The Future of Database and Data Warehousing Software
Data Management Meets the Future

DB2 v10 AND InfoSphere Warehouse v10

Luis Reina – IBM Sofware Group

DB2 10 and InfoSphere Warehouse 10

Improved Cost Efficiencies

Higher Performance

Increased Productivity

Improved Cost Efficiencies

- Adaptive Compression
- Multi-Temperature Data Management
- SQL Compatibility

Adaptive Compression

"We have already compressed our data, how can this help ***us****?* "

Adaptive Compression

DB2 9.7

- Uses a single, static compression dictionary
- Compresses data based on recurring patterns that appear in the table
- A classic table re-organization is necessary to improve compression ratios if a significant number of records in a table have been updated, or if a large amount of new data has been inserted

DB2 10

- Multiple page-level dictionaries in addition to a single table-level dictionary
- Compression dictionary contains locally frequent patterns, with one dictionary stored on every page
- When a page becomes full, page-level compression is applied, immediately freeing up more storage in that page
- Reduced need for table reorgs

NEW! Adaptive Compression Shrinks your Data Storage Needs

Lower Storage Costs; Lower Administration Costs

- **Higher performance**
  - Faster queries for I/O-bound environments
  - Faster backups
- **Lower costs**
  - Postpone upcoming storage purchases
  - Lower ongoing storage needs
  - Easier administration with reduced need for table re-orgs

"Page-level dynamic compression is one of the new DB2 features that will reduce planned outages and increase storage savings by up to 2X over DB2 9.7%."
—Jessica Tatiana Flores Montiel, DAFROS Multiservicios

"Our migration from Oracle Database to DB2 resulted in a 40% storage savings. Upgrading to DB2 9.7 and index compression brought our average savings to 57%. Now adaptive compression brings our **average savings to 77%**, dramatic savings!"
—Andrew Juarez, Lead SAP Basis / DBA, Coca Cola Bottling Company.

Multi-Temperature Data Management

"How can we speed up queries without having to invest in more SSD storage? "

NEW! Multi-Temperature Data Management

Increase Ability to Meet SLAs; Postpone Hardware Upgrades

- Storage pools for different tiers of storage
  - For range partitions, policy-based automated movement of data

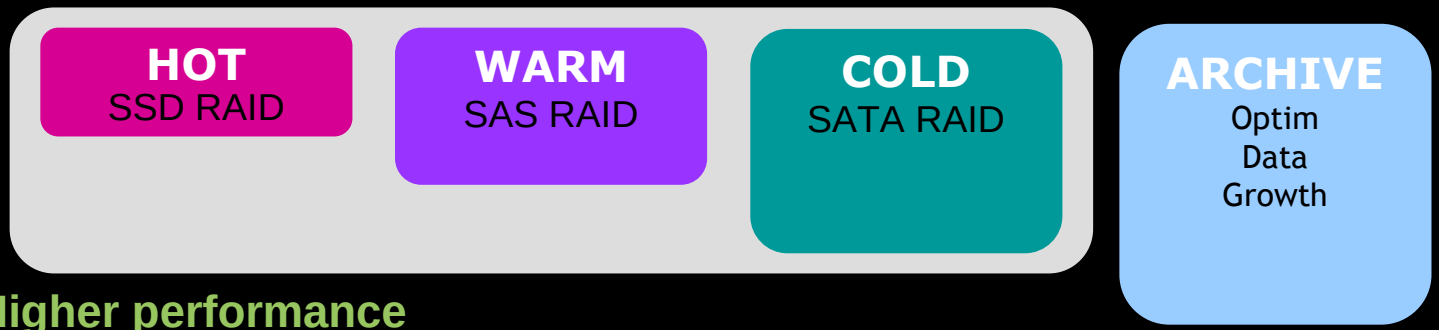

- **Higher performance**
  - Improved ability to meet SLAs while retaining greater amount of data for analysis
- **Lower costs**
  - Embrace new lower-cost storage technology
  - Further reduces the cost for meeting SLAs

ITGAIN

"The multi-temperature database management feature of DB2 V10.1 is great because the hardware world is not just RAM and hard disks. There are many types of storage options with different I/O speeds and prices. This feature allows administrators to make optimal use of these different devices, balancing expensive SSDs with cheaper SATA disks and everything in between. Using SSDs for indexes and logs and a SATA array for the data, we noticed fantastic improvements in I/O speeds, especially for synchronous reads. Additionally, the background movement of data between the storages groups is very fast." —Thomas Kalb, CEO ITGAIN GmbH

Multi-Temperature Data Management

Low Operational Costs

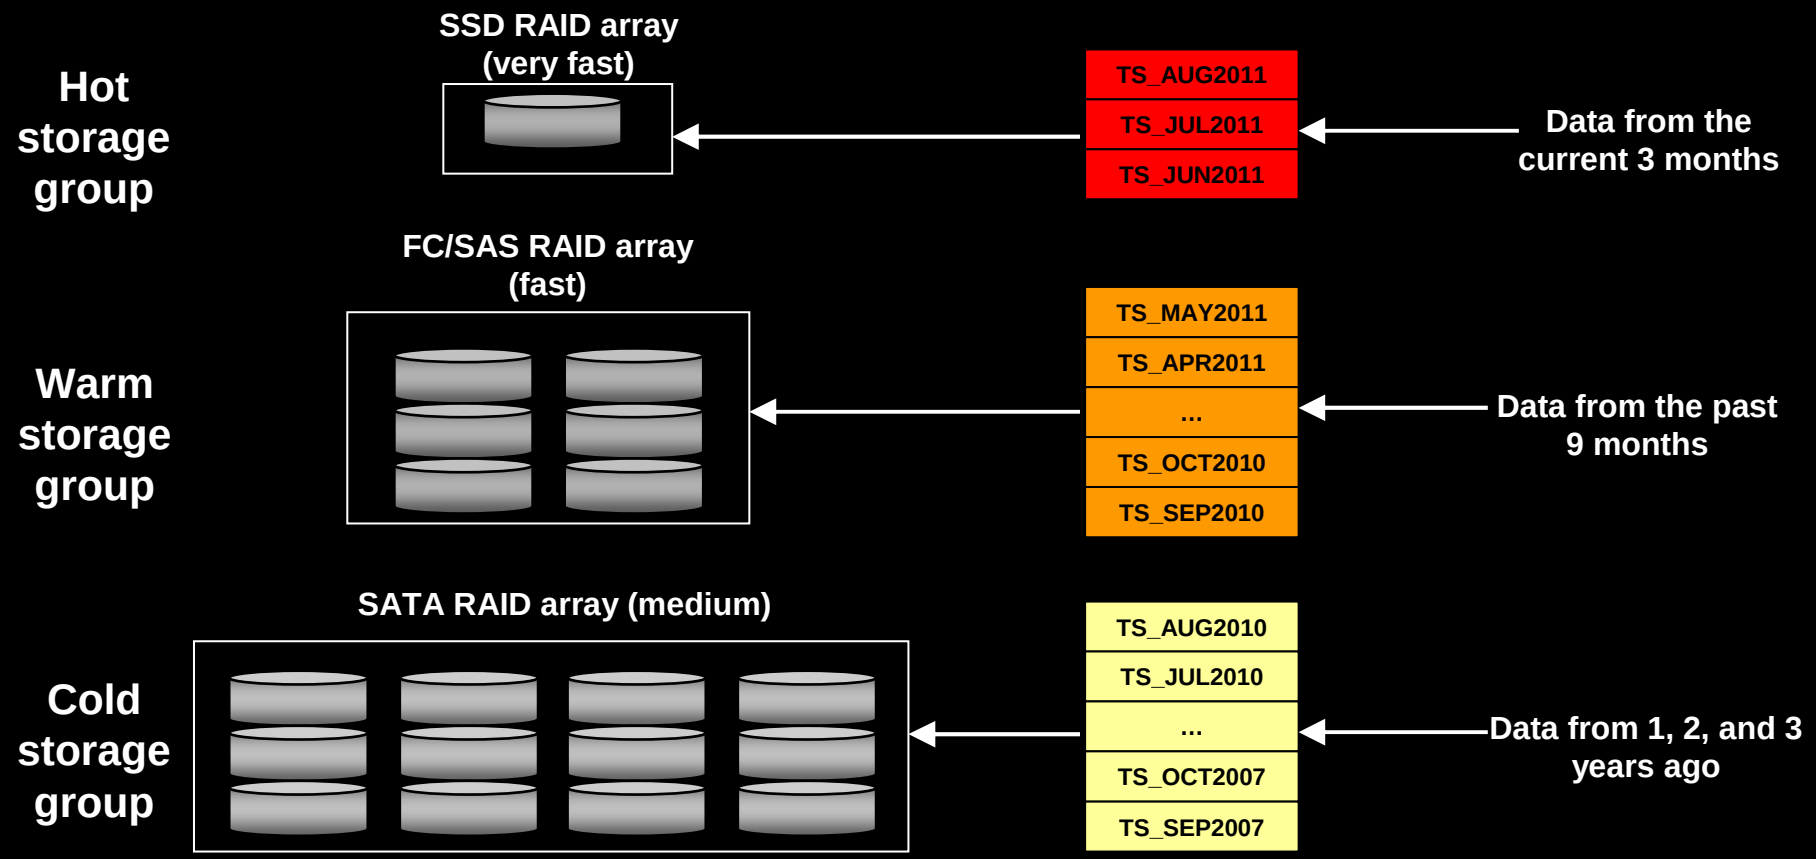

Multi-Temperature Data Management (cont.)

Moving Data Between Storage Tiers

Low Operational Costs

1 ALTER TABLE sales
ATTACH PARTITION TS_SEP2011
STARTING FROM ('9/1/2011') ENDING ('9/30/11')

2 ALTER TABLESPACE TBSP_JUN2011
USING STOGROUP SG_WARM

3 ALTER TABLESPACE TBSP_SEP2010
USING STOGROUP SG_COLD

4 ALTER TABLE sales
DETACH PARTITION TS_SEP2007
INTO TS_SEP2007_DETACHED

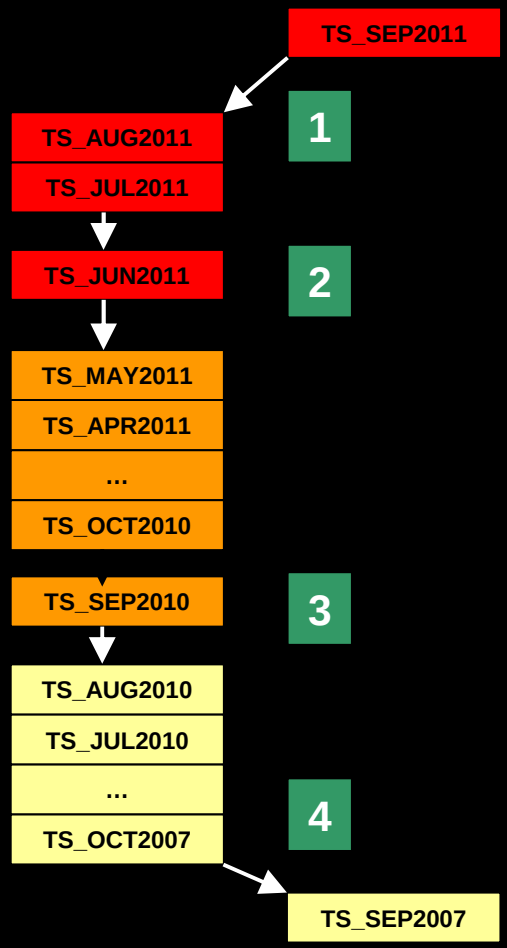

Multi-Temperature Data Management (cont.)

Integration With the DB2 Workload Manager

Low Operational Costs

- DB2 WLM work class and threshold DDL have been extended to support the new data tag attribute
- Work class sets are predictive based on compilation information
  - Sometimes there isn't enough information at compile time (for example, queries with parameter markers) to predict which table spaces will be touched. For this reason…
  - Activities can be mapped to a service subclass based on the data expected to be touched before the activity starts to run
- Data tag thresholds are reactive and use information that is available at runtime
  - Activities can be remapped to a service subclass based on the data touched by the activity as it is running

Multi-Temperature Data Management (cont.)
Optimizes Your Storage Use

Low Operational Costs

- Higher performance
  - Improved ability to meet SLAs
- Lower costs
  - Gracefully extend lifespan of current storage

*"Multi-temperature data management helps in **breaking the I/O thresholds** during intensive log writings in high OLTP environments."*
-- Ahtesham Akhtar, Information Management Consultant SBM.

*"The ability to prioritize access to hot data with DB2 10 Workload Manager and Multi-temperature Storage means a **significant storage cost savings** for our customers."*
-- Radu Parvu, Senior Systems Analyst – Solution Specialist, Accenture, Finnland.

In DB2 10 Early Access Program testing, DB2 obtained an average of 98% compatibility with Oracle PL/SQL

Easily Move from the More Expensive Oracle Database; Leverage Oracle Skills with DB2

Reliance Life Insurance

*"The total cost of ownership with DB2 running on IBM systems is **almost half the cost** of Oracle Database on Sun systems."*

Banco de Crédito del Peru

*"We switched from Oracle Database to IBM DB2 and **cut our costs in half**, while improving performance and reliability of business applications."*

JSC Rietumu Banka

- Moved from Oracle Database to IBM DB2
- Used "compatibility features"
- 3-30x faster query performance
- 200% improvement in data availability

9.7.1	SUB STRB	Increase compatibility
9.7.1	UDF Parameters: INOUT	Increase compatibility
9.7.1	FORALL/BULK COLLECT	Increase compatibility
9.7.1	Improve BOOLEAN	Increase compatibility
9.7.1	Conditional Compilation	Enhancement
9.7.1	Basic DPF Support	Broaden coverage
9.7.1	OCI Support	Broaden coverage
9.7.2	UDF Parameters: DEFAULT	Increase compatibility
9.7.2	Obfuscation	Enhancement
9.7.2	NCHAR, NVARCHAR, NCLOB	Increase compatibility
9.7.3	NUMBER Performance	Performance
9.7.3	Runtime "purity level" Enforcement	Increase compatibility
9.7.3	RATIO_TO_REPORT Function	Increase compatibility
9.7.3	RAISE_APPLICATION_ERROR	Increase compatibility
9.7.3	Small LOB Compare	Increase compatibility
9.7.4	Multi-action Trigger & Update Before Trigger	Increase compatibility
9.7.4	Autonomous Tx Improvements	Increase compatibility
9.7.4	LIKE Improvements, LISTAGG	Increase compatibility
9.7.4	ROW & ARRAY of ROW JDBC Support	Increase compatibility
9.7.5	Pro*C Support	Increase compatibility
9.7.5	Nested Complex Objects	Increase compatibility
10	Local Procedure Definitions	Increase compatibility
10	Local Type Definitions	Increase compatibility
10	PL/SQL Performance	Performance

Higher Performance

- Query Performance
- Index Management
- Continuous Data Ingest

Faster and More Reliable Query Performance

"I'm not really looking for more features that I am going to have to hand tune…

Anything out-of-the-box?"

Multiple Instances of at least 3x Faster Query Performance

Increase Ability to Meet SLAs; Postpone Hardware Upgrades

- Multi-core parallelism enhancements
- Performance improvements for:
  - Queries over star schemas
  - Queries with joins and sorts
  - Queries with aggregation
  - Hash joins
- **Higher performance**
  - Up to **35% faster** out-of-the-box performance
  - Multiple instances of at least **3x faster** when using new features *
- **Lower costs**
  - Postpone hardware upgrades

"IBM and Intel® have collaborated over a decade to optimize DB2 performance with Intel® Parallel Studio 2011, software development suite on Intel® Xeon® processors. We are excited to see a **~10x improvement** in query processing performance using DB2 10 over the previous DB2 version, running on IBM System x3850 using Intel® Xeon® Processor E7. Customers can now realize dramatically greater performance boost at lower cost per query running IBM DB2 10 on servers powered by Intel® Xeon® processors."
—Pauline Nist, GM Software Strategies, Intel's Datacenter & Connected Systems Group

* Based on both external tests by partners, as well as internal tests of IBM DB2 9.7 FP3 vs. DB2 10.1 with new compression features on P6-550 systems with comparable specifications using data warehouse / decision support workloads, as of 29 Mar 2012.

Index Management Re-defined

Increase Ability to Meet SLAs; Lower Administration Costs

- Jump Scan
- Smart Index Pre-fetching
- Smart Data Pre-fetching
- Predicate Evaluation Avoidance
- **Higher performance**
  - Faster index performance
- **Lower costs**
  - Fewer indexes to maintain
  - Dramatic reduction in index reorgs

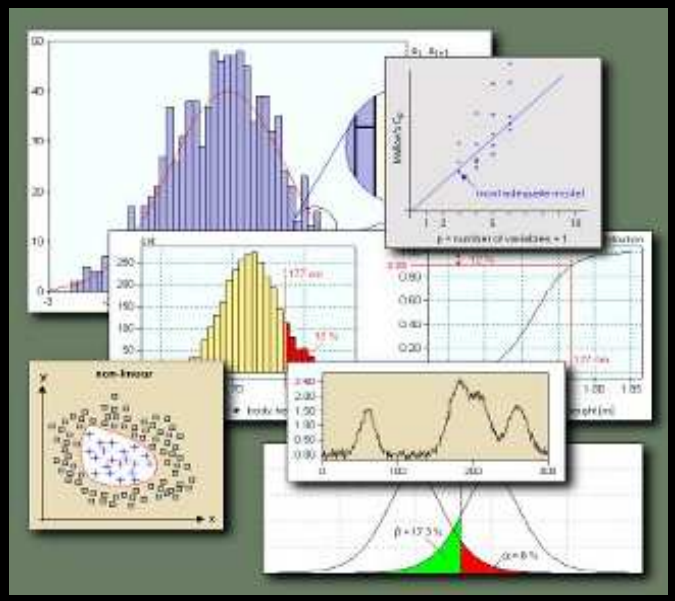

"Jump Scan optimizes buffer usage by **75 to 80%**, resulting in very good improvement in overall performance and saving the CPU cycles."
—Shanmukhaiah D, Cognizant Technology Solutions.

Smart Index Prefetching Means Easier Management

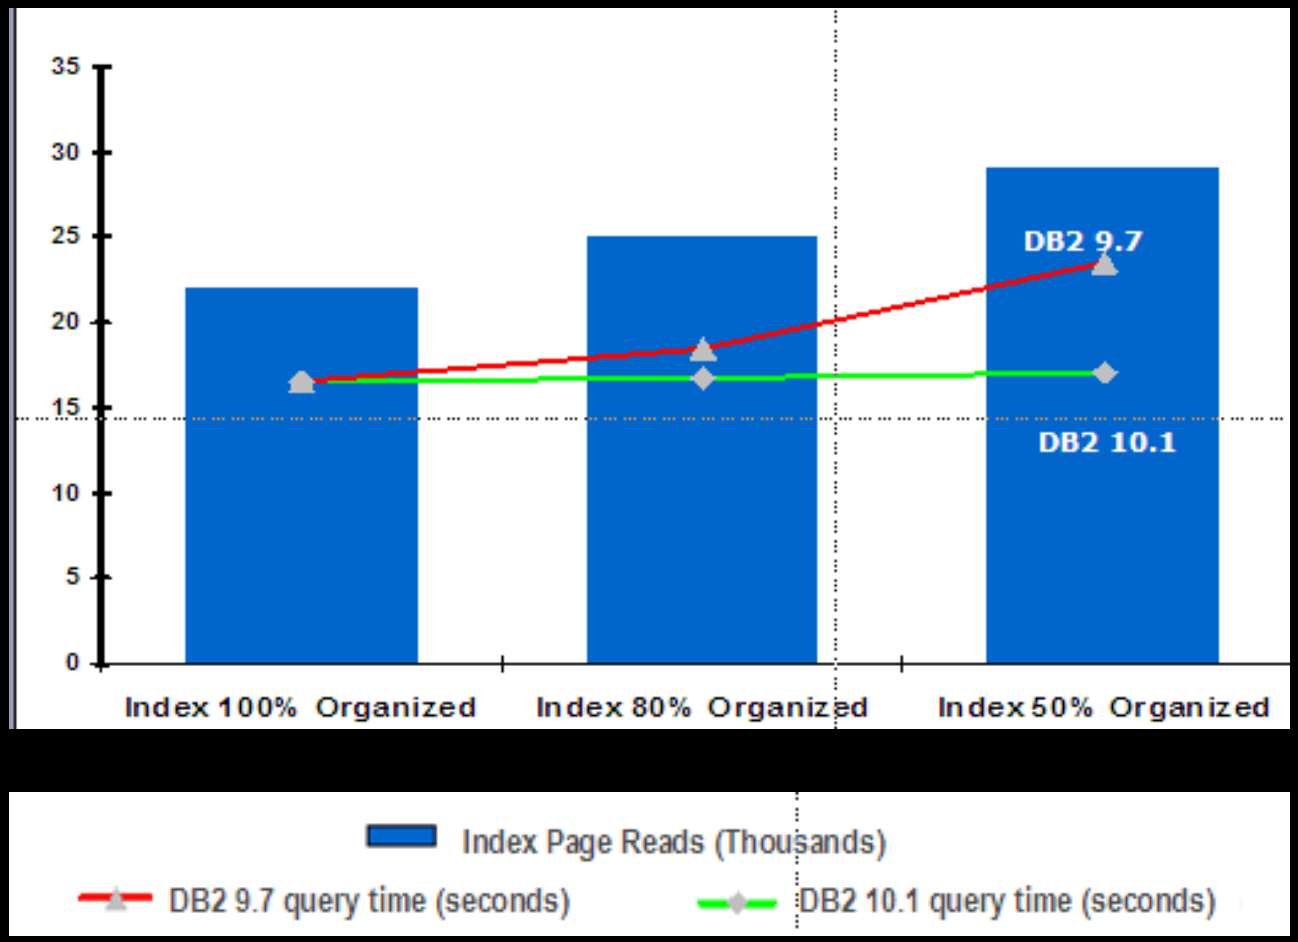

DB2 10.1 28% faster!

Low Operational Costs

Results based [illegible] IBM internal testing

Continuous Feeding of Data

Traditional Data Warehouses

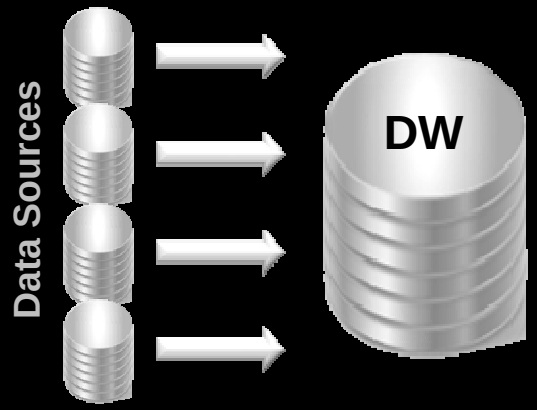

VS

24 Hrs

InfoSphere Warehouse 10

New

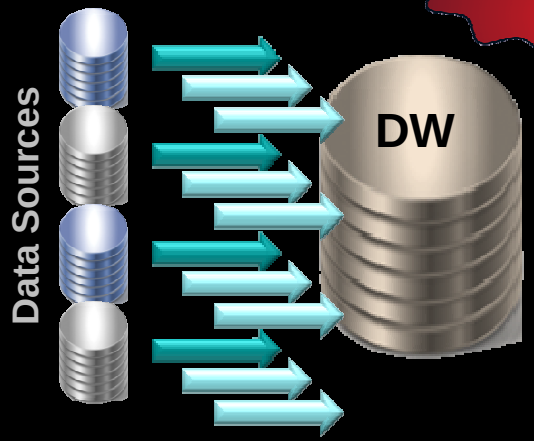

ETL- Extract, Transform, Load	ETL- Extract, Transform, Load
Frequency: Daily Loads, Slower Loads	**Up to the Second Loads, Faster Loads**

LIS.TEC ISV Distributor

"You can now continuously feed data into your data warehouse at a high rate even whilst you are running queries against the tables in your data warehouse. InfoSphere Warehouse 10 represents a greatly strengthened offering for the data warehouse market."
—Ivo Grodtke, LIS.TEC GmbH

INGEST Utility – What is it?

- New, high-speed, client-side utility that streams data from files or pipes into a target table
- Familiar with IMPORT or LOAD? Think of INGEST as a mix
  - Aim is to combine the best features of both, as well as new capabilities!
- INGEST characteristics
  - Fast! Uses an optimized multi-threaded design to process data in parallel and to ingest data into multiple database partitions in parallel
  - Available! Operates at the row level using row-level locking. Allows normal read/write access against the target table while it is running
  - Continuous! Not only can it run against pre-processed data directly from files, INGEST can run in a continuous manner by reading a stream from a pipe
  - Robust! Designed to be tolerant of unexpected failures. Can be restarted from last commit point
  - Flexible and Functional! Supports a number of input formats and has rich data manipulation capabilities

INGEST Utility – What is it? (cont.)

- Avoids use of a staging table when loading a table
  - Common scenario is to load into a staging table, followed by INSERT/SELECT from the staging table to the target table
- Allows other applications to do INSERTs, UPDATEs, or DELETES while the utility loads a table
- Load tables continuously as new data comes in
- More flexible SQL including UPDATE, MERGE, DELETE, function calls, and expressions
- When a recoverable error occurs (for example, connection failed, but re-established), recover automatically and continue

Increased Productivity

- Time Travel Query
- Row and Column access controls
- Data diversity beyond relational
- HADR with Multiple Standby Servers.

Time Travel Query

"We would like to travel to the past to see what our data said at a different point in time.

Can we do this without changing our apps?"

NEW! Time Travel Query

Easily Analyze Historical Trends and Predict Future Demand

- Temporal logic & analysis
- Valid time, transaction time, "AS OF" queries
- **Higher performance**
  - Native support for fast performance
- **Lower costs**
  - Eliminate need to maintain and update custom temporal implementations
  - Easy to administer (simply turn on for any table)

BitemporalData.com

"The use of standardized SQL syntax for temporal operations and the integration deep into the database engine, make DB2 a leader in second generation bitemporal data management - Bitemp 2.0!"
—Craig Baumunk, Principal at BitemporalData.com

IBM

Time Travel Query (cont.)

Ease of Development

- In what department was Ritu as of 12/01/1997?

Emp	Dept	System_start	System_end
Ritu	K25	11/15/1995	01/31/1998
Ritu	M24	01/31/1998	12/31/9999

These values are automatically maintained by DB2 when Ritu's dept is updated

- What account will Marcelo be assigned to on 9/15/2012?

Emp	Account	Bus_start	Bus_end
Marcelo	1	03/01/2011	06/30/2011
Marcelo	2	06/30/2011	12/10/2011
Marcelo	3	03/01/2012	12/31/2012

These values are maintained by the user or application

IBM

Time Travel Query

Reduces Code by up to 45x

Ease of Development

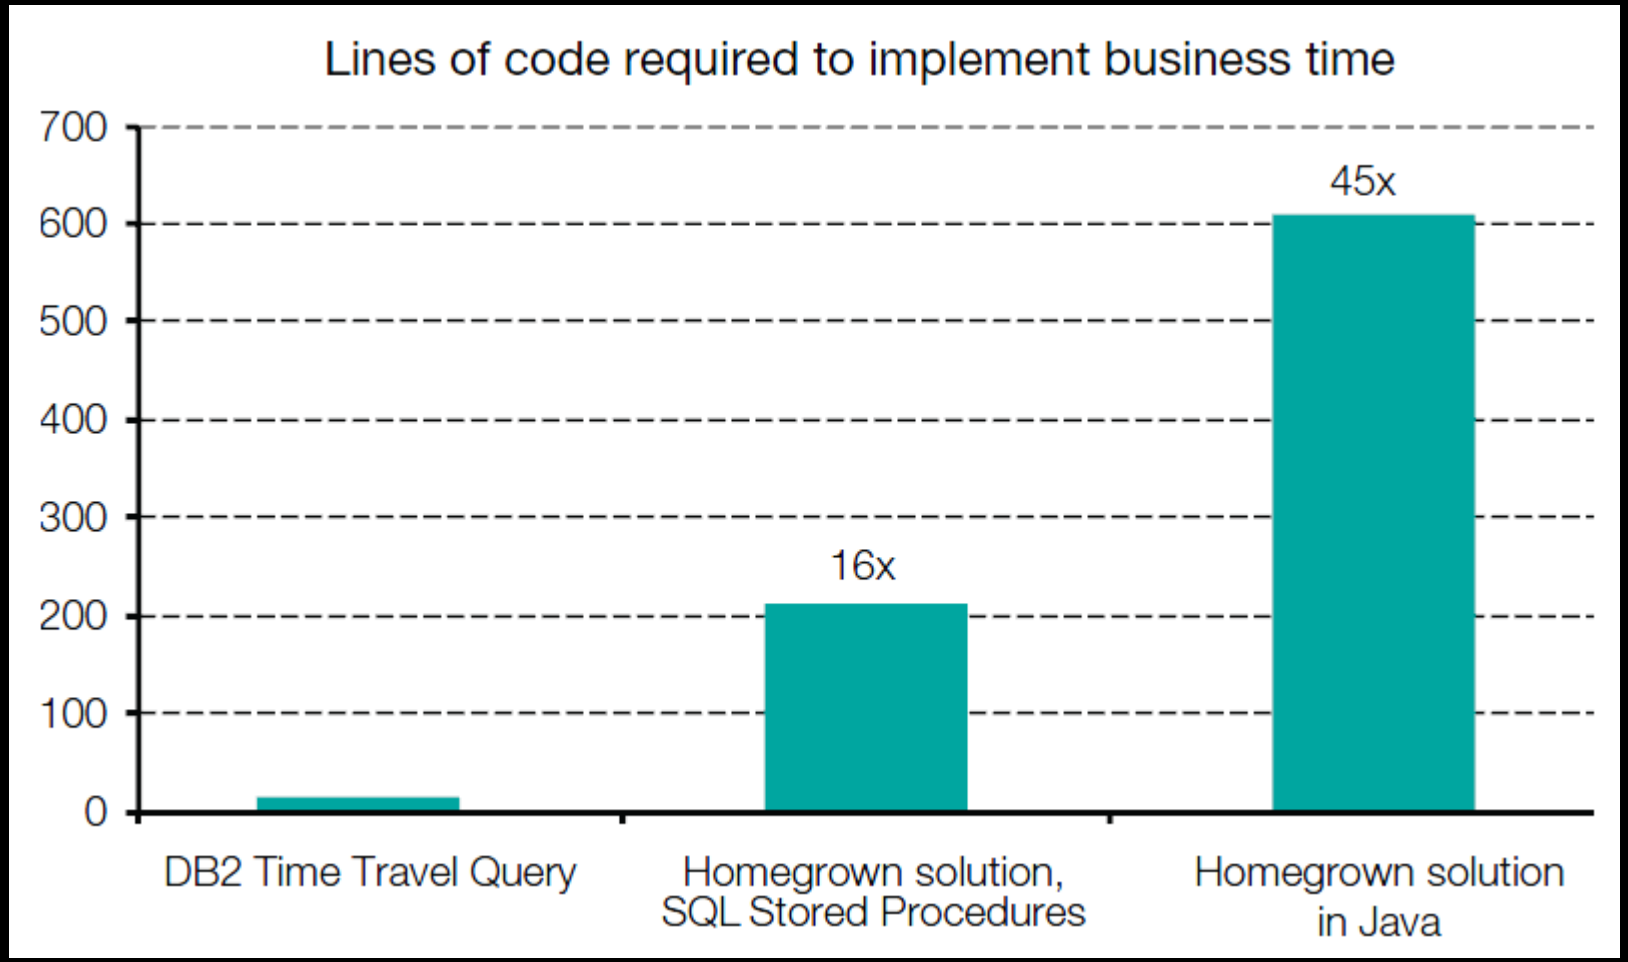

Row and Column Access Control (RCAC)

"How can we be sure that only the proper people, see the appropriate data?"

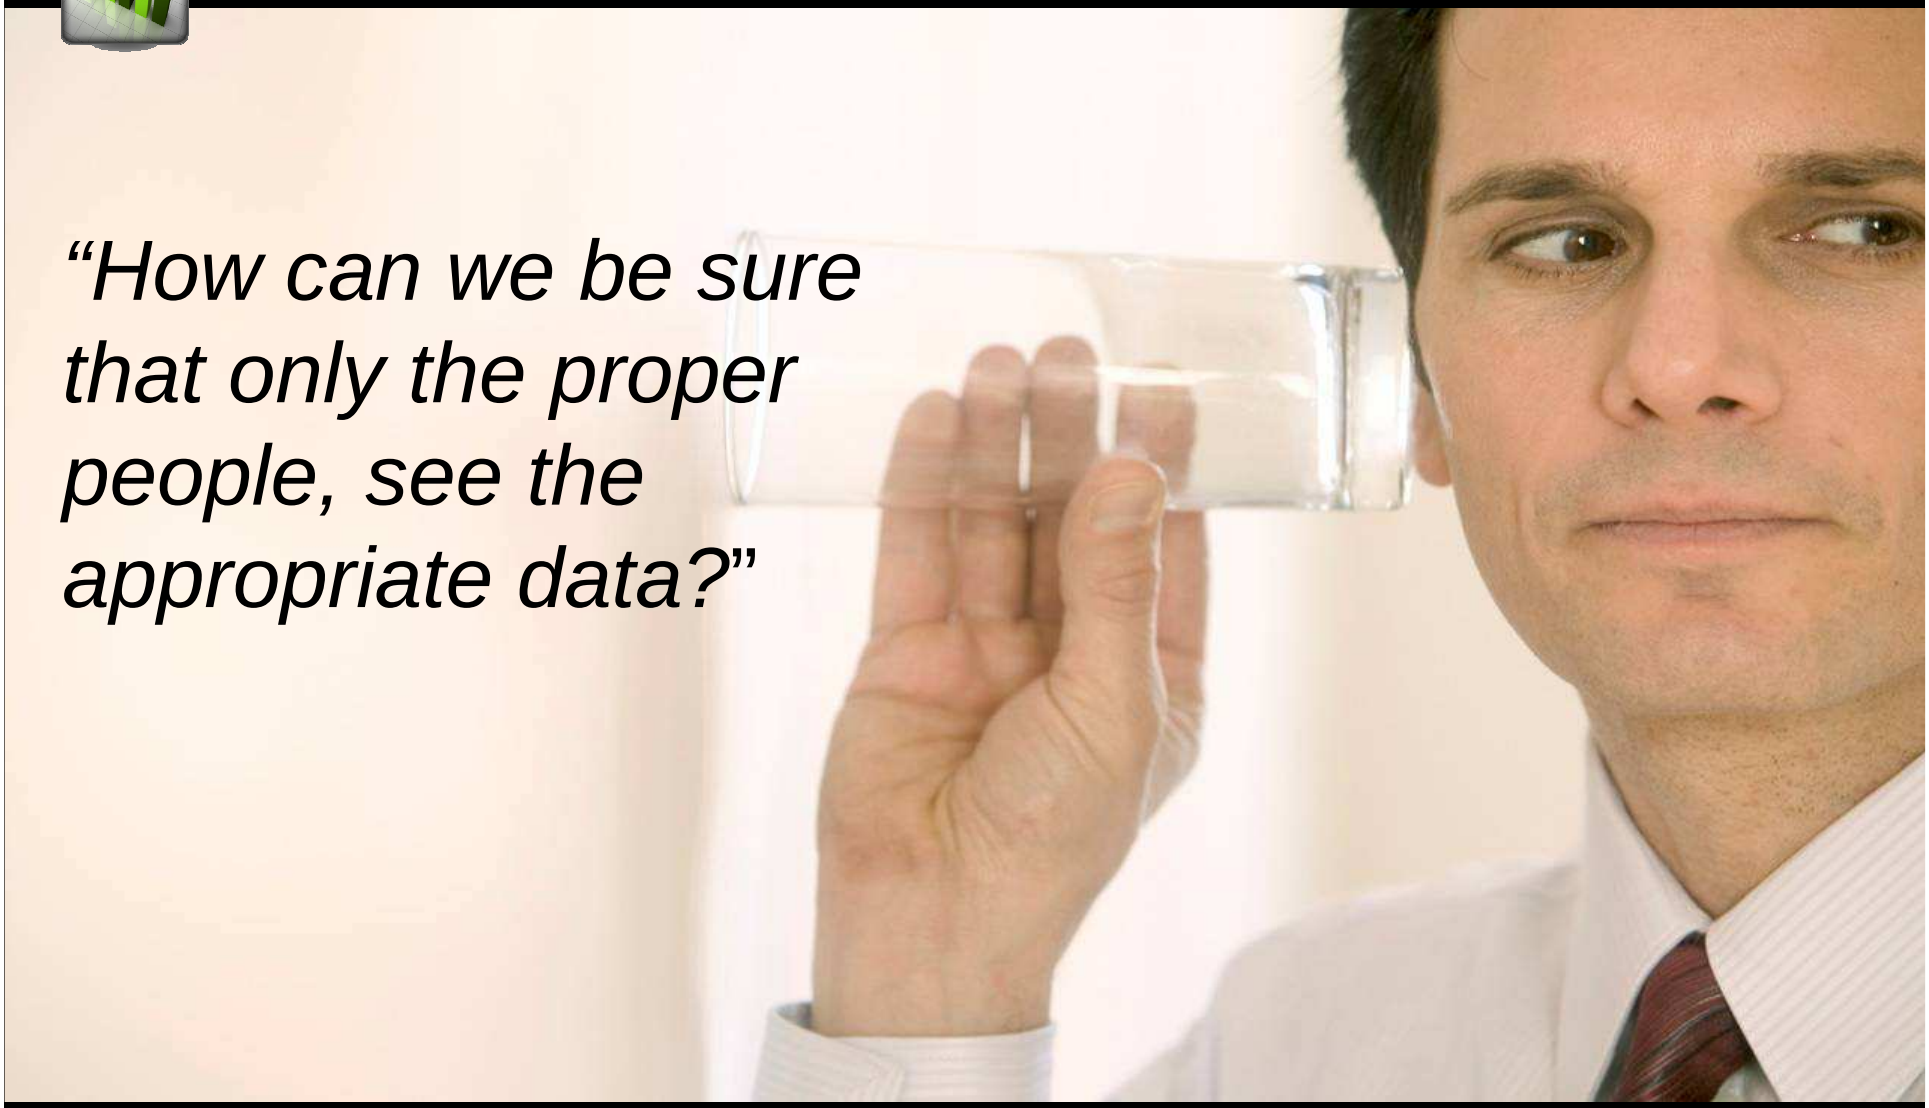

Why Use Row and Column Access Control?

- Currently, data access is restricted via views or application logic
- Users with direct access to databases can bypass these layers
  - Example: Users with DATAACCESS authority can still view all data
- DB2 10.1 for LUW provides a **new way to control data access** at row/column level
  - Set up rich security policies
  - Prevents administrators with DATAACCESS authority from accessing all data in a database
  - No dependency on application logic
  - Allows for data masking (column)
  - Facilitates table level multi-tenancy

How can we ensure that managers only see data for their own employees?

How can we ensure that a social security number column is masked out for unauthorized users?

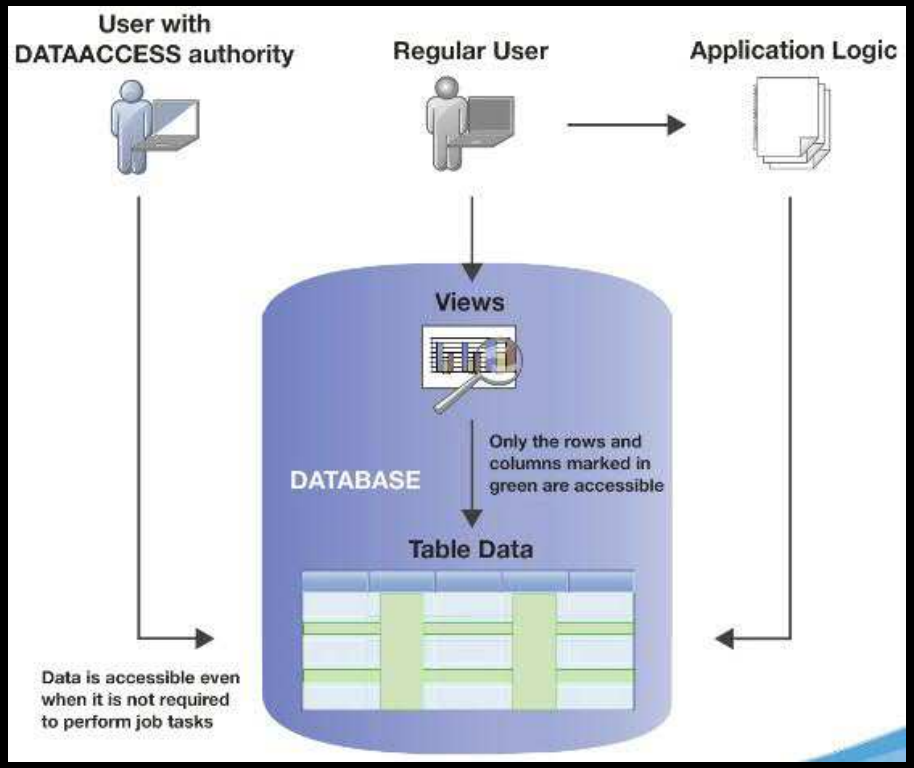

What is Row and Column Access Control?

- **Additional layer of data security** introduced in DB2 10.1 for LUW
- **Complementary** to table level authorization
- Allows **access only to subset of data** useful for job task
- Controls access to a table at the row, column, **or both**
- **Two sets of rules**
  - Permissions for rows
  - Masks for columns

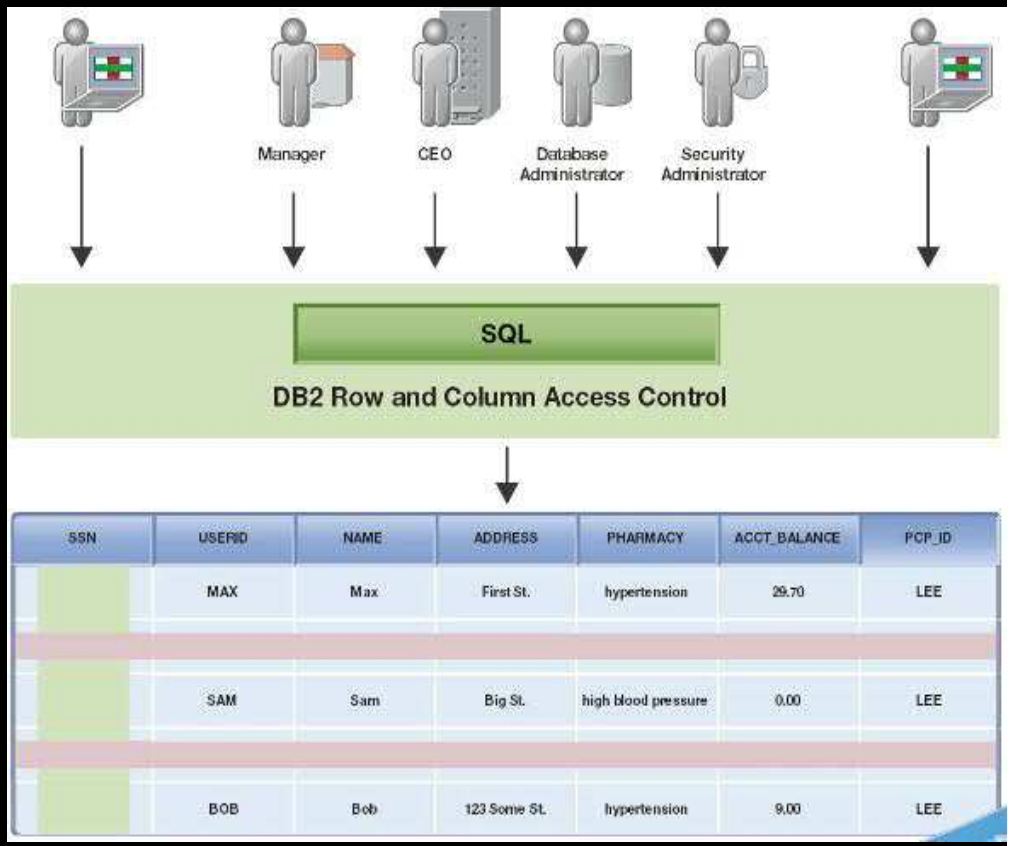

Create Permission

- To create a permission governing access to rows
  1) CREATE the permission with access rule defined by search condition
     - Choose to enforce for all DML or simply select
  2) ENABLE or DISABLE the permission
     - If enabled, this access rule will be implemented when row access control is ACTIVATEd for the affected table
  3) ALTER table to activate row access control

```
CREATE PERMISSION p_name ON table/view FOR ROWS
WHERE search condition ENFORCED FOR ALL ACCESS {disable/enable};

ALTER TABLE/VIEW table/view ACTIVATE ROW ACCESS CONTROL;
```

WHERE clause

ACTIVATE the row access control

Determines if permission will be ENABLED when access control is ACTIVATED for table

Create Column Mask

- To create a mask for a column
  1) CREATE the mask with visibility of column value determined by case expression
  2) ENABLE or DISABLE the permission, determining if this access rule will be implemented when column access control is enabled for the affected table
  3) ALTER table to ACTIVATE column access control

```
CREATE MASK m_name on t_name FOR COLUMN c_name RETURN
case-expression {disable/enable}

ALTER TABLE/VIEW table/view ACTIVATE COLUMN ACCESS CONTROL;
```

Result of case expression is returned in substitute of column value

Determines if mask will be enabled when access control is ACTIVATEd for table

ACTIVATE column access control

NEW! Graph Store

Rapid Application Development

- Optimized way to store graph triples in DB2

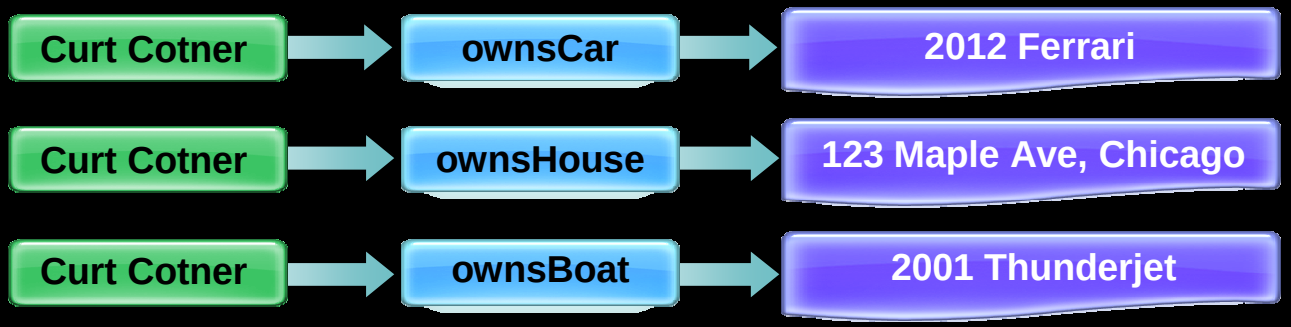

- Supports SPARQL 1.0 query language
- **Higher performance**
  - Accelerates leading open source semantic Web framework by up to 3.5x
- **Lower costs**
  - Rapid development with schema-less approach
  - Easy adaption as needs evolve
  - Simpler data management for triples

NEW! HADR now Supports Multiple Standby Servers

Increase Ability to Meet SLAs; Disaster Recovery

- HADR now supports more than one stand-by server
- If Primary Server fails, Principal Standby takes over
- If Principal Standby then fails, can switch to Auxiliary Standby
- Auxiliary Standby can provide complete offsite availability, while maintaining speed of local standby

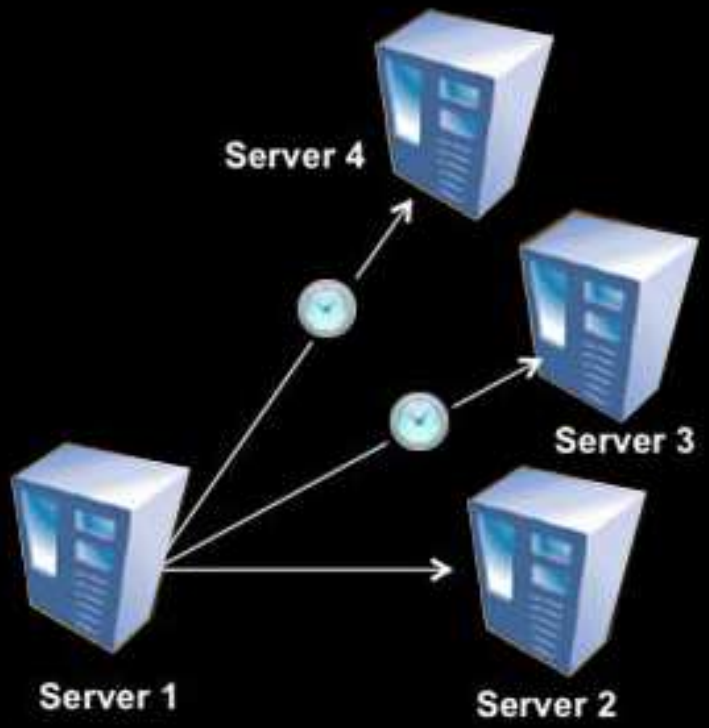

Cost/Benefit Case for Advanced Enterprise Server Edition and InfoSphere Optim Operational DBA Tools *- ITG Report*

- Average **41% lower** license and support costs
- Average **20% lower** infrastructure costs, including costs for database servers and disk storage systems
- Average **26% lower** DBA personnel costs
- Average **37% lower** combined costs

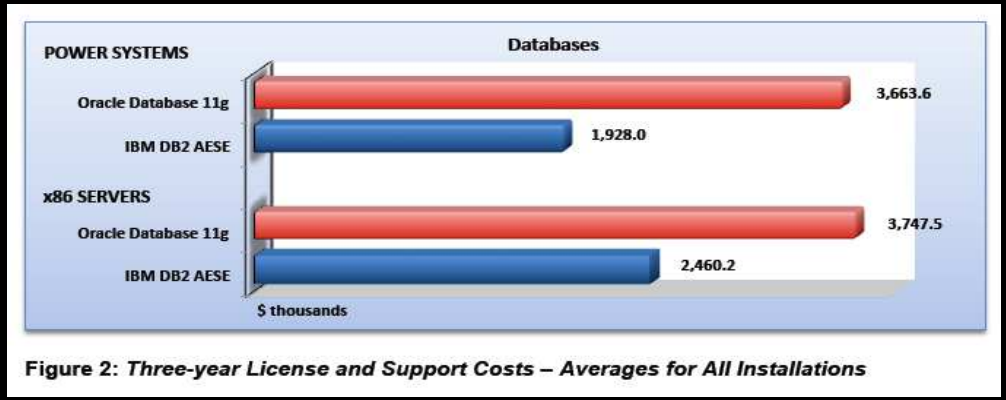

Figure 2: ***Three-year License and Support Costs – Averages for All Installations***

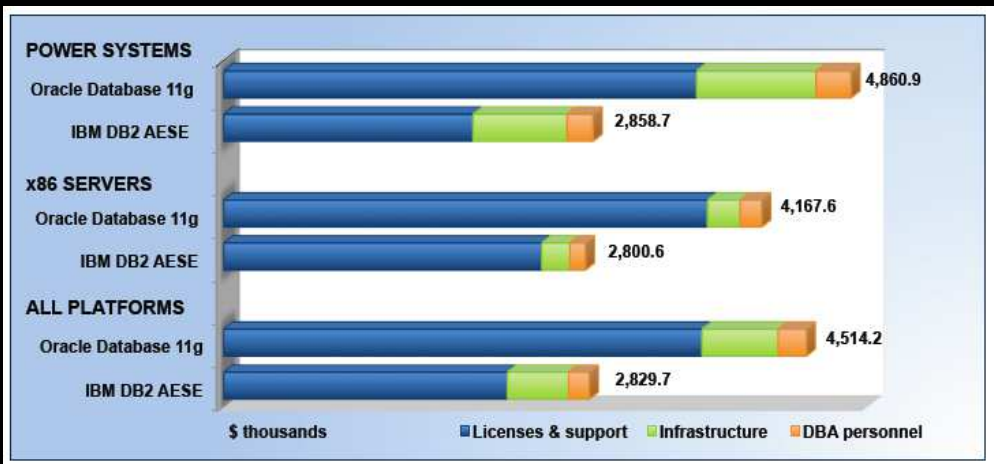

Figure 3: ***Overall Three-year Costs – Averages for All installations***

* January 2012 Management Brief - ITG Report Cost/Benefit Case for IBM DB2 Advanced Enterprise Server Edition and InfoSphere Optim Operational DBA Tools - Comparing Costs and Capabilities with Oracle Database 11g

Competitive Pricing DB2 AESE vs Oracle Enterprise Edition

DB2 AESE Delivers Better Value per Core

Function	Included with Oracle TopLink Workbench	Included with DB2 AESE
Java Data Access	Oracle Toplink Workbench	Optim pureQuery Runtime for DB2
SQL Performance Tuning	Oracle Tuning Pack	Optim Query Workload Tuner
Stored Procedure Monitoring	Oracle Stored Procedure Monitoring	n/a
Development	Oracle Internet Development Suite	Data Studio
Graph Store	Oracle Graph Store	DB2 RDF Client Package
Multi-site Active Active Replication	Oracle Golden Gate	DB2 Q-Replication
Provisioning	Oracle Provisioning	n/a
Data Partitioning	Oracle Partitioning	DB2 Table / Range Partitioning
Multi-level Security	Oracle Label Security	DB2 Advanced Access Control
Data Compression	Oracle Advanced Compression	DB2 Storage Optimization with Adaptive Compression
Data Modelling	Oracle SQL Developer Data Modeller	InfoSphere Data Architect

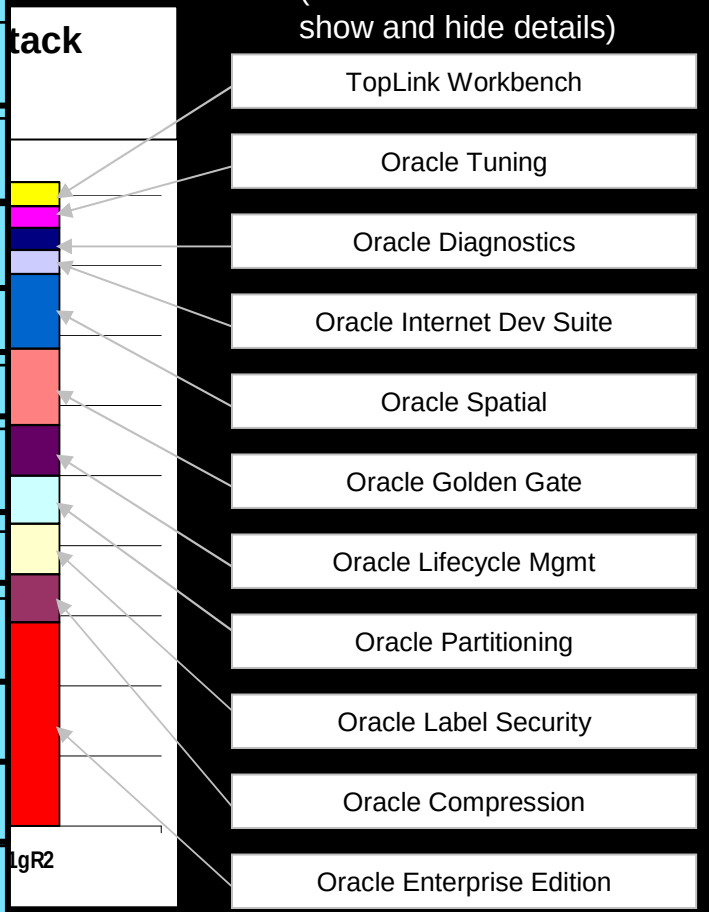

* Based on publically available US $ list prices as of March 15, 2012 including year 1 S&S. DB2 AESE price assumes 100 PVU / core. Oracle price assumes 1.0 multiplier and is calculated by summing the separately purchasable packages and features highlighted on this chart.

Advanced Enterprise Server Edition For DB2 10 and InfoSphere Warehouse 10 Delivers

Improved Cost Efficiencies

- Adaptive Compression
- Multi-Temperature Data Management
- On-going SQL compatibility

Higher Performance

- Query Performance
- Index Management
- Continuous Data Ingest

Increased Productivity

- Time Travel Query
- Row and Column access controls
- Data diversity beyond relational
- HADR with multiple standby servers

THANK YOU

Backup Slides

Database Management Solutions Included in Advanced Enterprise Server Edition

	DB2 10 AESE	ISW 10 Advanced Departmental Edition	ISW 10 Advanced Enterprise Edition
Design	InfoSphere Data Architect *	• Design Studio • InfoSphere Data Architect *	• Design Studio • InfoSphere Data Architect*
Development & Administration	Data Studio	• Data Studio • SQL Warehousing Tool	• Data Studio • SQL Warehousing Tool
Performance Management	• InfoSphere Optim Query Workload Tuner • InfoSphere Optim Configuration Manager • InfoSphere Optim pureQuery Runtime • InfoSphere Optim Performance Manager Extended Edition	• InfoSphere Optim Query Workload Tuner • InfoSphere Optim Configuration Manager • InfoSphere Optim Performance Manager Extended Edition	• InfoSphere Optim Query Workload Tuner • InfoSphere Optim Configuration Manager • InfoSphere Optim Performance Manager Extended Edition
Advanced Recovery	SPP	SPP	• DB2 Merge Backup • DB2 Recovery Expert • InfoSphere Optim High Performance Unload

NEW! Row and Column Access Control

Easy Compliance with Privacy and Sensitive Data Requirements

- Fine-grained access control
  - Hide rows from unauthorized users
  - Mask the value of columns for unauthorized users
- Policy-driven security, with flexible policies
- Does not require classification

Account	Name	Income	Branch
1111-2222-3333-4444	Ana	22,000	A
2222-3333-4444-5555	Bob	71,000	B
3333-4444-5555-6666	Celia	123,000	B
4444-5555-6666-7777	Dinesh	172,000	C

Teller
Amy sees

Account	Name	Income	Branch
2222-3333-4444-5555	Bob	71,000	B
3333-4444-5555-6666	Celia	123,000	B

Telemarketer
Pat sees

Account	Name	Income	Branch
xxxx-xxxx-xxxx-4444	Ana	22,000	A
xxxx-xxxx-xxxx-5555	Bob	71,000	B
xxxx-xxxx-xxxx-6666	Celia	123,000	B
xxxx-xxxx-xxxx-7777	Dinesh	172,000	C

NEW! Row and Column Access Control

Easy Compliance with Privacy and Sensitive Data Requirements

- **Higher performance**
  - Less data duplication than using "Views" to mask data
  - More secure than using "Views" to mask data
- **Lower cost**
  - Easier to implement and maintain
  - Easier compliance with privacy and sensitive data requirements
  - Easier to maintain that using application code to mask data

"Because we deal with sensitive securities and financial information, the privacy of that information is a top priority. Row and Column Access Control will help enhance our security solutions and help us meet strict regulatory guidelines."
—Shi Jin Li, China Securities Depository and Clearing Corporation Ltd.

"Row and Column Access Control help us to improve data confidentiality and security in production environments."
—Jessica Tatiana Flores Montiel, DAFROS Multiservicios

Competitive Pricing DB2 Advanced Enterprise Server Edition vs Oracle Database Enterprise Edition

DB2 AESE Delivers Better Value per Core

DB2 AESE 10.0

- **Advanced Database**
- **Rich Tools and Features**
- **One Low Cost**
- **One Part**

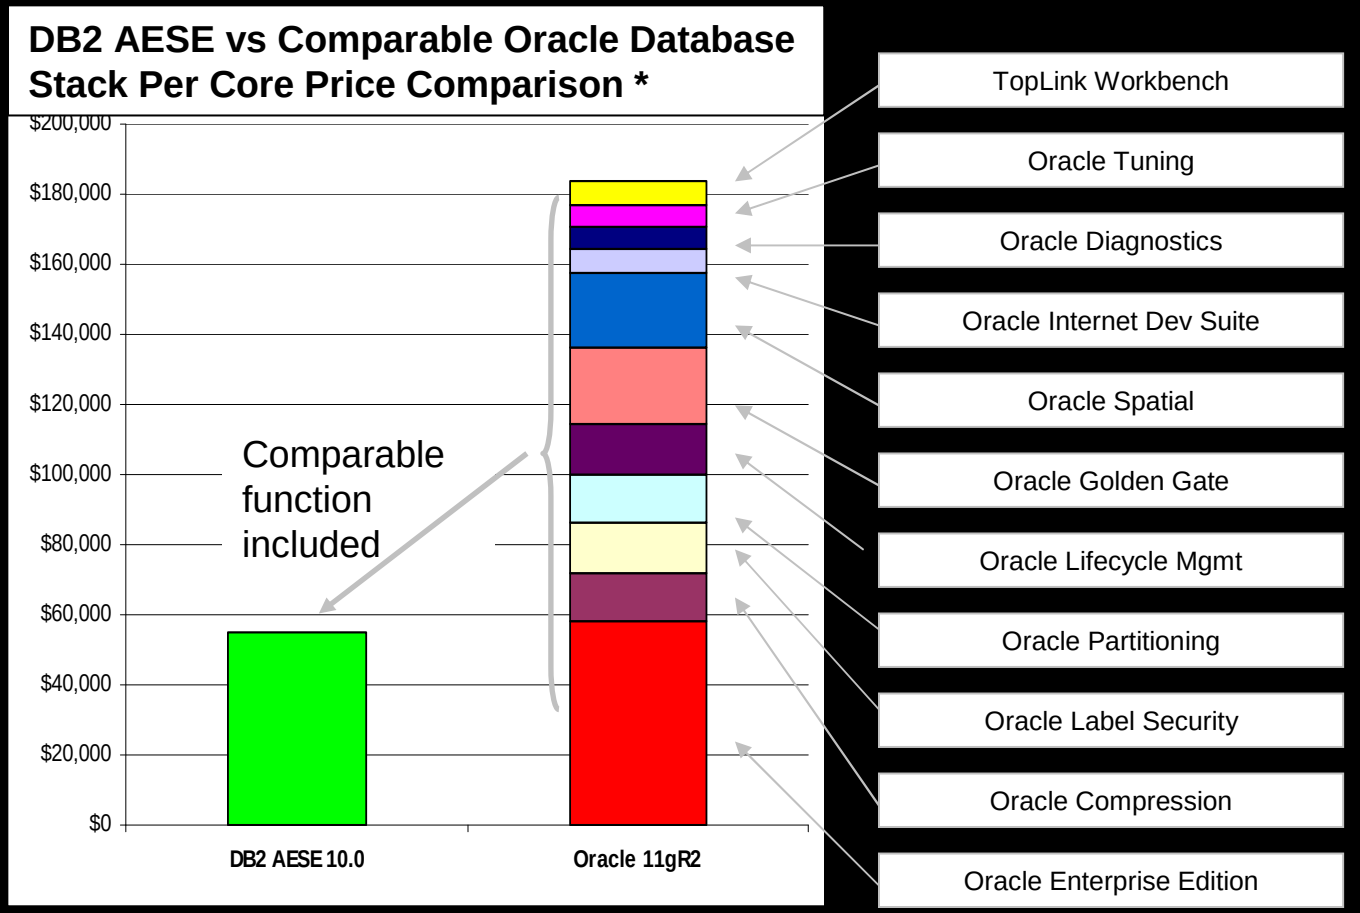

* Based on publically available US $ list prices as of March 15, 2012 including year 1 S&S. DB2 AESE price assumes 100 PVU / core. Oracle price assumes 1.0 multiplier and is calculated by summing the separately purchasable packages and features highlighted on this chart.

Competitive Pricing ISW Advanced Enterprise vs Oracle Database Enterprise Edition

ISW Advanced Edition Delivers Better Value per Core

ISW Advanced Enterprise 10.0

- **Advanced Data Warehouse**
- **Rich Tools and Features**
- **One Low Cost**
- **One Part**
- **Terabyte Pricing Available**

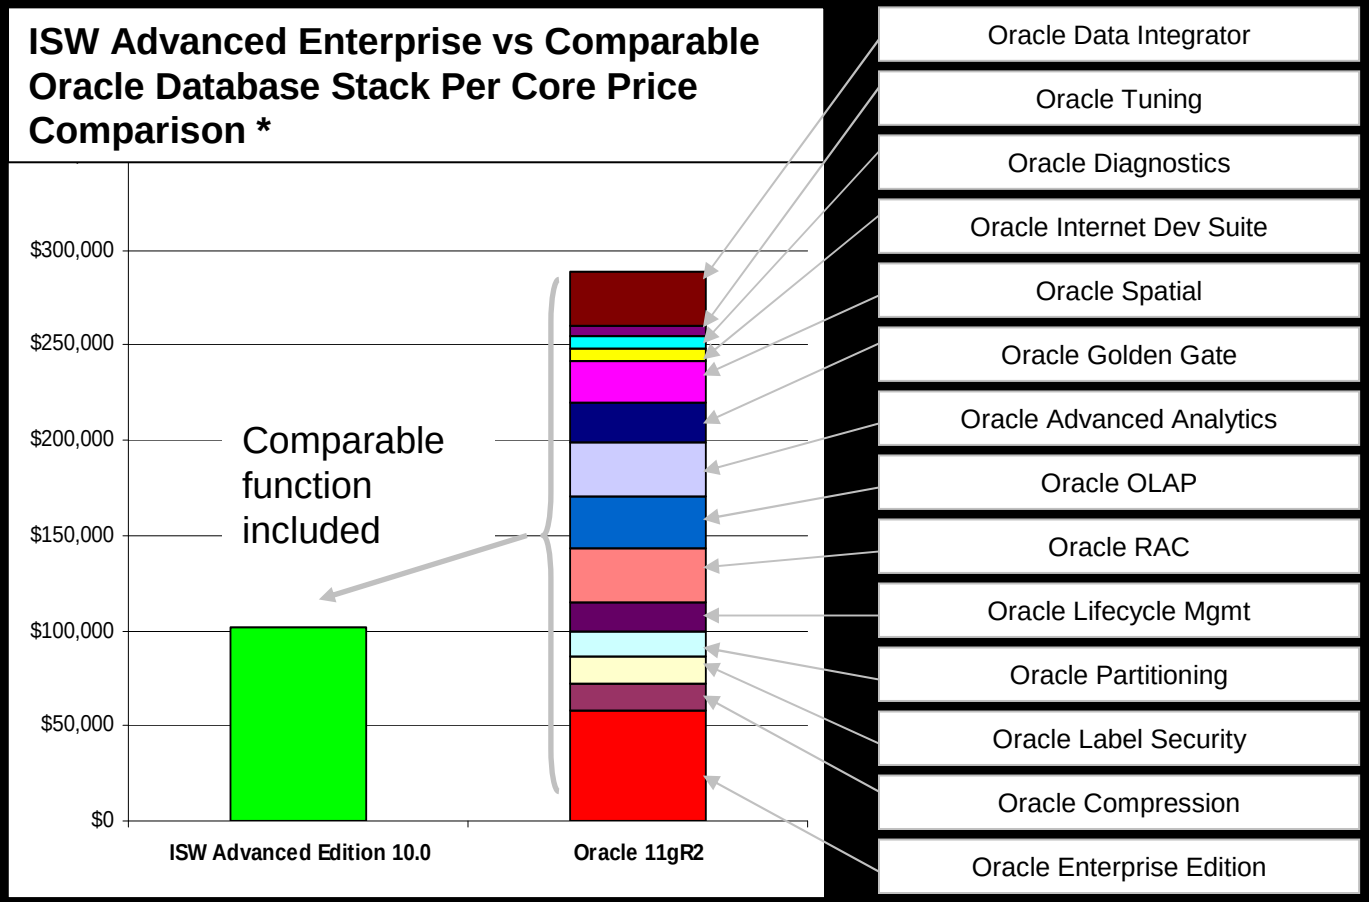

* Based on publically available US $ list prices as of March 15, 2012 including year 1 S&S. ISW Advanced Enterprise price assumes 100 PVU / core. Oracle price assumes 1.0 multiplier and is calculated by summing the separately purchasable packages and features highlighted on this chart.