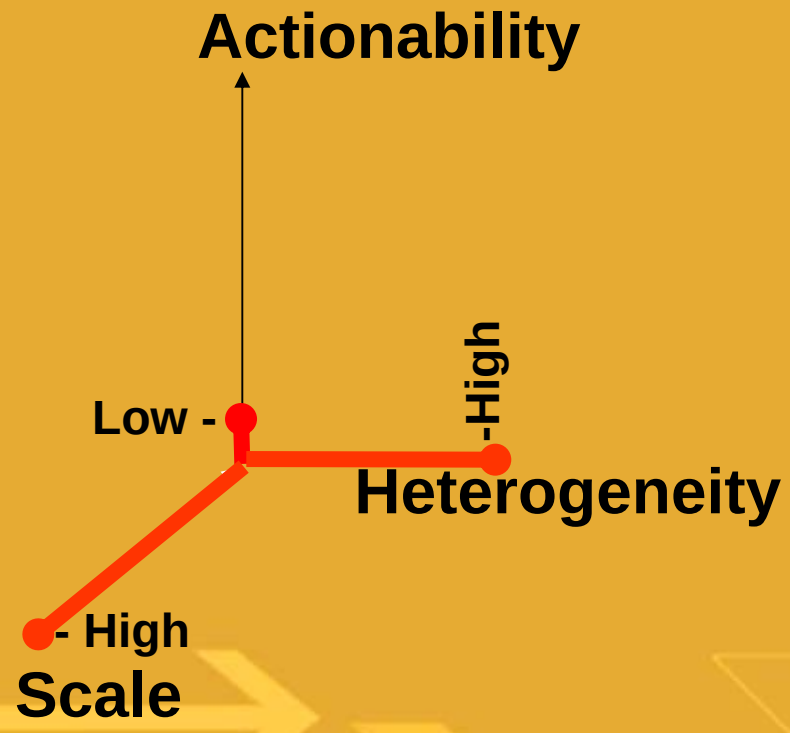

Mining for gold – maybe it's there????

Goal for the Future:
Make ALL Data Actionable

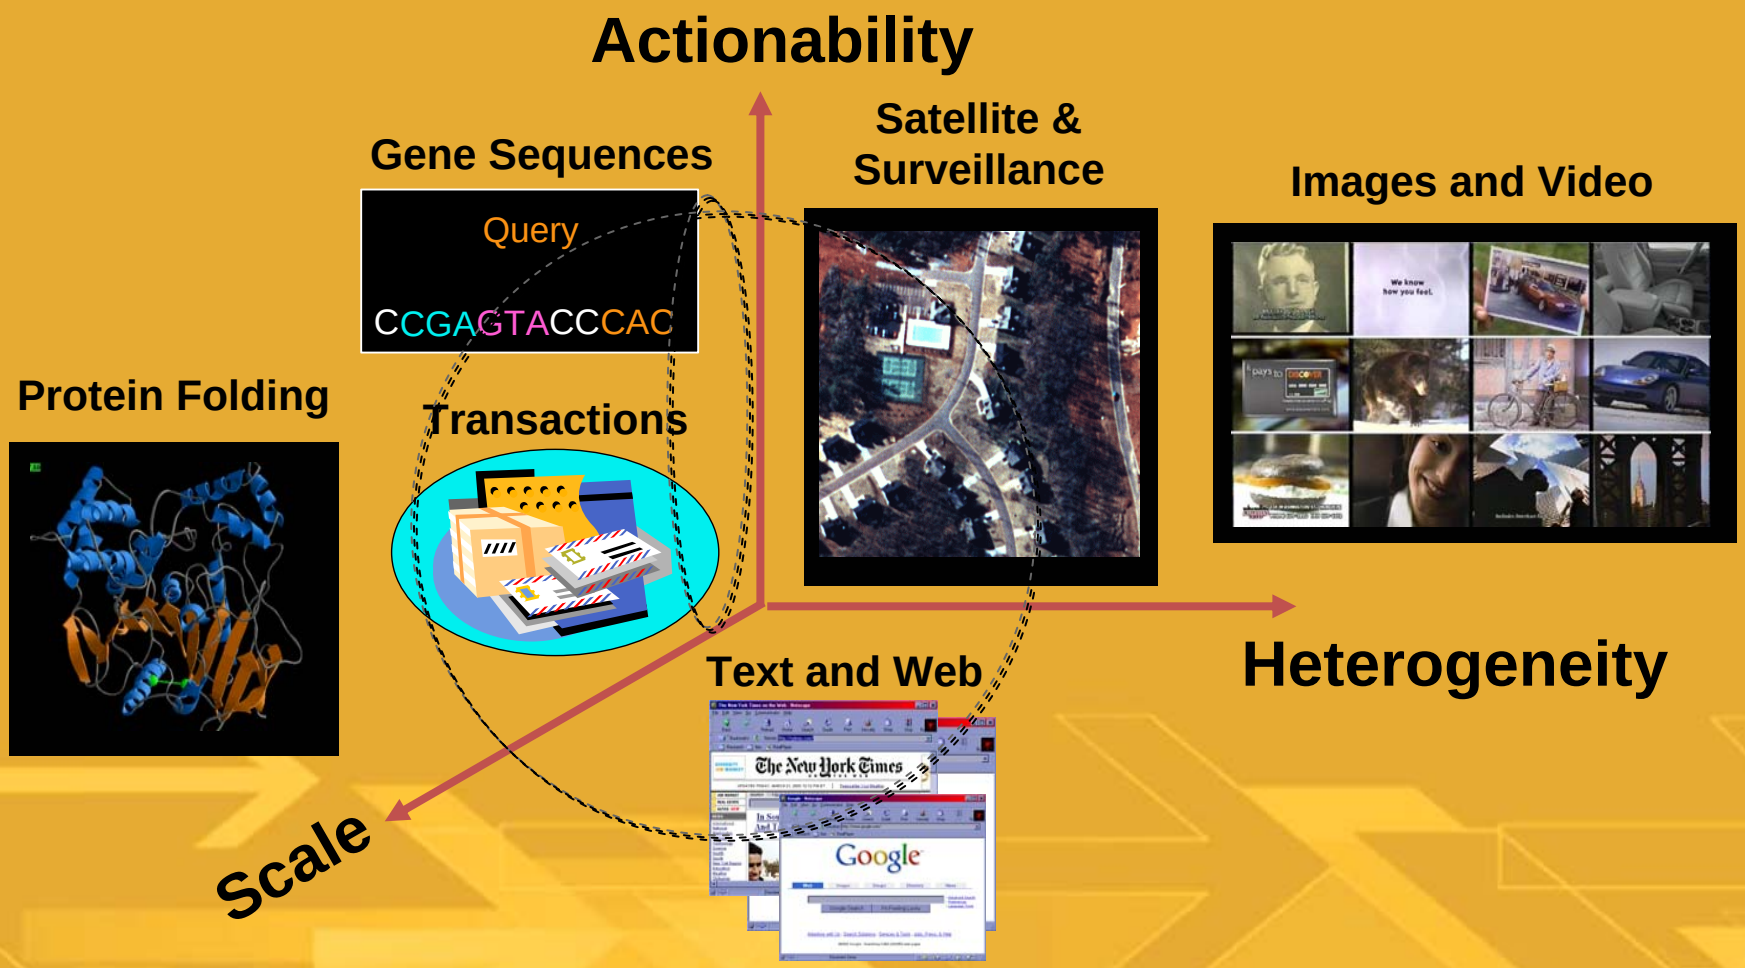

Content Management Solutions

Content Solutions

Information Integration

Workflow/Business Process Management/Collaboration

Archiving	**Document Management**	**Web Content Management**	**Output/Report Management**
Multimedia Management			**Imaging**
Digital Asset Management			**Content Integration**

IBM Content Management Portfolio

Regulatory Compliance / Records Management

Digital Rights Management

Learn about XML

Learn about Content Management

- ***…nagers provide …ata …t" services***
  - …al and referential in…
  - …nized query, Scalable st…
- ***D…gers want more "Content Management" se…es***
  - …ck-in, check-out and versioning
  - …grated hierarchical storage management
  - …-normal (i.e. hierarchical) metamodel
- ***XML will drive this change***
  - Sometimes it's data – other times it's content

IBM Data Server Innovation

Milestones in Innovation…Over 3,000 Patents

1968	1980	2006
First Hierarchical Data Server	***First IBM Relational Data Server***	***First Hybrid Data Server***

Continuous IBM innovation

IBM Data Server Innovation

Defining the Industry…. Setting the Standard

- System R
- Relational Database
- SQL
- R* D
- Obje
- Para
- DB2
- Rec
- Auto
- Auto
- Mate
- Parallel Database
- Automatic Summary Tables
- Materialized Query Tables

- Index Advisor
- Heterogeneous Database

l Clustering

er

er

on

er

- XQuery Standard
- pureXML
- Automatic Schema Mapping

pureXML

2006

Deep Compression

Pat's Observation #2: Nature of "Data Warehousing" is Changing

Old Warehouses

- Weekly batch update
- "Nice to have" – helps with marketing decisions
- Separate system

New Warehouses

- Continuous feed
- Mission critical – drive key business decisions
- Integrated and connected – query results drive transactional workload

Pat's Recommendation #2
Learn about Dynamic Warehousing

Every Person, Every Transaction, Every Asset

Information On Demand
Optimize Each Transaction
as It Happens

Dynamic Warehousing

OLAP

Dynamic Warehousing Attributes:

1. *Real-time access – in context*
2. *Analytics – part of a business process*
3. *Unstructured information – extracted knowledge*
4. *Extended database capabilities*

Query
Unde
Happened

g

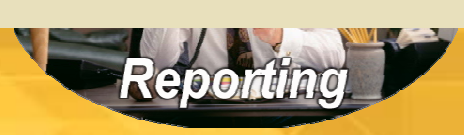

E.G….Manufacturing

Making a Decision…or a Good Decision?

S A L E S

Unified, Trusted, Insightful Information, In Real Time

Production

- Overtime Shifts
- Raw Materials @ Premium
- MRO Impacts

Supply Chain

- Ripple Effect Strains Suppliers & Distributors
- Cost & Relationship Impact

Clients

- Lost Faith, Lost Orders
- Lost Clients

Finance

- Costs Increases
- Revenue Losses

Dynamic Warehousing -- What do you need to learn?

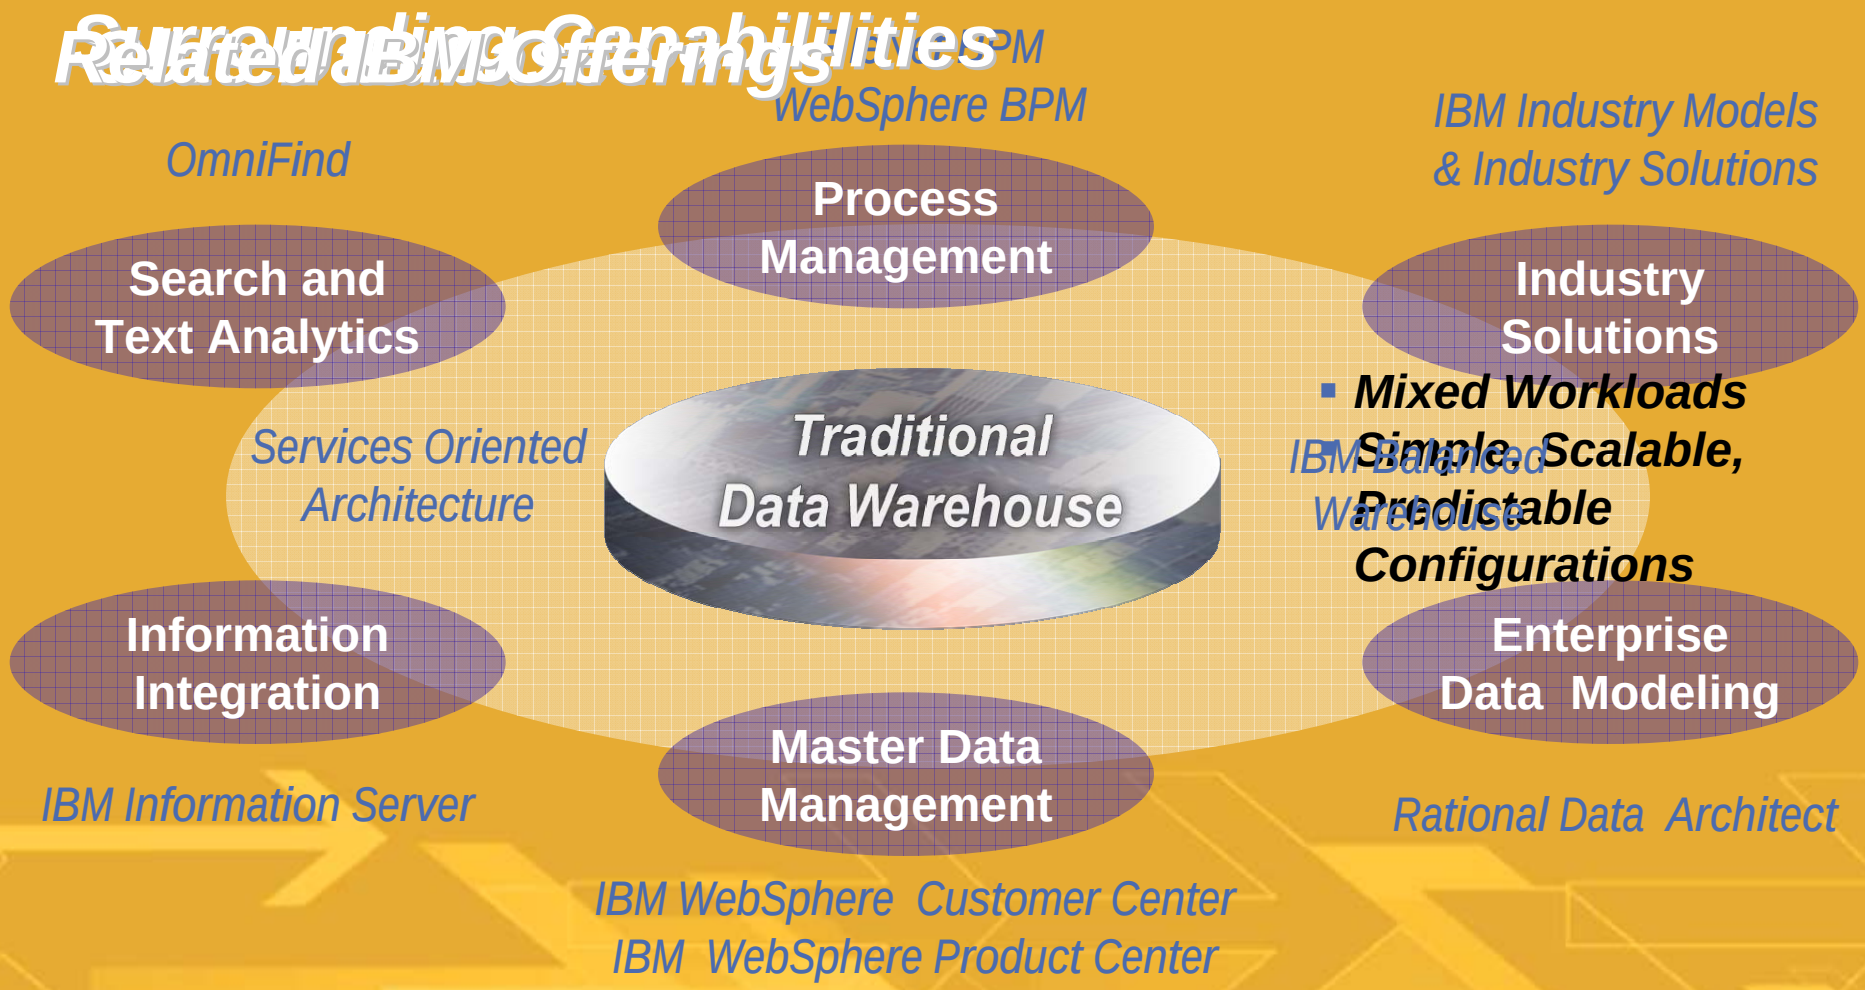

IBM Information On Demand

Tying it all together

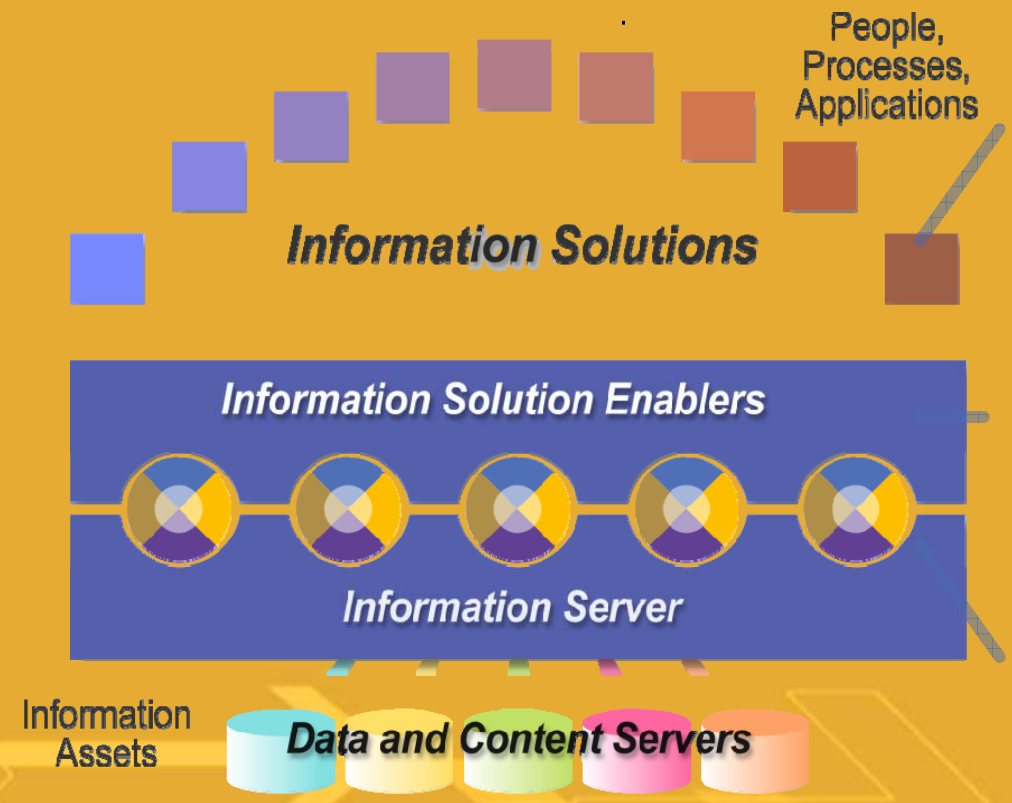

Information Solutions
Create new business value
and innovation
from information

Master Data Management &
Industry Models
Define and synchronize critical
information for operations and
business insight

Information Server
Understand, Cleanse,
Transform, Deliver;
Metadata Driven

24 Years…

DB2 Means More Than Ever Before…
And There's Plenty More to Come

Databases:

Yesterday, Today, and Tomorrow

Thank You!

Databases:

Yesterday, Today,

and Tomorrow

Thank You!