

# Database Growth: Problems & Solutions

Thomas Boyd  
Hewlett-Packard



© 2006 Hewlett-Packard Development Company, L.P.  
The information contained herein is subject to change without notice


# Agenda

- The Data Explosion
- Effects of Database Explosion
  - End User
  - IT Operations
  - Lab Tests
  - Real-world experiences
- Solutions
- Q & A



# The Data Explosion

- “Data Explosion”: 166,000 Google Hits
- “Information Life Cycle Management”: 1.3M Google Hits
- Numerous analyst papers
- Archiving vendors



# Data Explosion: Largest OLTP on Unix Databases



- Source: Winter Corporation; Top 10 Survey ([www.wintercorp.com](http://www.wintercorp.com))
- 2003: Max: 5.4 TB, Average 2.5 TB
- 2005: Max: 16.4 TB, Average 6.8 TB

**Max: 200% Increase**

**Average: 175% Increase**

# Drivers for Data Explosion

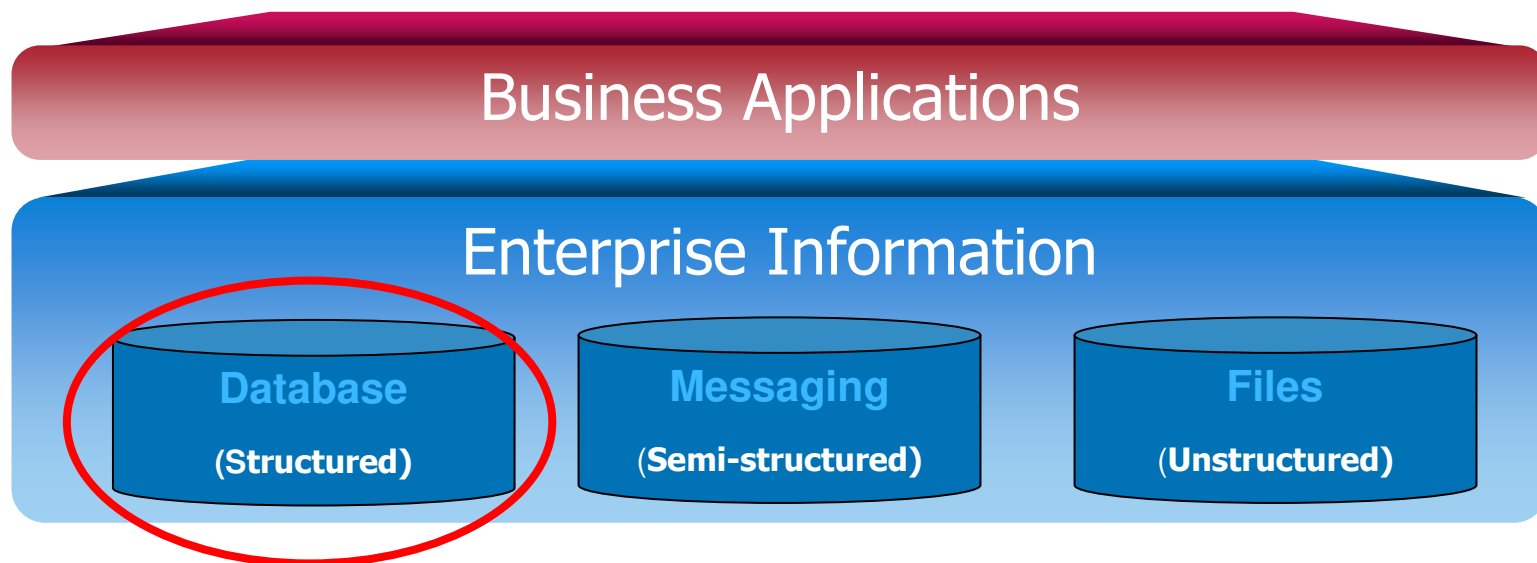
- Increased regulatory retention requirements
- Increased business intelligence and trend analysis
- Increased detail with-in business transactions
- Consolidation of systems
- Underlying growth in business volume

# Data Explosion: Extremes

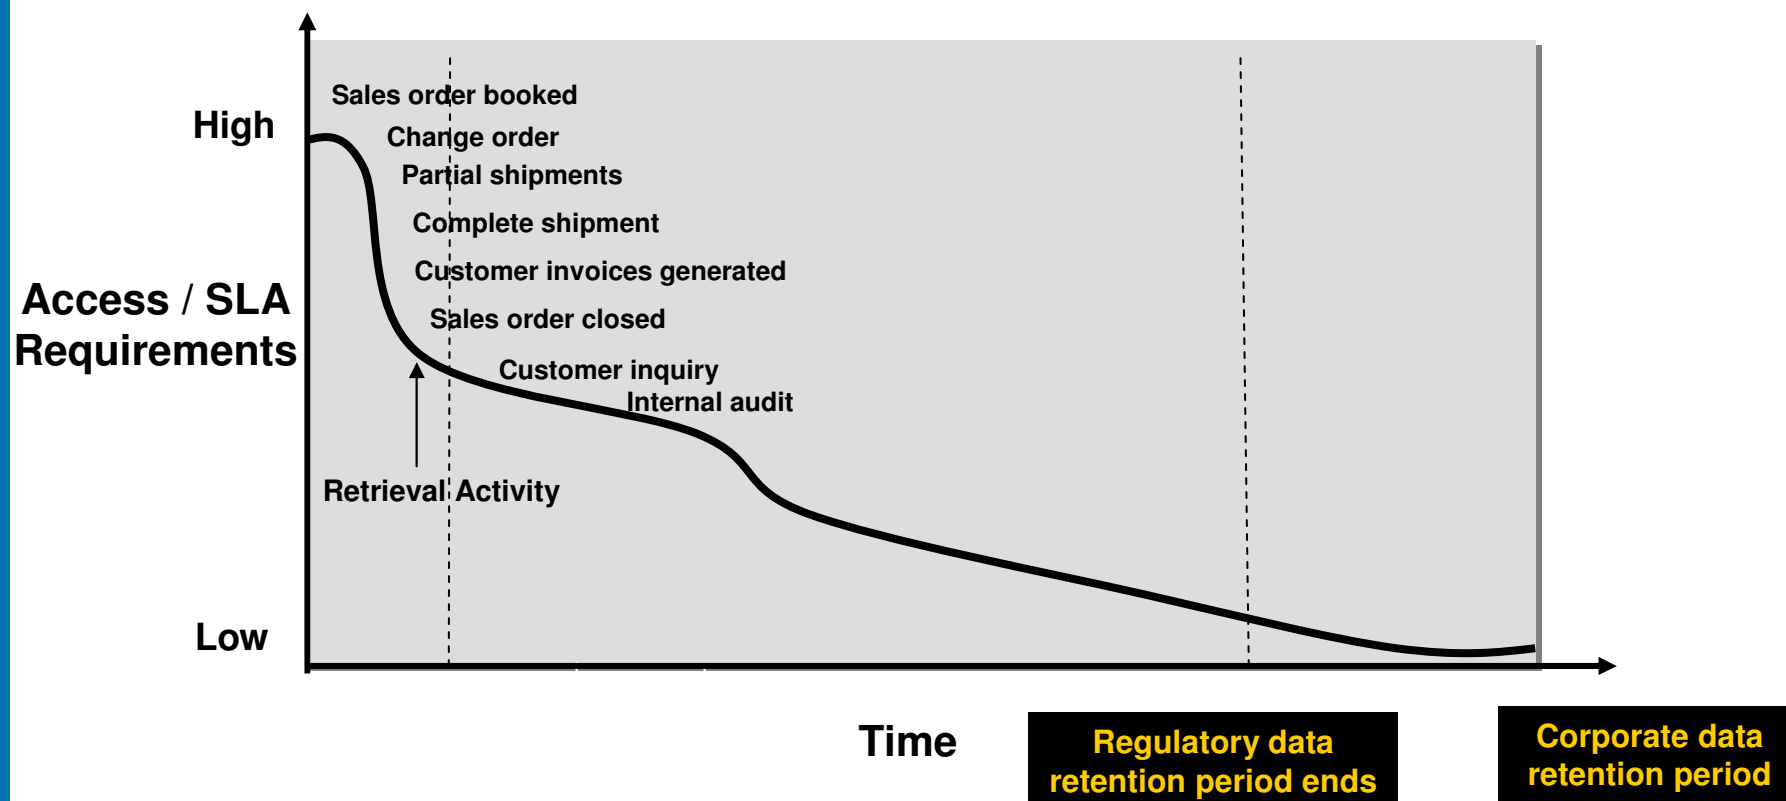
- India: 6.6 new mobile subscribers for month of October, 2006
- USA: 25-35% growth in 2006 holiday online retail sales, industry-wide
- Europe: BASEL II
- Hewlett-Packard: Consolidating 85 data centers into 6

# Data Explosion: Types of Data

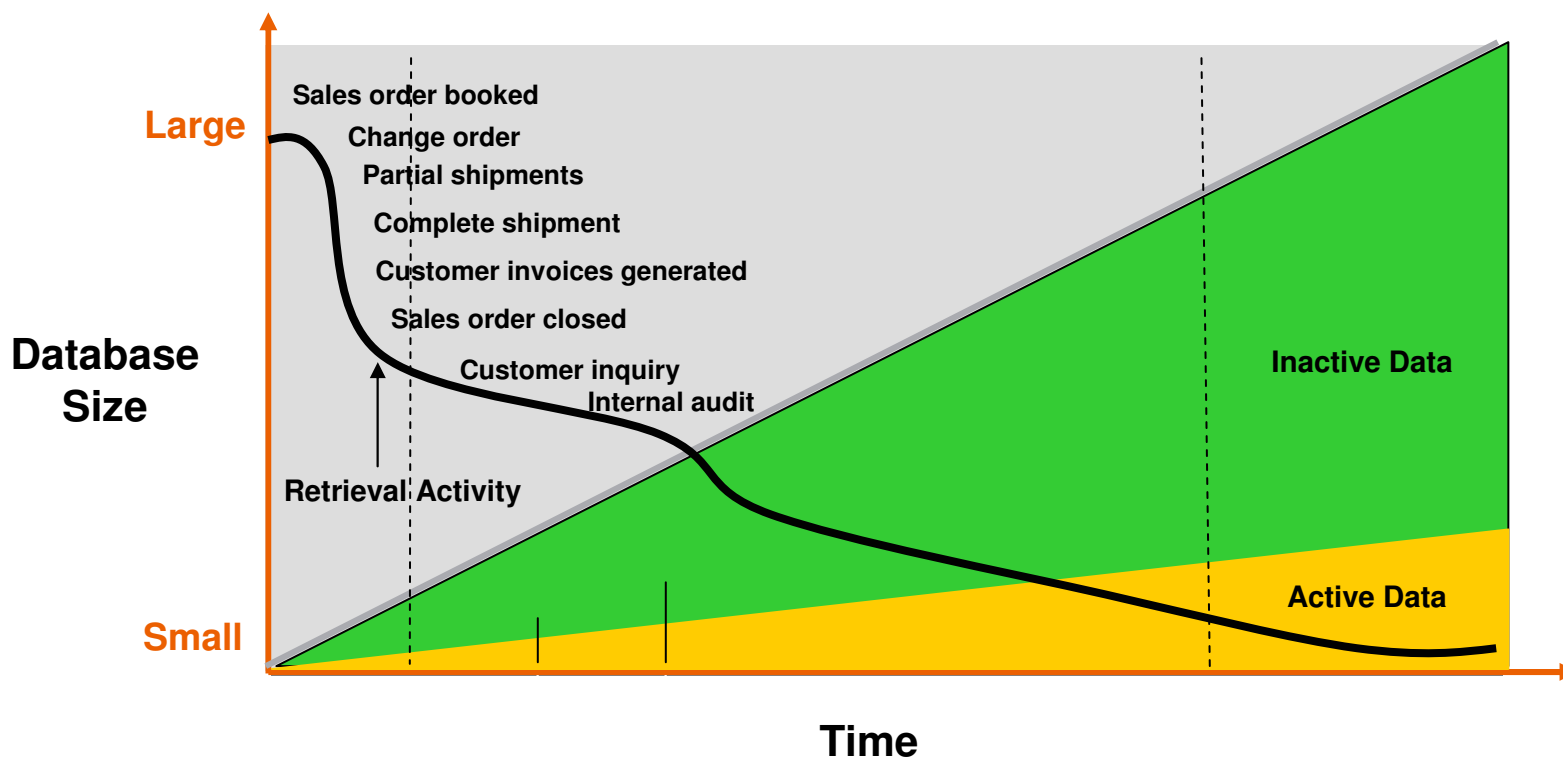
- Structured, semi-structured, and unstructured



# Data has a Lifecycle




# Inactive Data



# Agenda

- The Data Explosion
- Effects of Database Explosion
  - End User
  - IT Operations
  - Lab Tests
  - Real-world experiences
- Solutions
- Q & A



# Effects of Database Growth

- End-User
  - Declining application performance
  - Decreased application availability
  - Information overload
    - Is your application presenting irrelevant information to end-user, causing extra work?

# Effects of Database Growth

- IT Operations
  - Storage requirements
  - Server requirements
  - Backup and recovery windows
  - Maintenance and upgrades windows
  - Cloning operations
  - Database and Application Tuning

# Lab Tests

- Single Table Query, non-unique index range scan (good selectivity)

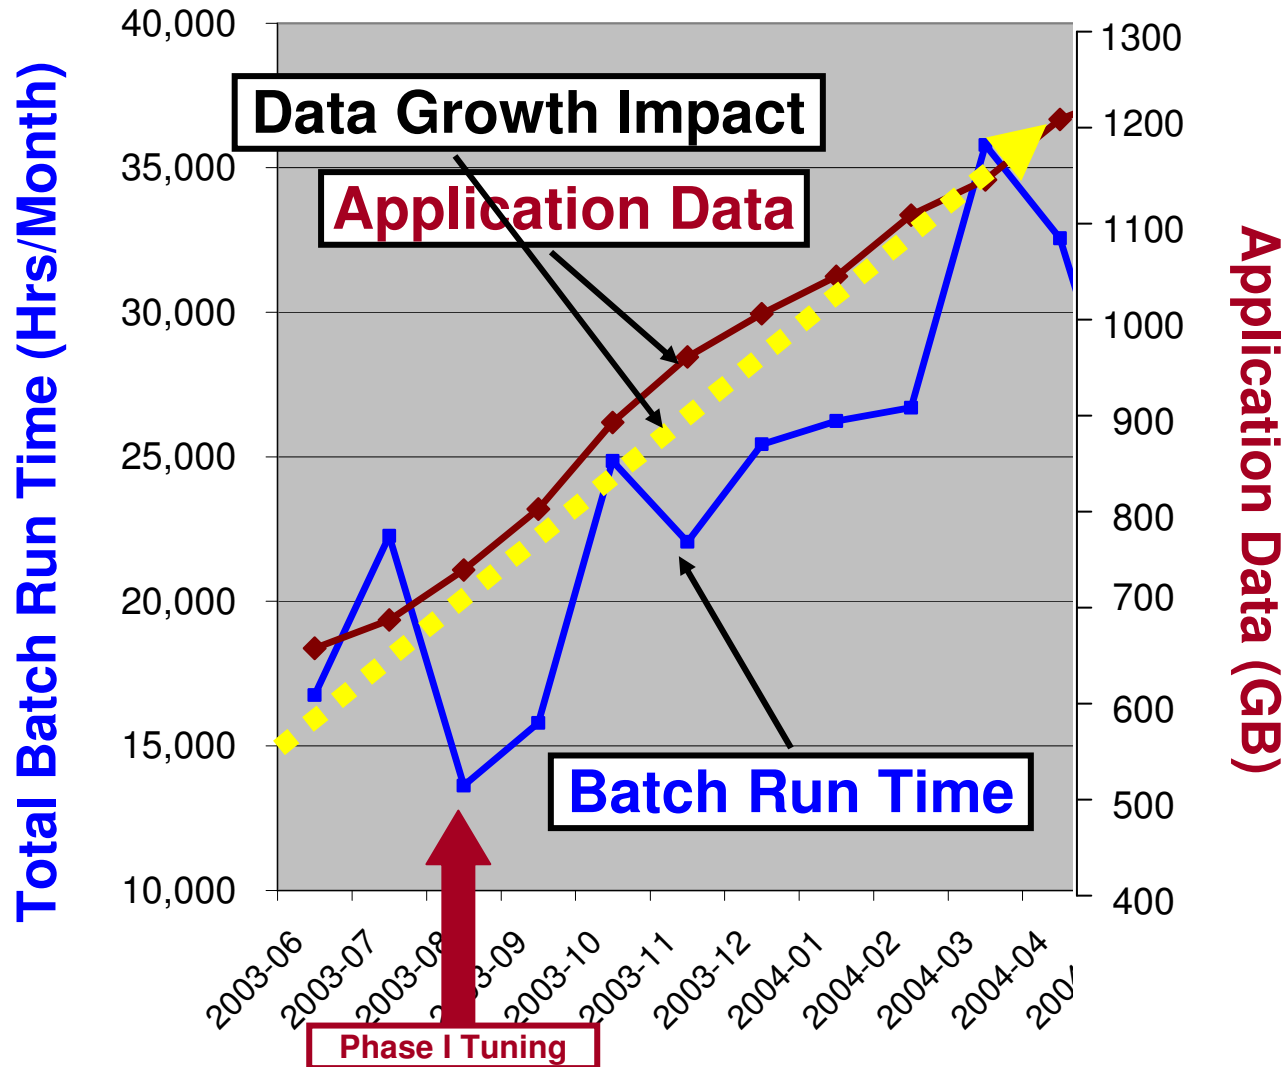
Table Rows	Consistent Gets
200,000	?
400,000	?
2,000,000	?

# Lab Tests: Future

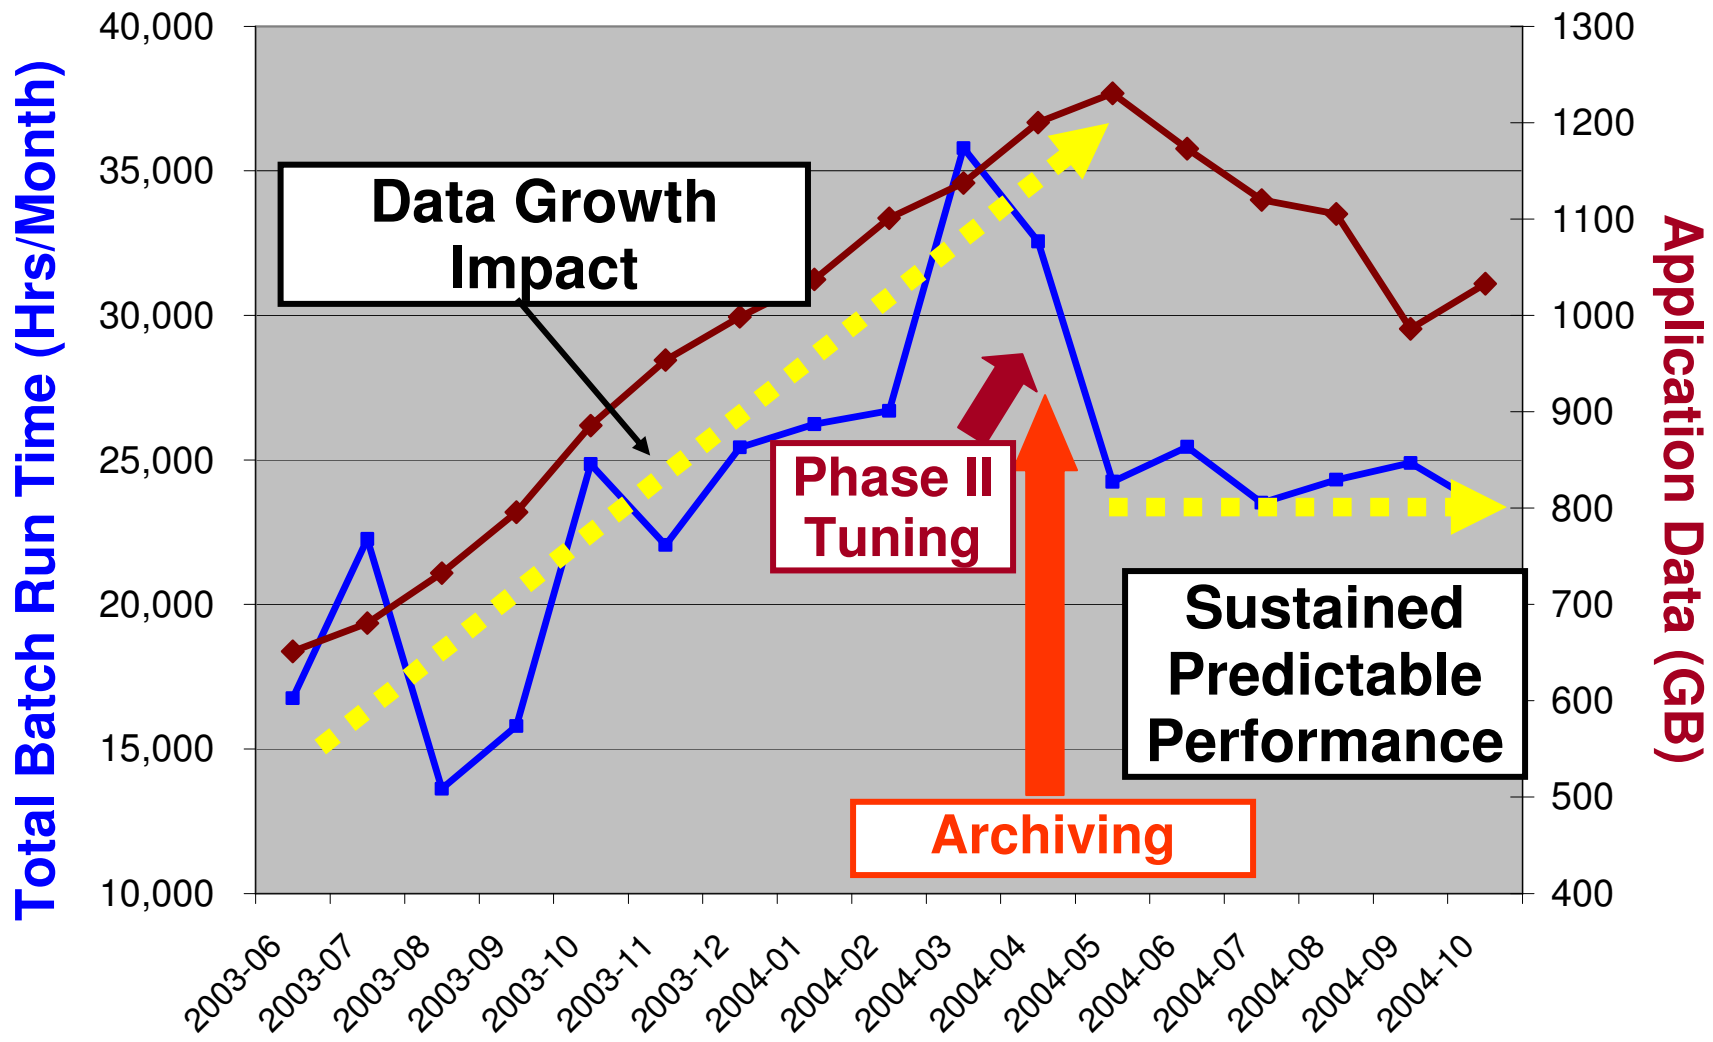
- Single Statement Testing
  - Track all relevant statistics
  - Insert and delete statements
  - Complex joins
- Full scale stress testing
  - Simulate Oracle EBS load
  - Varying machine and storage configurations
- Redo, Rollback analysis
- Backup (RMAN) analysis



# Real World: Data Growth vs. Performance



# Real World: Data Growth versus Performance



# Agenda

- The Data Explosion
- Effects of Database Explosion
  - End User
  - IT Operations
  - Lab Tests
  - Real-world experiences
- Solutions
- Q & A



# Solutions: Control Data Generation

- Look for setups or configurations that are contributing to data growth
  - Persistent failures in interfaces
  - Persistent failures/notifications from Workflow
  - Excess use of Oracle Alerts
  - Inappropriate inventory tracking levels
- **Monitor**

# Solutions: Purging

- Some data is not subject to regulatory requirements and has no ongoing business value
- Candidates:
  - Concurrent request data
  - MRP runs
  - Workflows
  - Temporary and Interface Tables
  - Others

# Solutions: Hardware Upgrades

- Storage
- Server
  - Scale up
  - Scale out (RAC)
- Network



## Solutions: No Consolidation

- Sometimes managing two medium sized databases is easier than managing one large database
- Gives up on all the advantages of consolidation
  - IT Operations
  - Consolidating reporting, etc...

# Solutions: Partitioning

- Powerful Oracle feature for managing large tables
- Improves performance **if** queries are able to use partition keys
- “Life cycle” of data may leave active data in same partition as inactive data

# Solutions: Archiving

- Remove data from OLTP Databases while retaining ability to access
  - Database Backups/Snapshots
  - Exports, CSV files
  - XML Files or other self-describing format
  - Relocation to secondary database
    - Access via reporting tools or business analytics
    - Access via native application

# Agenda

- The Data Explosion
- Effects of Database Explosion
  - End User
  - IT Operations
  - Lab Tests
  - Real-world experiences
- Solutions
- Q & A



# Conclusion

- Plan for it!
- Understand its effects
- COLLABORATE 07: Updated Presentation and White Paper
  - Full lab results

