IBM Software Group

DB2 9 Bootcamp
Technology Overview

SWIT – Information Management Technical Presales

Michele Benedetti

ON DEMAND BUSINESS

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Agenda

- DB2 9 Strategy and Key Investment Areas
- pureXML: native XML Support
- Reducing the Total Cost of Ownership
- Expanding Database Capacity and Removing Limits
- Table Compression
- Granular Security

IBM Database Technology Strategy

- Continued Focus on Performance, Scale, Availability
- Reduce TCO and Accelerate Time-to-value
- Support for New Data Types
- Deep Cross-middleware Integration

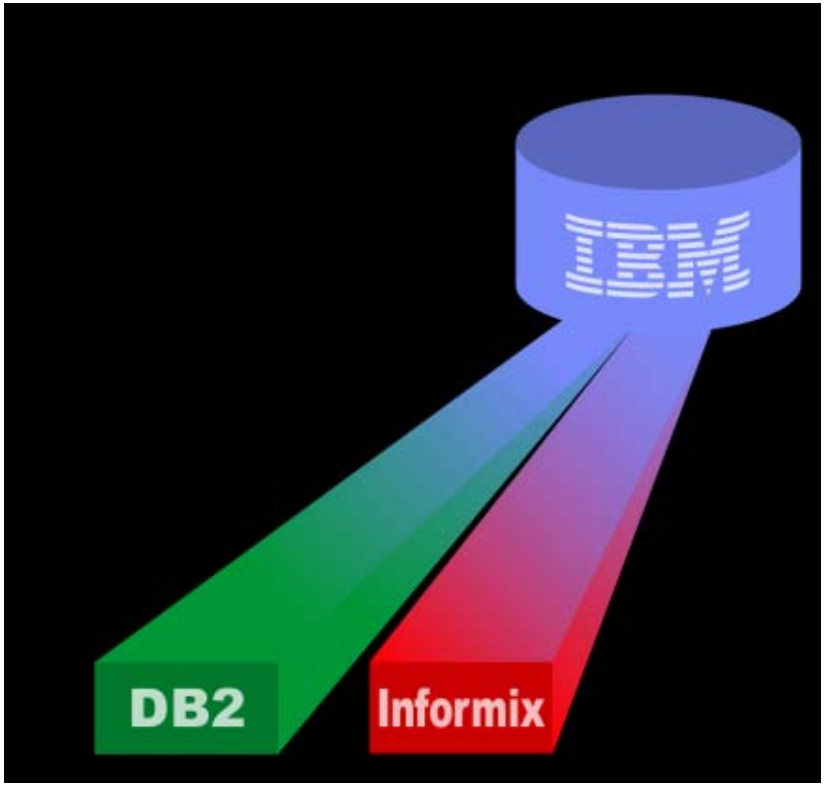

Commonality Across DB Servers

IBM Software Group

pureXML® Technology

Integrating XML into the DB2 Engine

ON DEMAND BUSINESS™

Evolution in “Database” Technology: XML has changed the rules!

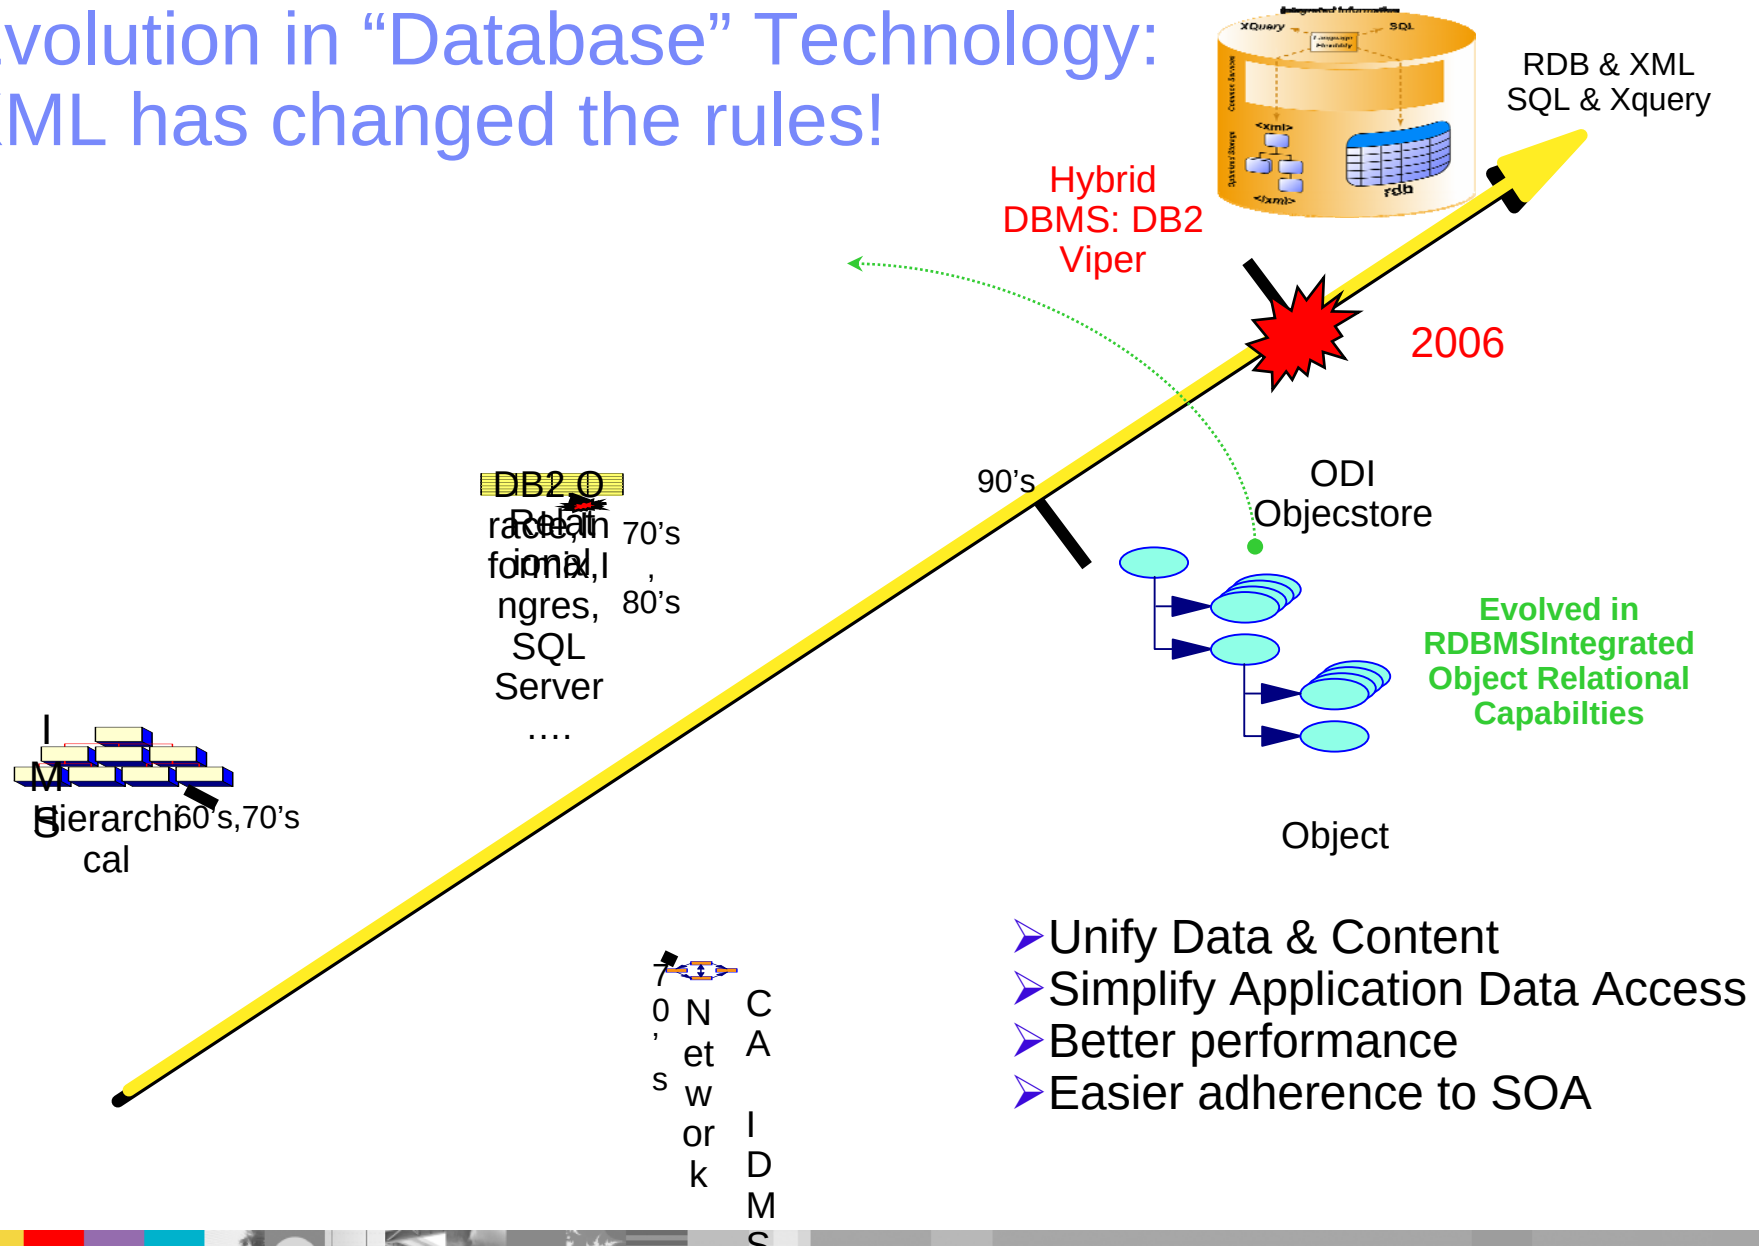

- Unify Data & Content
- Simplify Application Data Access
- Better performance
- Easier adherence to SOA

What is XML?

- XML Technology
  - XML = Extensible Markup Language
  - Self-describing data structures
  - XML Tags describe each element and their attributes

```
<? xml version="1.0" ?>
<purchaseOrder id='12345" secretKey='4x%$^'>
    <customer id="A6789">
        <name>John Smith Co</name>
        <address>
            <street>1234 W. Main St</street>
            <city>Toledo</city>
            <state>OH</state>
            <zip>95141</zip>
        </address>
    </customer>
    <itemList>
        <item>
            <partNo>A54</partNo>
            <quantity>12</quantity>
        </item>
        <item>
            <partNo>985</partno>
            <quantity>1</quantity>
        </item>
    </itemList>
</purchaseOrder>
```

- Benefits
  - Extensible
    - No fixed format or syntax
    - Structures can be easily changed
  - Platform Independent
    - Not tied to any platform, operating system, language or software vendor
    - XML can be easily exchanged
  - Fully Unicode compliant

XML Example: Financial Data (FIXML)

- Buying 1000 Shares of IBM Stock..

8=FIX.4.2^9=251^35=D^49=AFUNDMGR^56=ABROKER^34=2
^52=20030615-01:14:49^11=12345^1=111111^63=0^64=2003
0621^21=3^110=1000^111=50000^55=IBM^48=459200101^22=
1^54=1^60=2003061501:14:4938=5000^40=1^44=15.75^15=USD
^59=0^10=127

Old FIX Protocol

New FIXML Protocol

- extensible
- lower appl development & maintenance cost

```
<FIXML>
  <NewOrdSingle  ClOrdID="123456"
    Side="2"
    TransactTm="2003-06-15T01:14:49-05:00"
    OrderType="2"
    Price="93.25"
    Acct="26522154">
    <Header  Sent="2001-06-21T01:31:28-05:00"
      PosDup="N"
      PosRsnd="N"
      SeqNum="521">
      <Sender  ID="AFUNDMGR"/>
      <Target  ID="ABROKER"/>
    </Header>
    <Instrument  Symbol="IBM"
        ID="459200101"
        IDSrc="1"/>
    <OrderQuantity  Qty="1000"  Cur="USD"/>
  </NewOrdSingle>
</FIXML>
```

XML Market Projections

- **XML Storage is a high growth area**

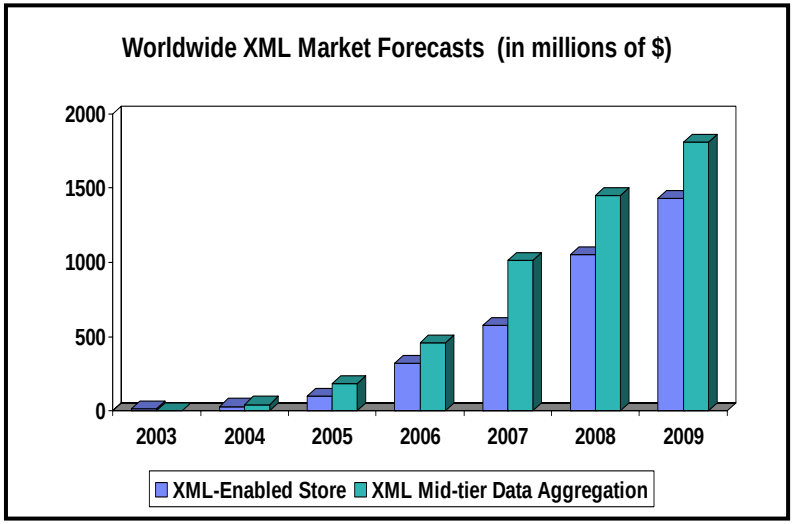

Figure VI.2: Market Size by XML Data Store Solution Type

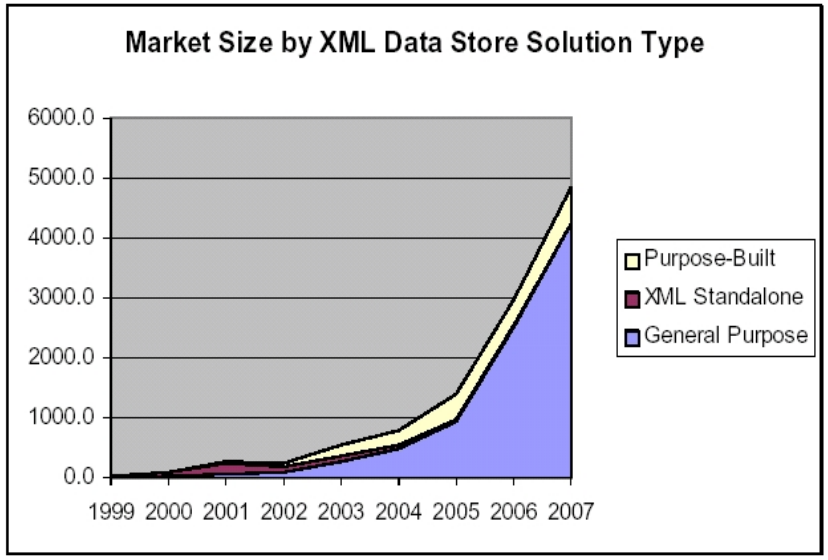

- **XML database revenue to grow at twice the rate of the total database market**

- IDC

Worldwide Enterprise Database Management Systems
Software Forecast Update, 2003-2007

Chart Sources: XML Market Opportunities, Forecasts and Strategies, 2004-2009
Wintergreen Research Inc. ZapThink

XML – A Very Different Data Model

- Relational is a data model:
  - Relations (tables)
  - Attributes (columns)
  - Set based w/ some sequences
  - Strict schema

POID	CustomerID	ItemID
12	1	2
162	3	4
162	3	5

Id	LastName	FirstName	Street	City	State	Zip
1	Pirahesh	Hamid	1 Harry Rd	San Jose	CA	95141
3	Selinger	Pat	555 Bailey Ave	San Jose	CA	95141

ItemID	Name
2	#6 wire nut
5	Small Walrus
4	Apollo moon rocket

- XML is a data model:
  - Nodes (elements, attributes, comments, etc.)
  - Relationships between nodes
  - Sequence based w/ some sets
  - Flexible schema

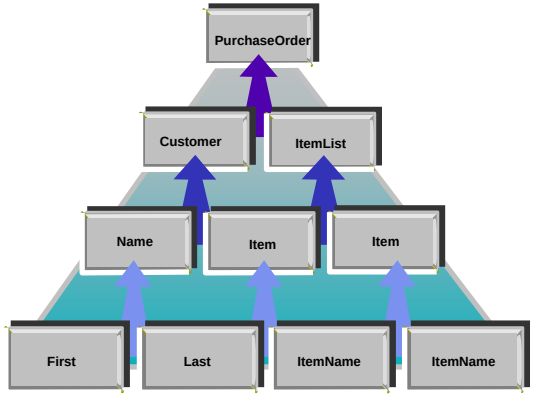

pureXML® : Native XML Storage

- Must store XML in parsed hierarchical format
  (similar to the DOM representation of the XML infoset)

```
create table dept (deptID char(8),…, deptdoc xml);
```

- Relational columns are stored in relational format (tables)
- XML columns are stored natively
- XML stored in UTF8

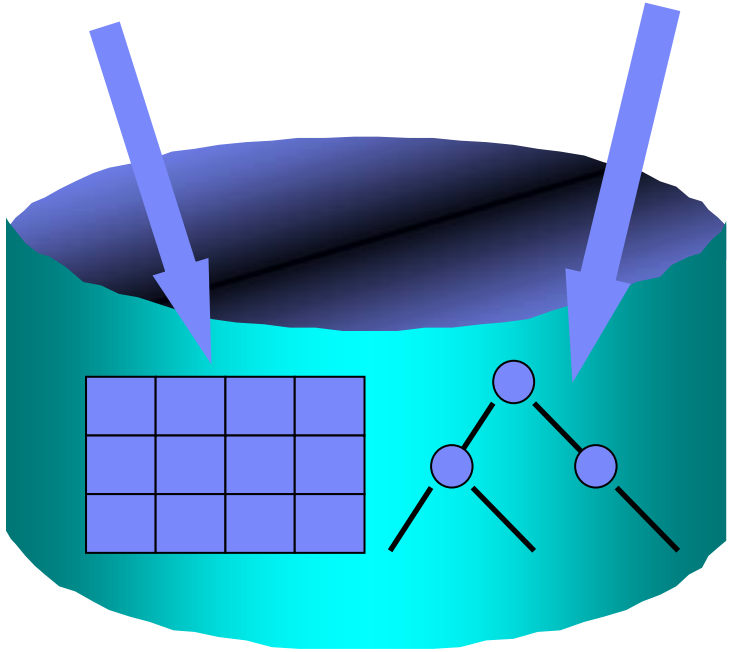

XML in DB2 9

SQL Person… "I see a world class RDBMS that also supports XML"

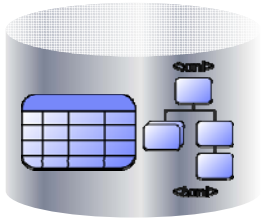

DB2 with XML Support

XML Person… "I see a world class XML repository that also supports SQL"

XML integrated in all facets of DB2!

New XML applications benefit from:

- Ability to seamlessly leverage relational investment
- Proven Infrastructure that provides enterprise-class capabilities

XML In DB2 9
SQL and XML on Equal Footing

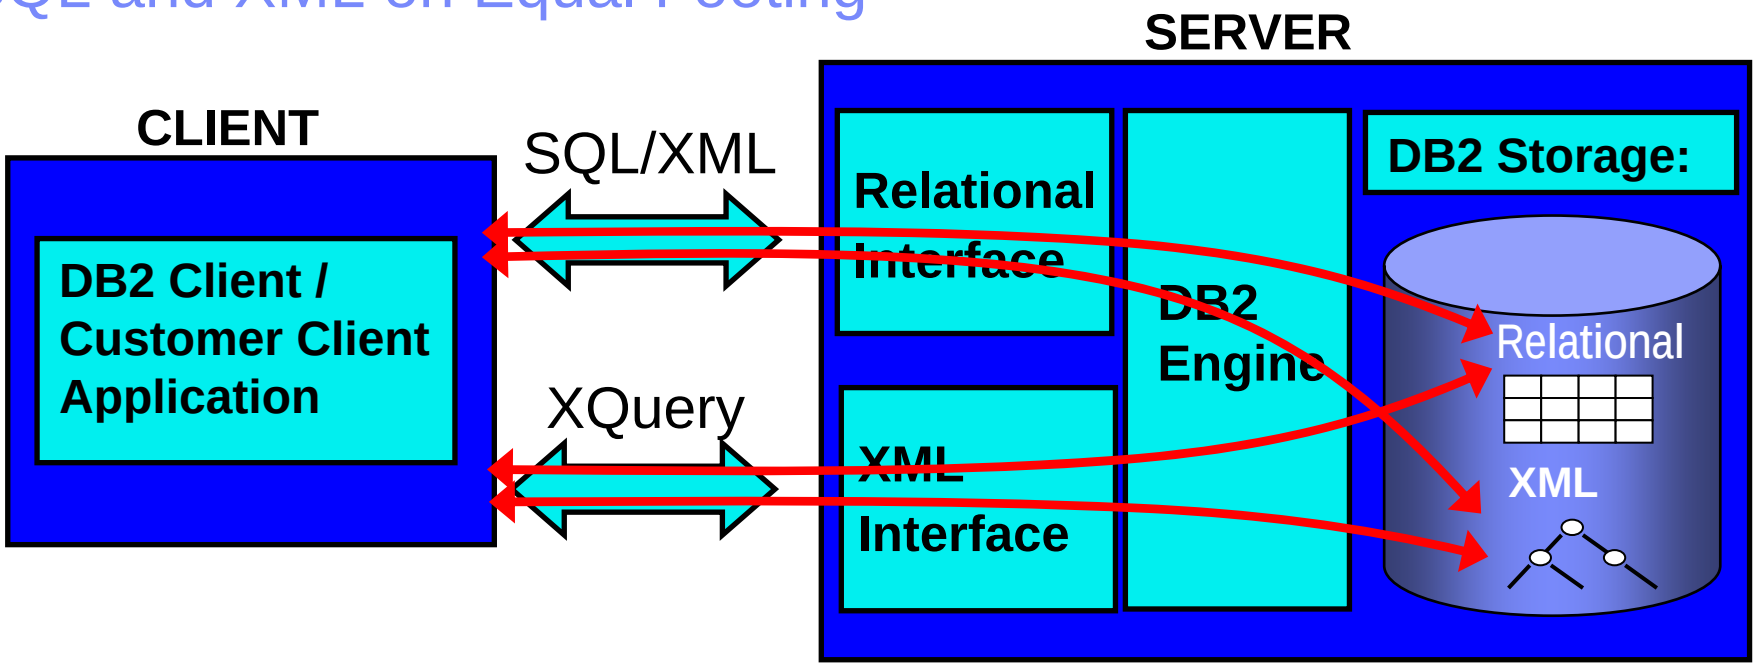

XML capabilities built into DB2

Deep enabling and native support for XML means high performance!

Optimized data store

New storage and indexing techniques for XML

Multiple Query Interfaces

SQL/XML and XQuery

⇒ Both languages have full access to stored data

⇒ Pick the view of the data that best suits the application

Querying XML Data in DB2

Four options are supported:

- XQuery/XPath as a stand-alone language

- XQuery with embedded SQL

- SQL/XML: SQL with embedded XQuery/XPath

- Plain SQL for full-document retrieval

Automation Automatically!

- Enable many of the DB2 autonomic computing features by default.
- Examples:
  - Configuration Advisor (2 second tuning)
  - Adaptive Self Tuning Memory
  - Automatic data statistics collection.
  - Automatic I/O Parameters Configuration (IO_CLEANERS, IO_SERVERS)

Adaptive Self-Tuning Memory

- DB2 9 introduce a revolutionary memory tuning system called the Self Tuning Memory Manager (STMM)
  - Works on main database memory parameters
    - Sort, locklist, package cache, buffer pools, and total database memory
  - Hands-off online memory tuning
    - Requires no DBA intervention
  - Senses the underlying workload and tunes the memory based on need
  - Can adapt quickly to workload shifts that require memory redistribution
  - Adapts tuning frequency based on workload

STMM in Action – Dropping an Important Index

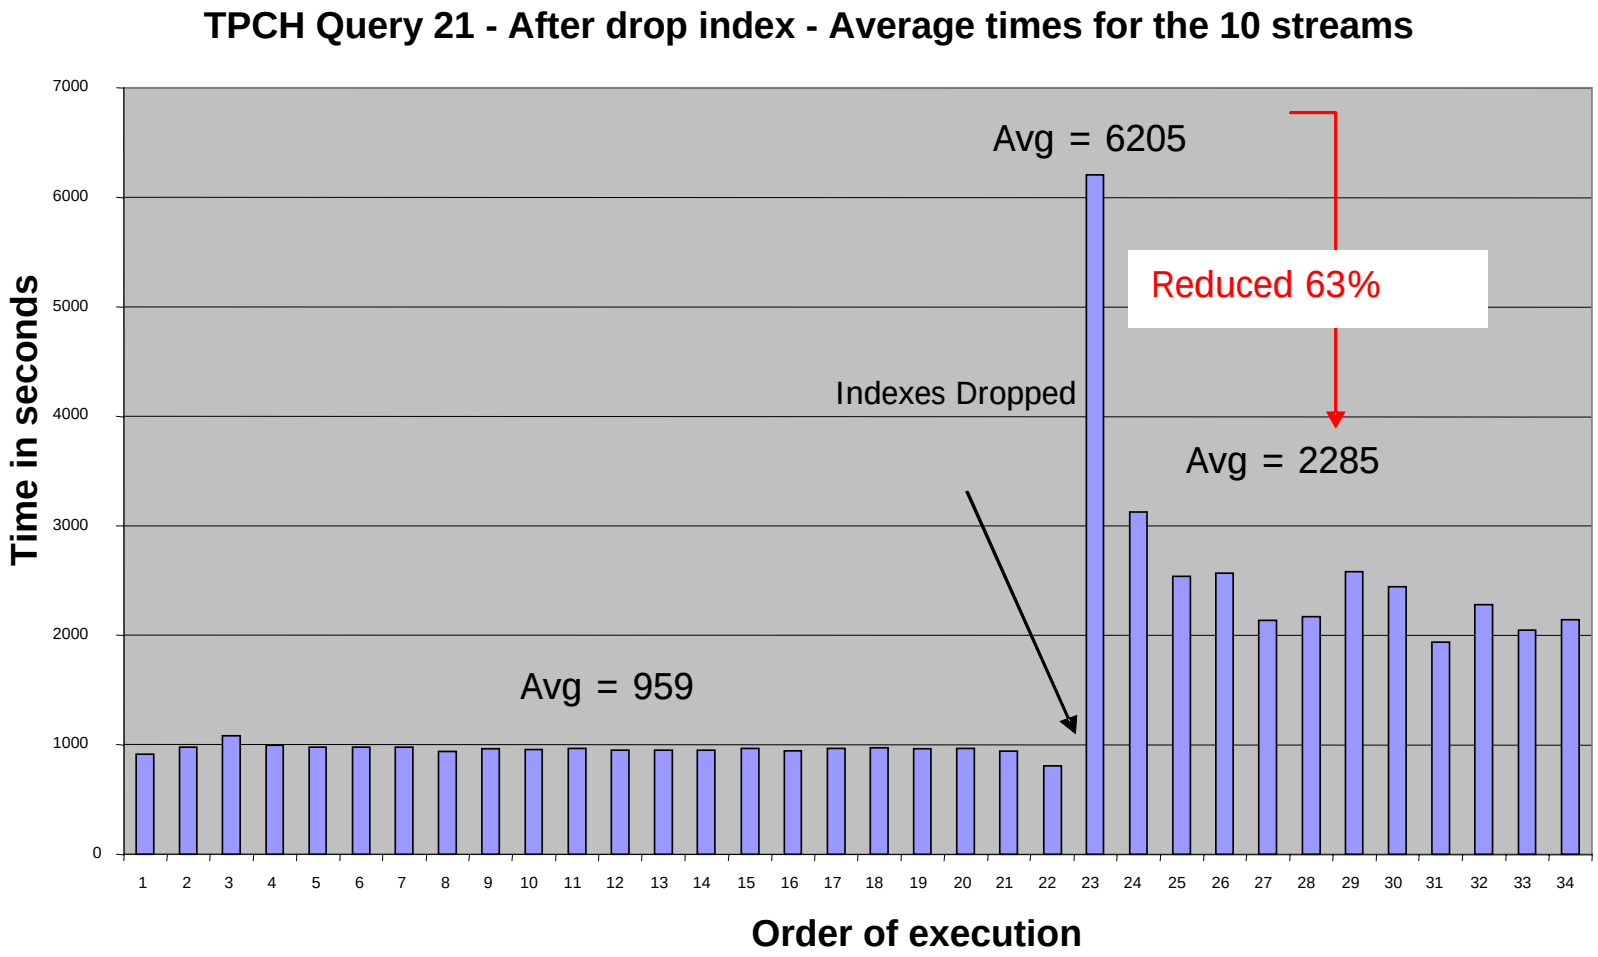

STMM in Action – Comparing Different Configurations

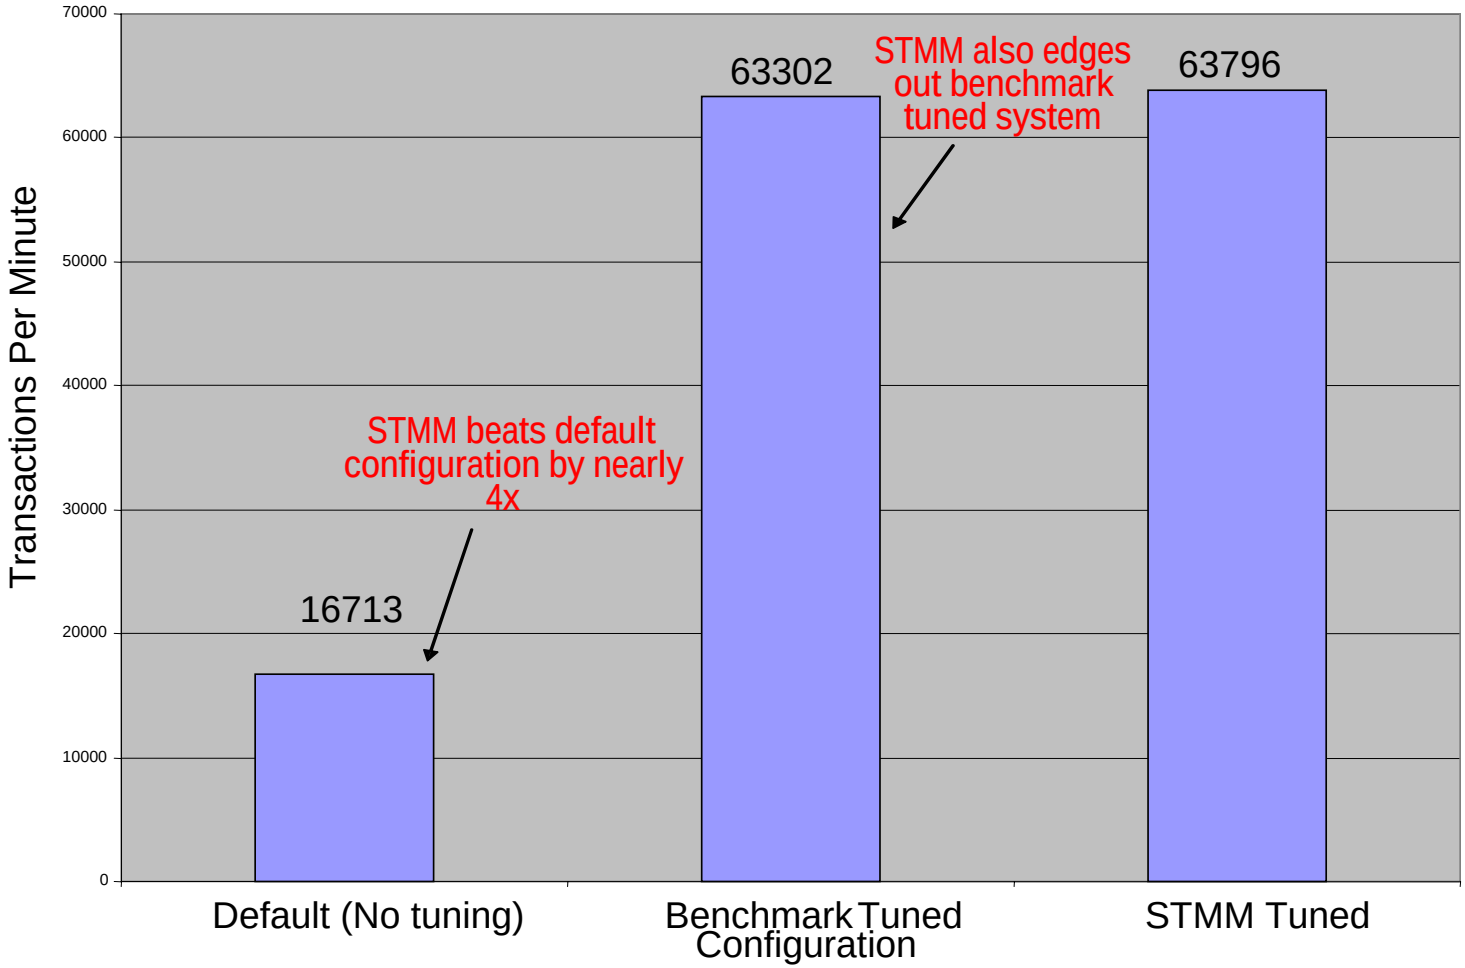

DB2 Automatic Storage Provisioning

- User specifies a group of storage devices for DB2, DB2 allocates and grows table consumption of storage on demand.
  - New to the “Saturn” release of DB2
  - Intended as a “single point of storage management” for table spaces
  - Create a database and associate a set of storage paths with it
- AUTOMATIC STORAGE table spaces
  - No explicit container definitions are provided
  - Containers automatically created across the storage paths
  - Growth of existing containers and addition of new ones completely managed by DB2
- Built around DMS storage model
- Add storage paths to the database afterwards
- Redefine those storage paths during a database RESTORE

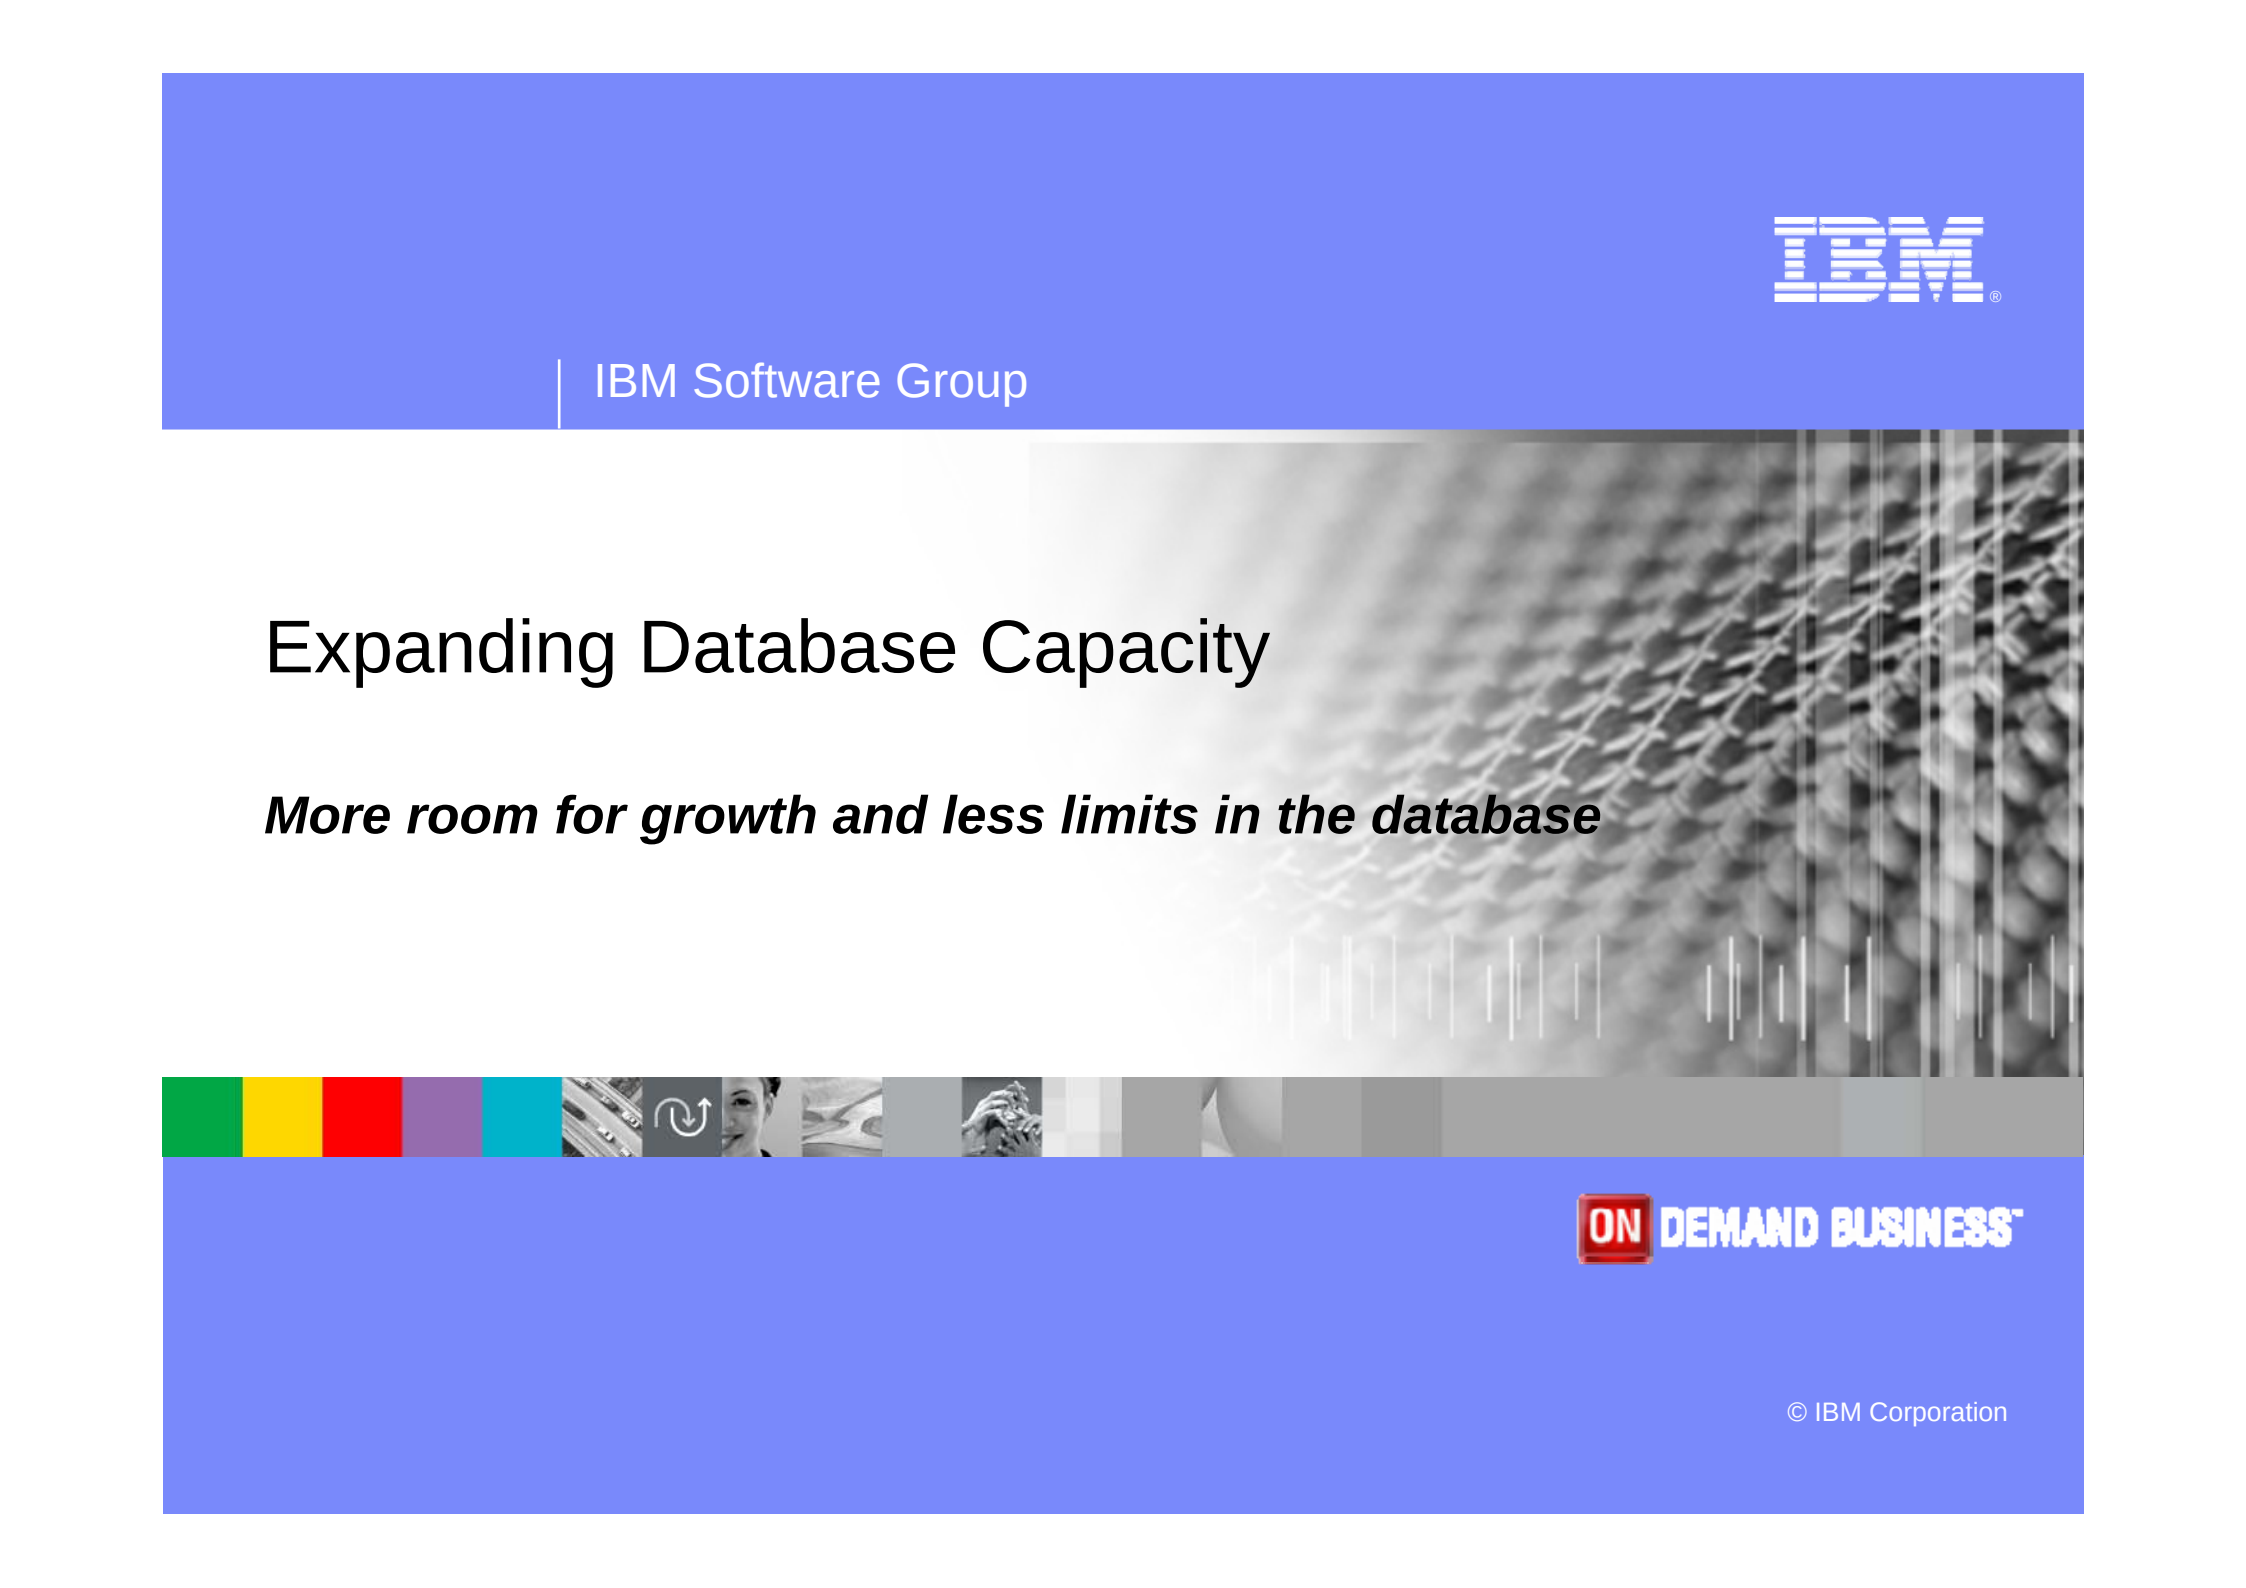
IBM
IBM Software Group
Expanding Database Capacity
More room for growth and less limits in the database
ON DEMAND BUSINESS
© IBM Corporation

Table Partitioning

- What is Table (Range) Partitioning ?
  - Storing a table in more than one physical object, across one or more table spaces
  - Each table space contains a range of the data that can be found very efficiently
- Why?
  - Increase table capacity limit
  - Increase large table manageability
  - Improve SQL performance through partition elimination
  - Provide fast & online data roll-in and roll-out
  - Converge towards Informix functionality
  - Family compatibility with DB2 on zOS and IDS

Table Partitioning : Benefits

Without Partitioning

Without Partitioning

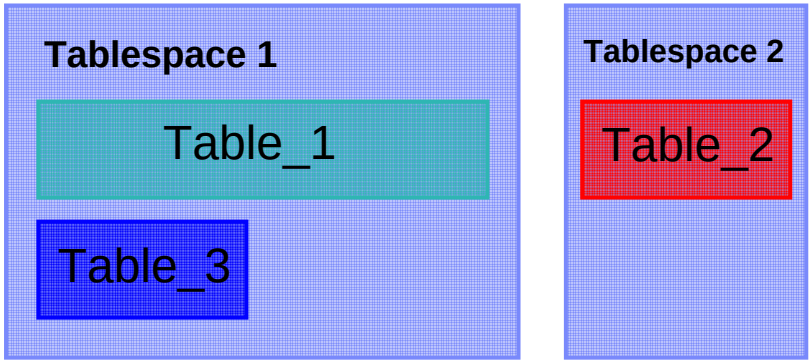

With Partitioning

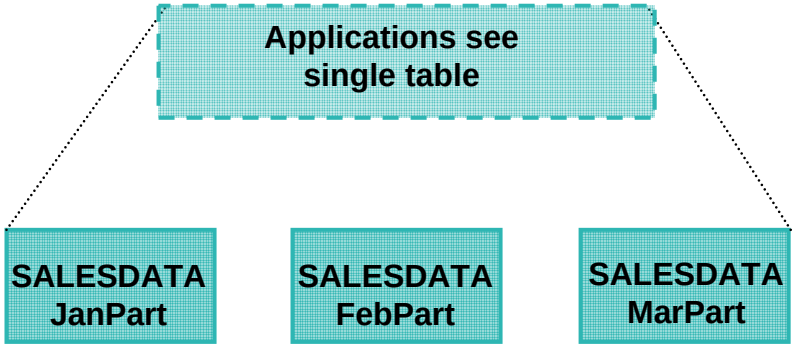

With Partitioning

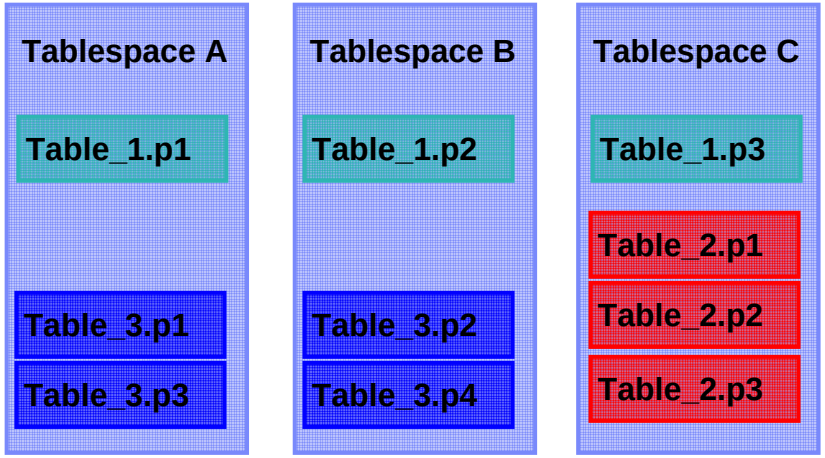

Table Partitioning

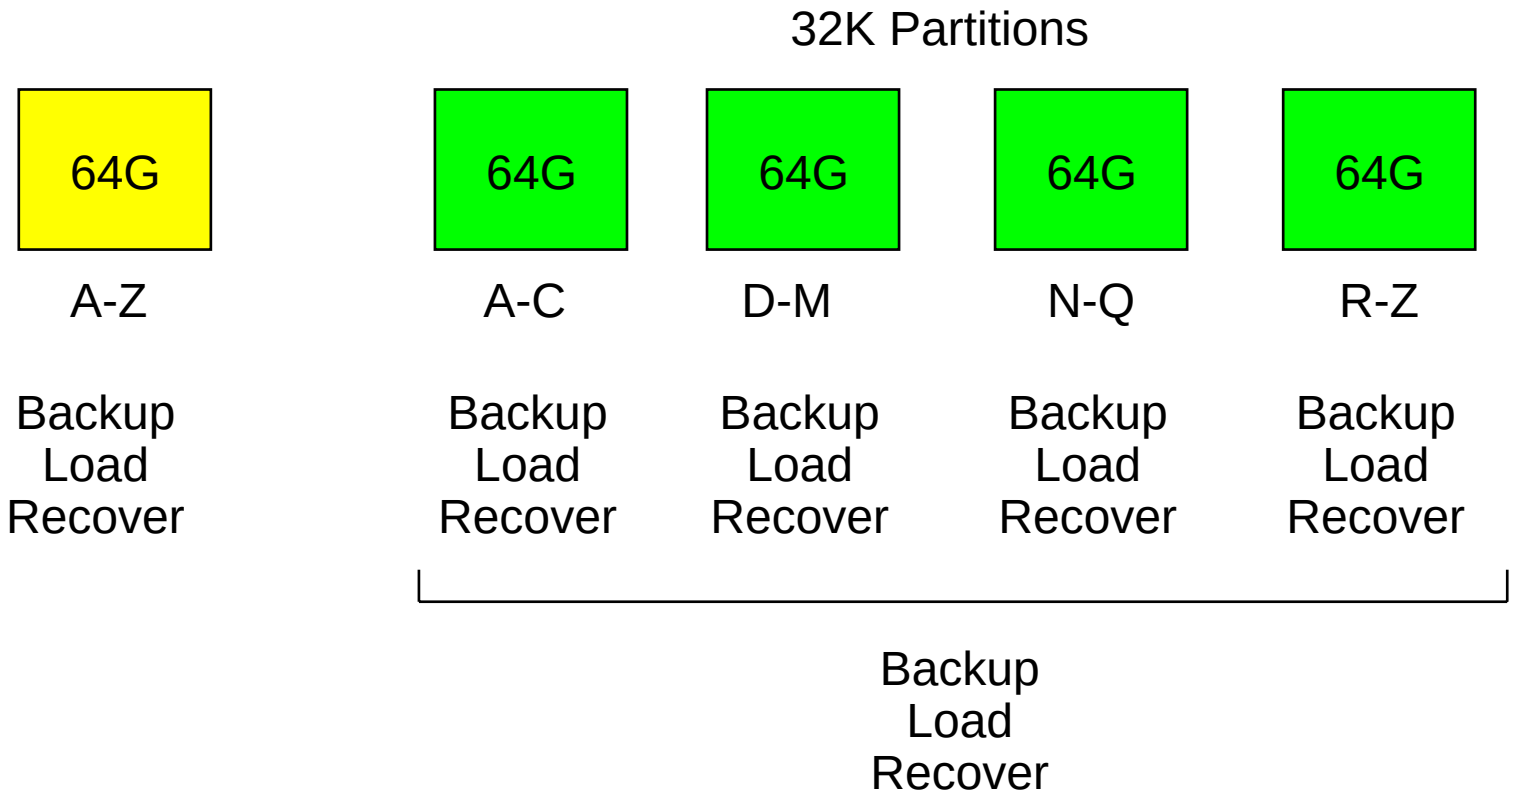

Table Partitioning

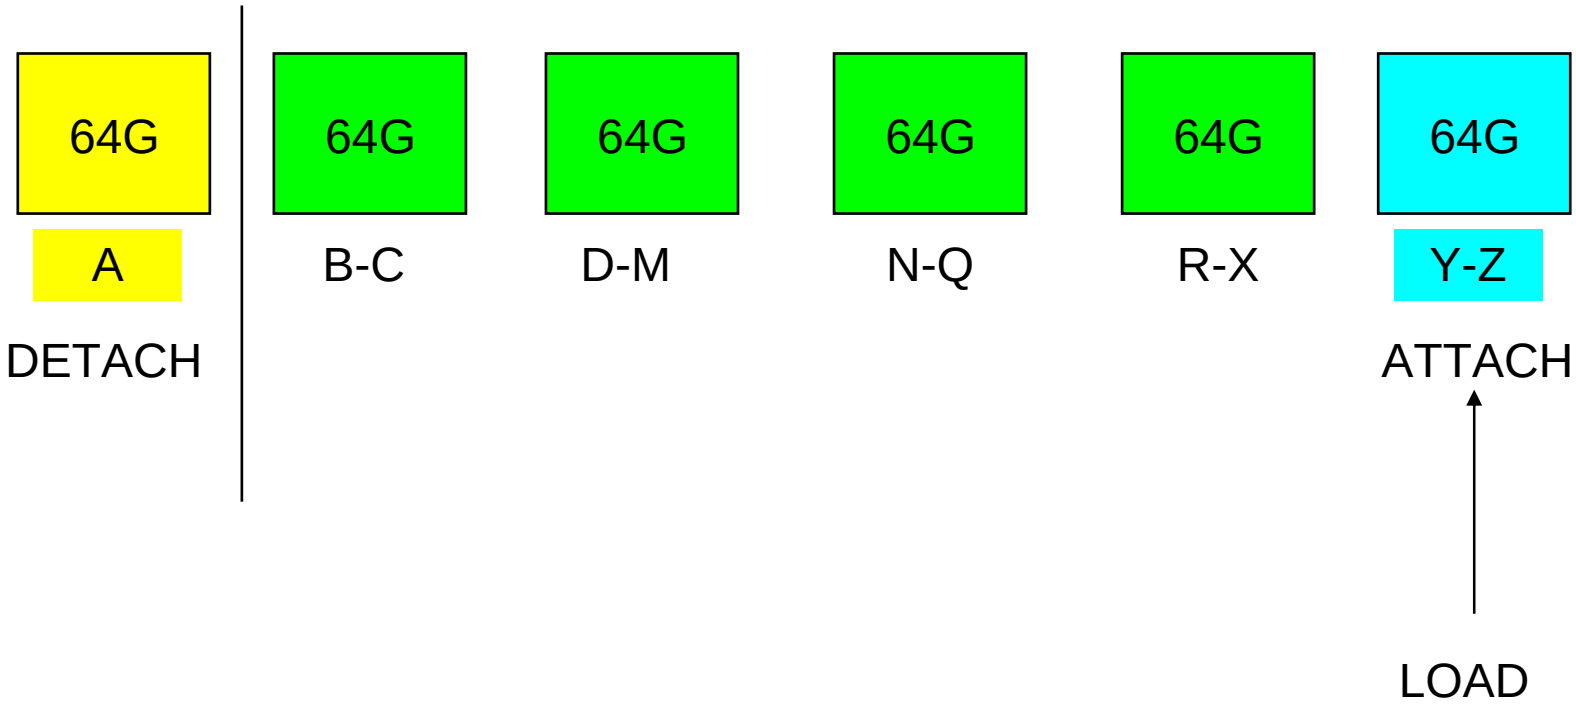

Storage Mapping: Indexes are Global in DB2 9

- Indexes are ***global***
- Each index is in a separate storage object
  - By default, in the same tablespace as the first data partition
  - Can be created in different tablespaces, via
    - INDEX IN clause on CREATE TABLE (default is tablespace of first partition)
    - New IN clause on CREATE INDEX
- Recommendation
  - Place indexes in LARGE tablespaces

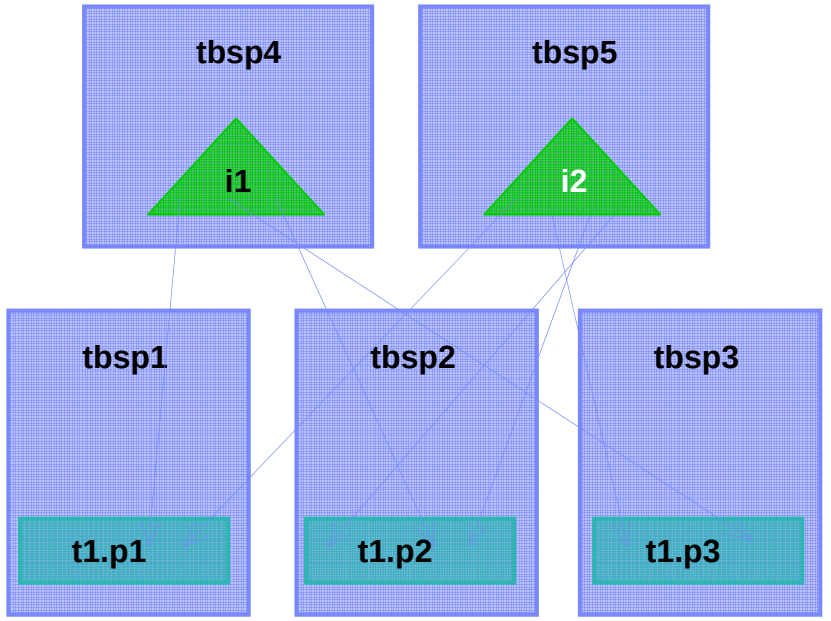


```
CREATE TABLE t1(c1 INT, c2 INT, …)
 IN tbsp1, tbsp2, tbsp3
 INDEX IN tbsp4
 PARTITION BY RANGE(a)
 (STARTING FROM (1) ENDING (100)
   EVERY (33))
CREATE INDEX i1(c1)
CREATE INDEX i2 (c2) IN tbsp5
```

Hybrid Partitioning

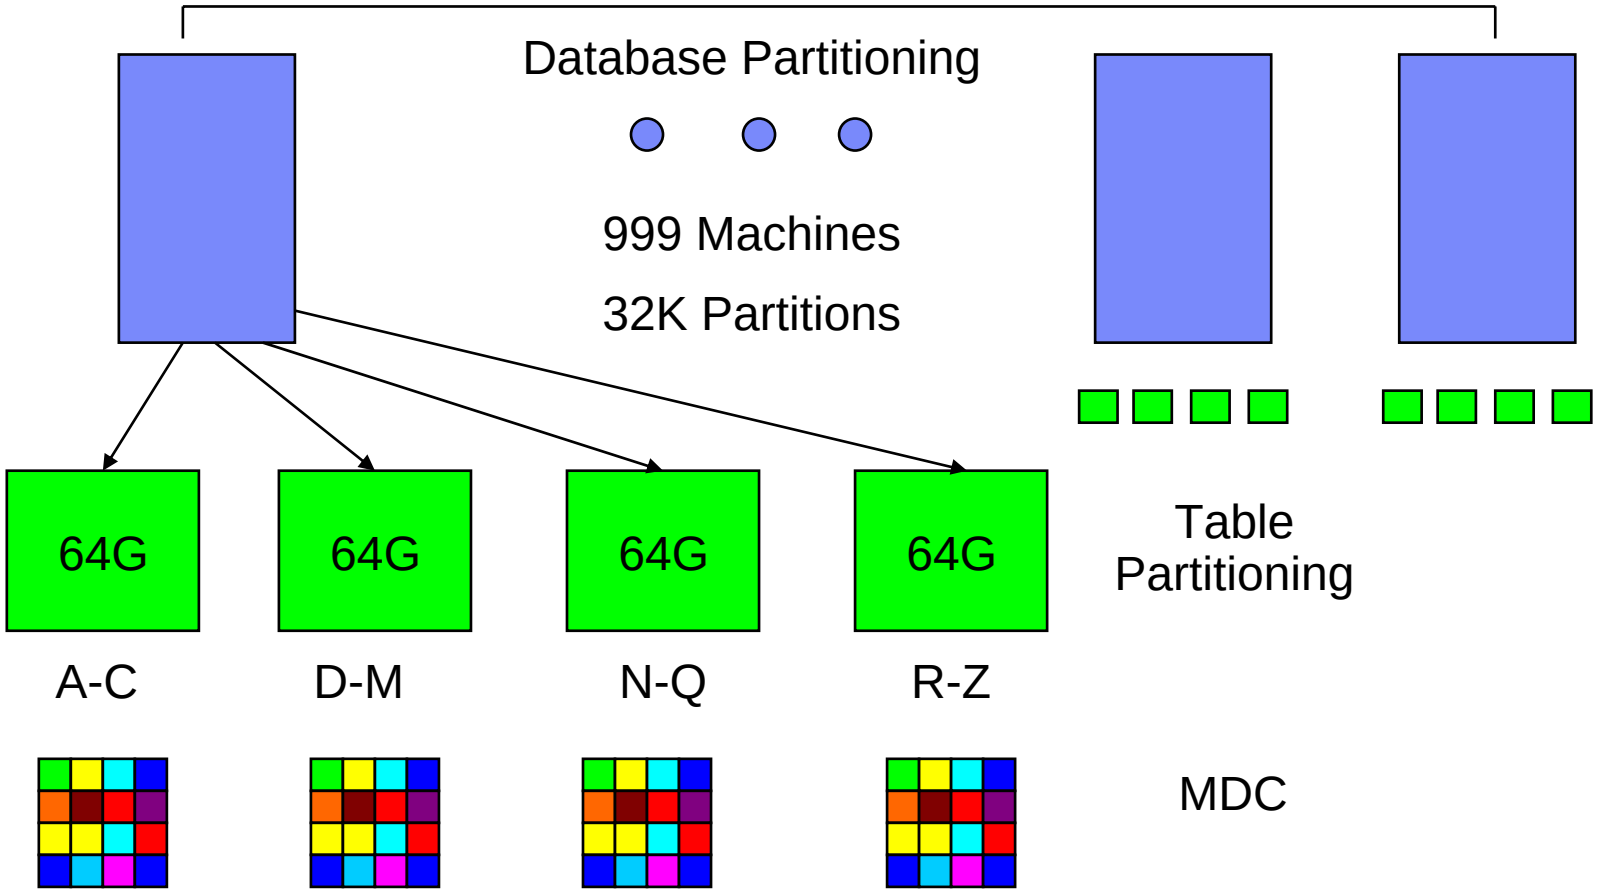

Large Row Identifiers

- Increase In table size limits and rows per page
  - Tablespace level definition
  - DMS Tablespace only
- ALTER TABLESPACE <name> CONVERT TO LARGE
  - Tablespace is locked, definition is modified and catalogues are updated
  - Indexes will need to be reorganized
    - Every index for every table in the converted tablespace needs to be reorganized or rebuilt to convert the RID entries from regular to large

Previous Table Space Design

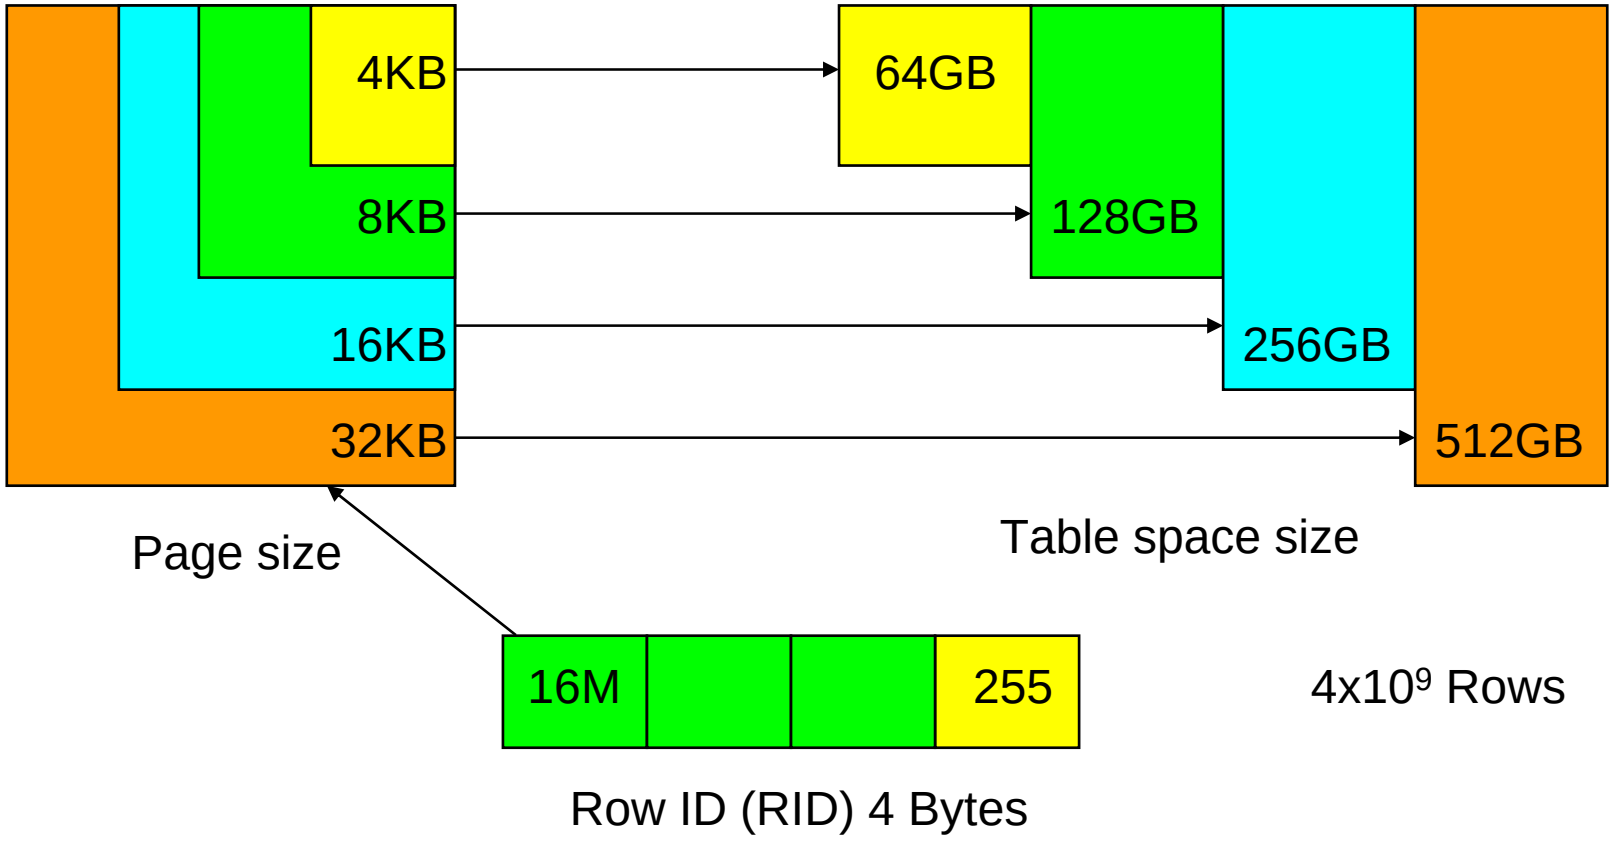

For tables in all table spaces (regular, temporary, DMS, SMS)

New “Large Table Space” Design

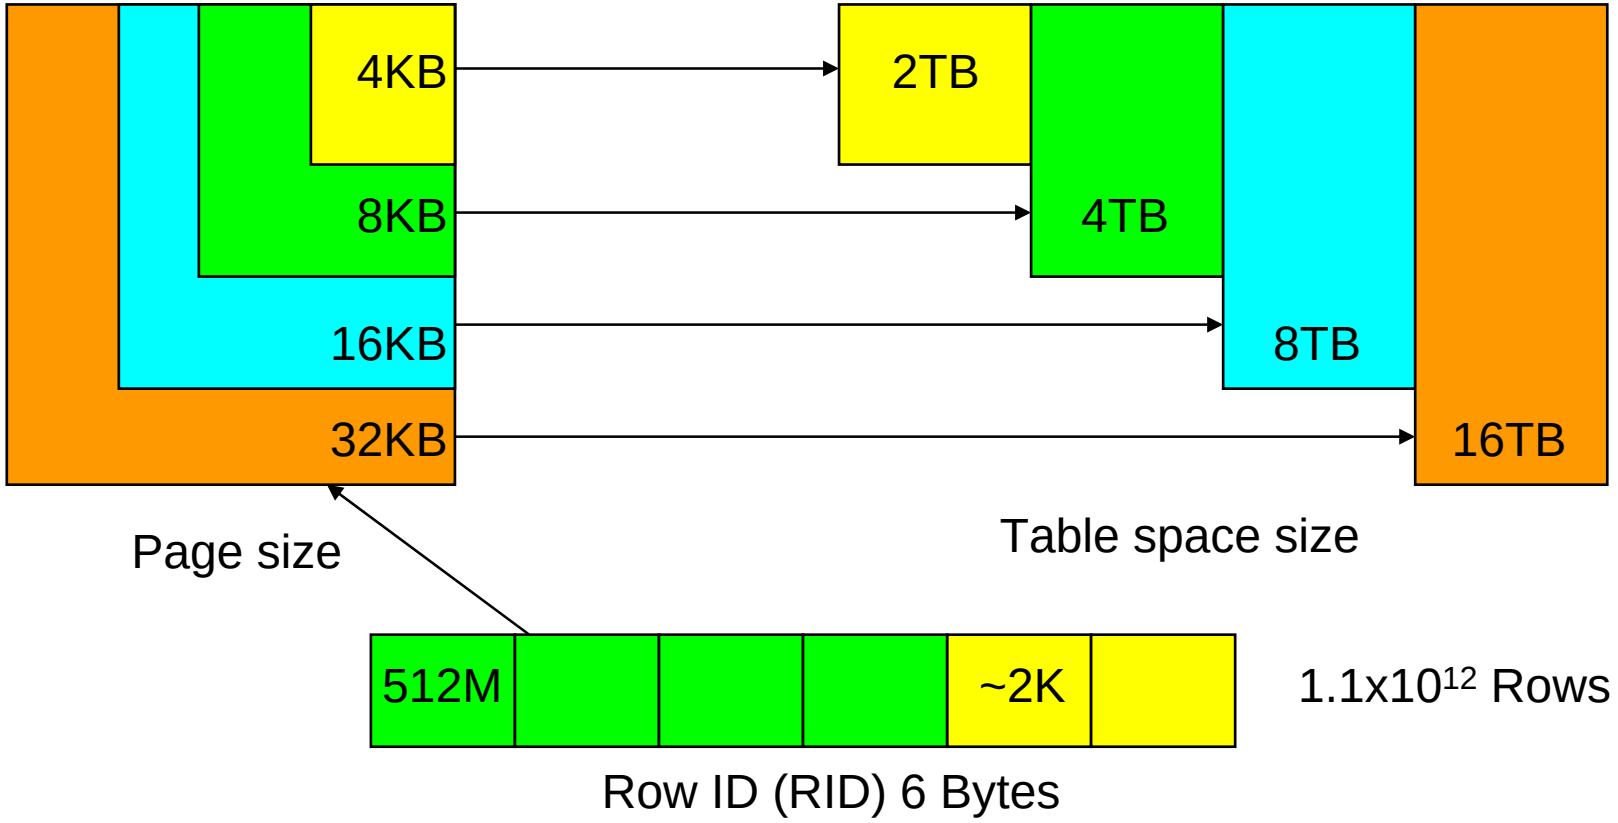

For tables in LARGE table spaces (DMS only)
Also all SYSTEM and USER temporary table spaces

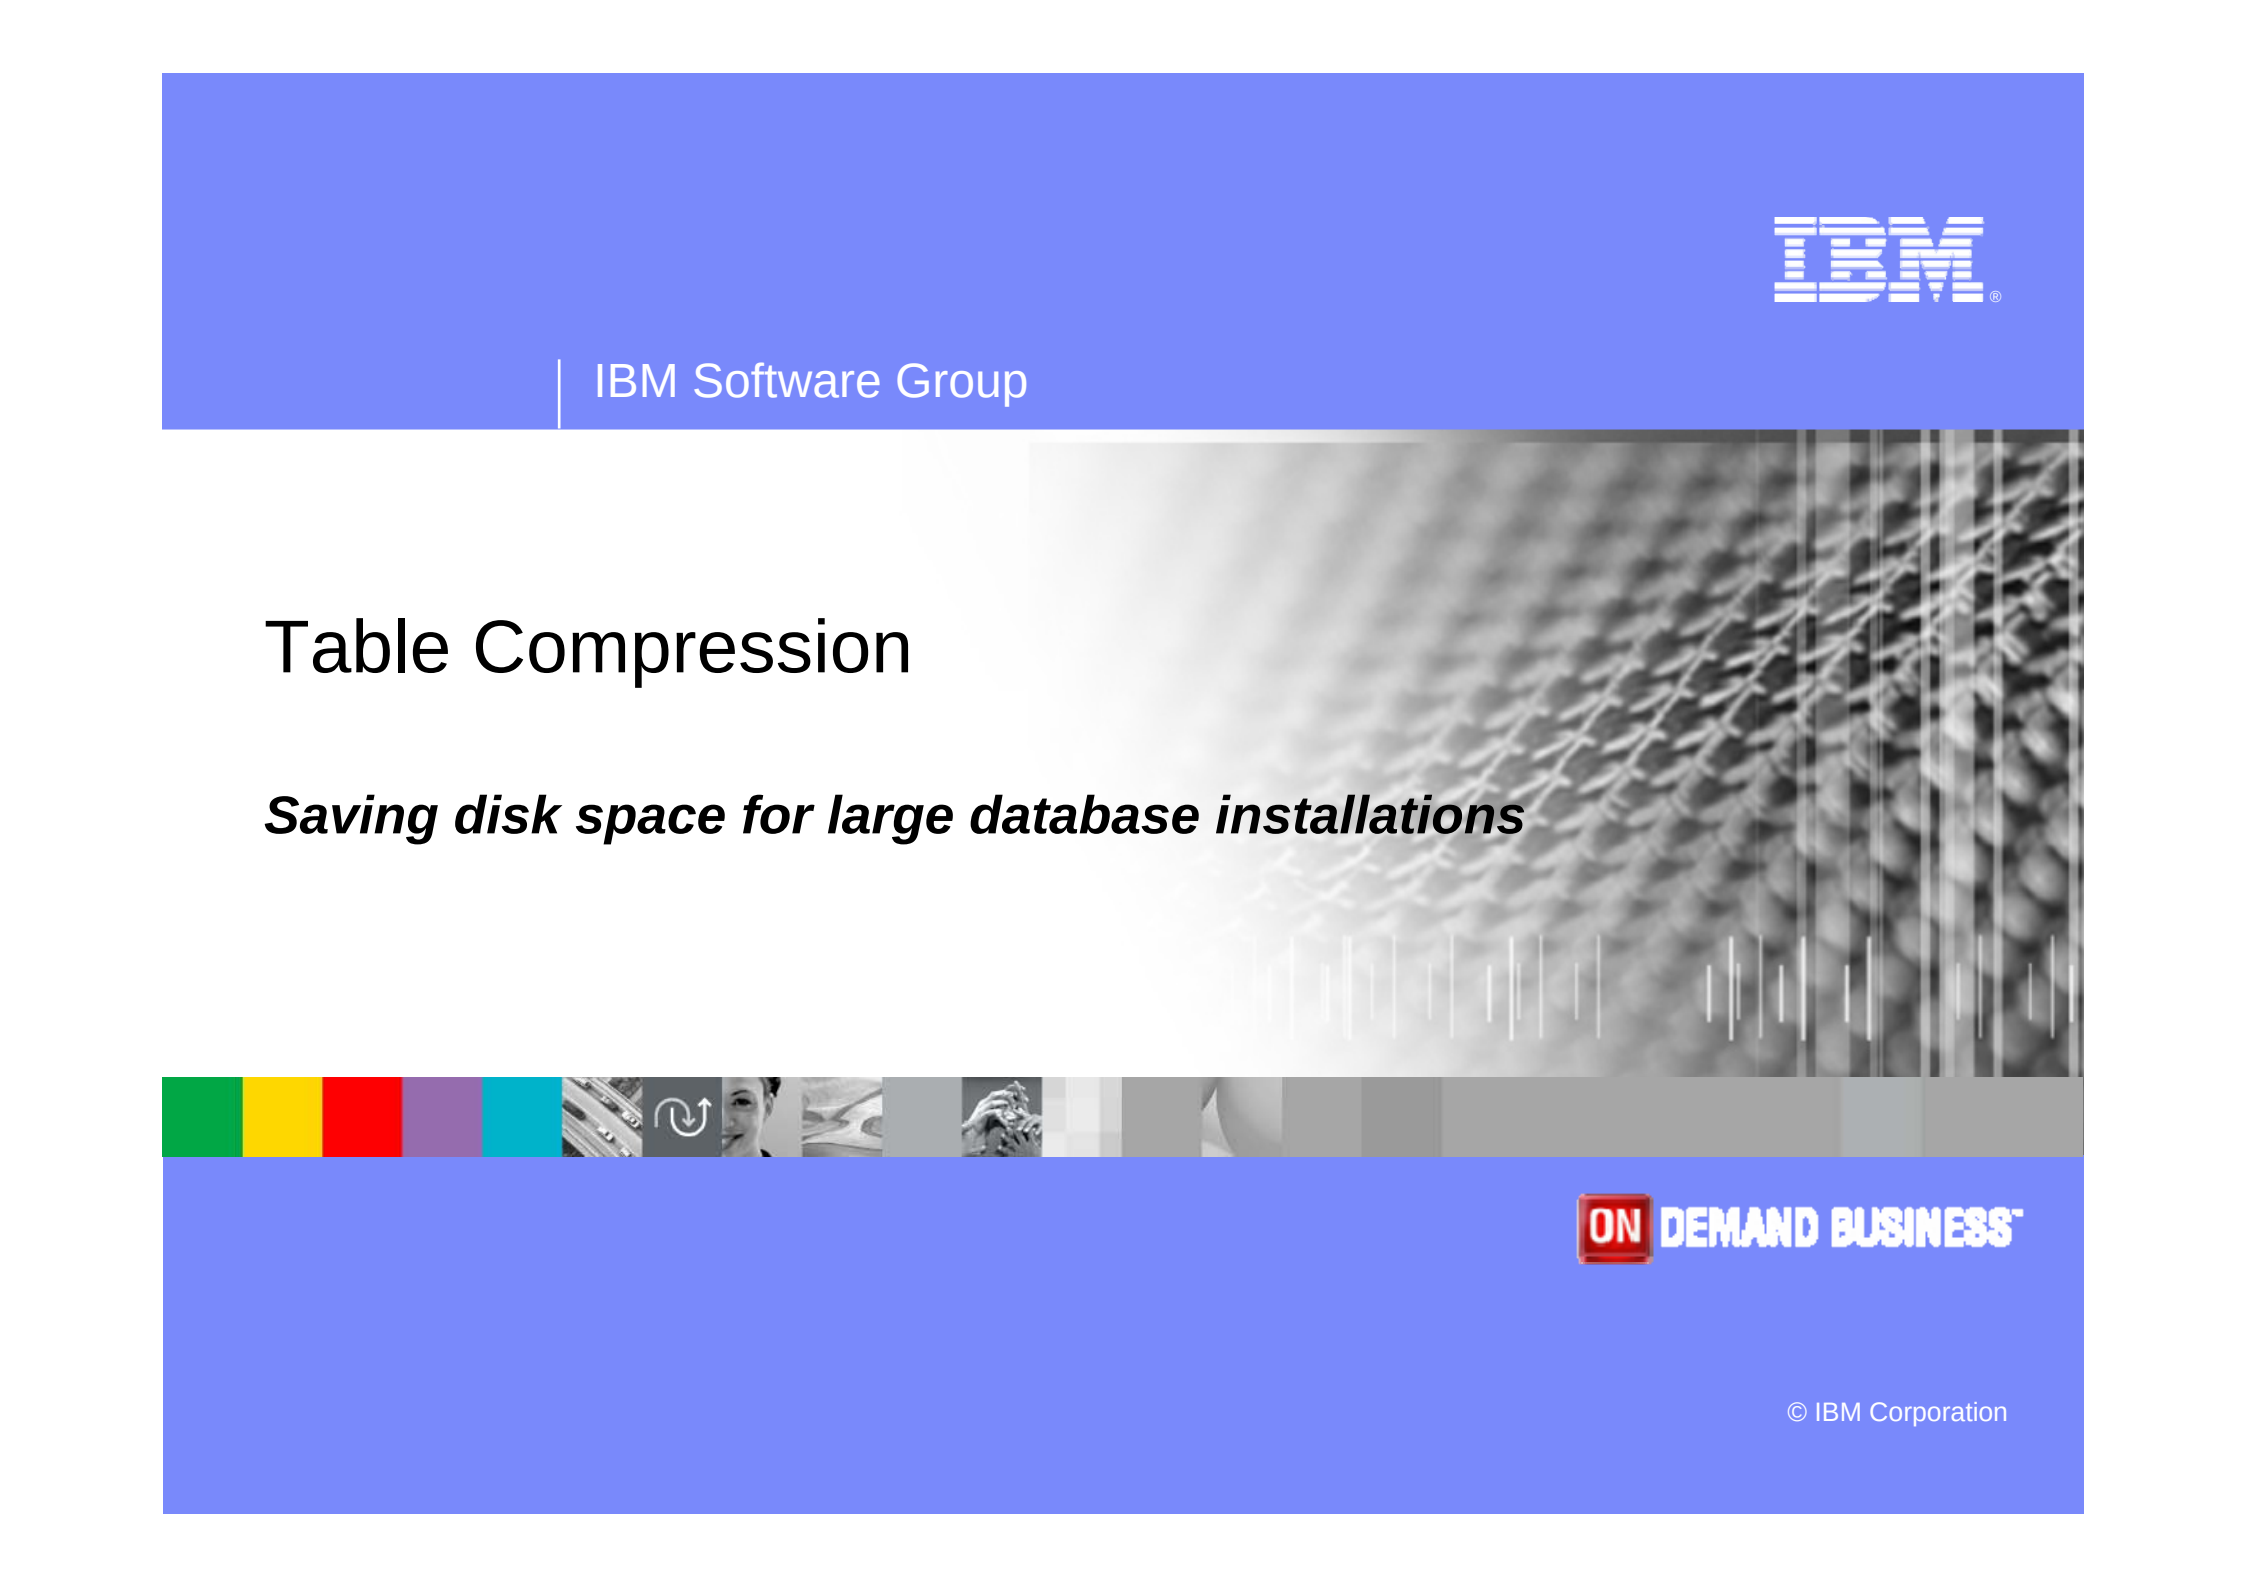

DB2 Compression

- NULL and Default Value Compression (V8 GA)
  - No disk storage consumed for NULL column values, zero length data in variable length columns and system default values
- Multidimensional Clustering (V8 GA)
  - Significant index compression can be achieved through block indexes
    - One key per thousands of records (vs one key per record with traditional indexes)
- Database Backup Compression (V8 FP4)
  - Smaller backup images; compress index and lf/lob tablespaces
- Data Row Compression (DB2 9)

Row Compression Using a Compression Dictionary

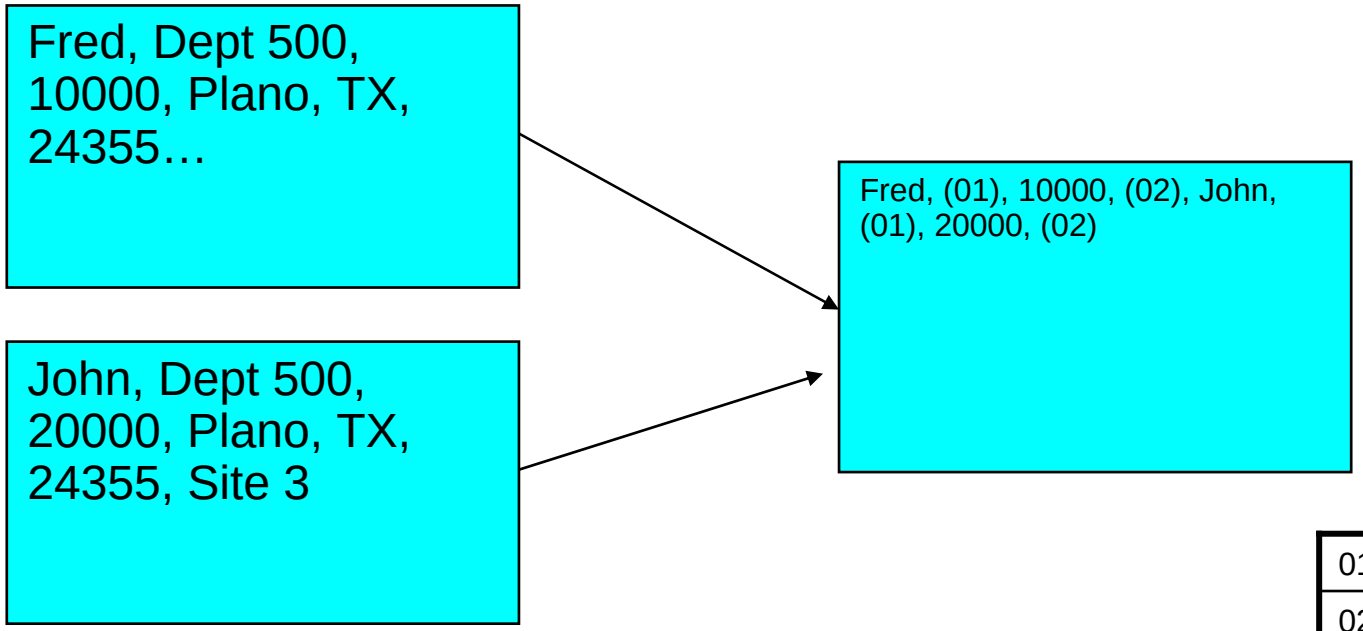

01	Dept 500
02	Plano, TX, 24355
…	…

Side tables contain repeated information from the rows

- Can be across column boundaries or within columns

Compression Ratio – Customer Data

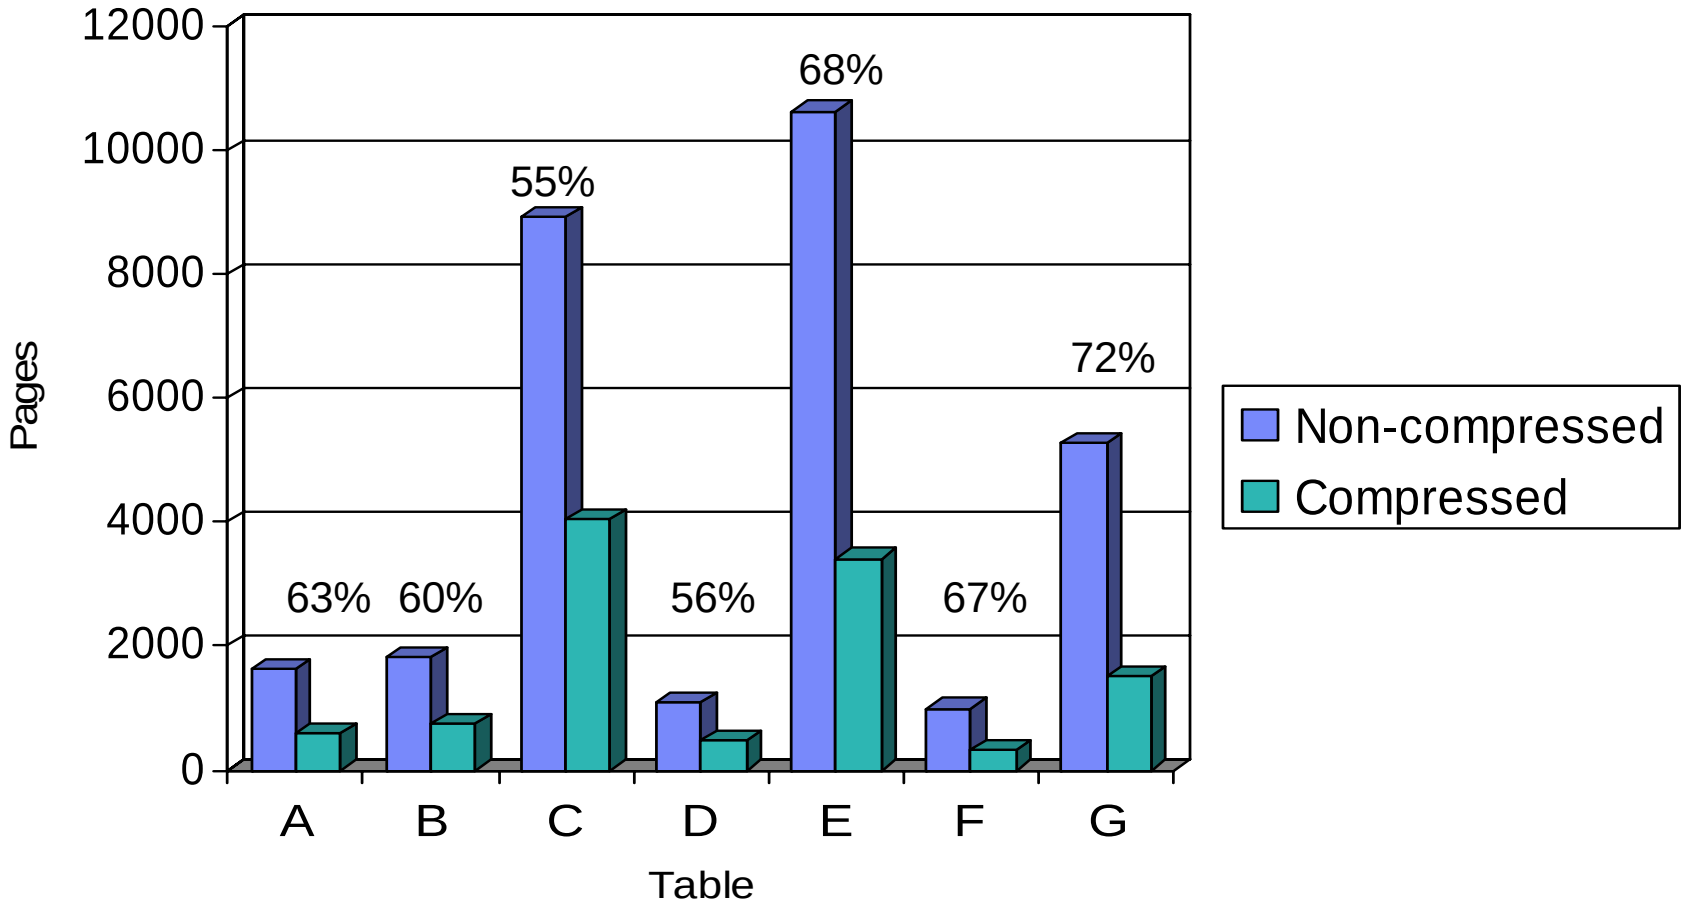

Security - Label Based Access Control

- Label Based Access Control (LBAC)
  - A “label” is associated with both user sessions and data rows or columns
  - Rules for comparing users and data labels provide allow access controls to be applied at the row level
- Labels may consist of multiple components
  - Hierarchical, group or tree types
  - Row labels appear as a single additional column in a protected table, regardless of the number of label components
  - User labels are granted by a security administrator
- Similar to the label security support in DB2 for z/OS v8

LBAC Hierarchy – Tree

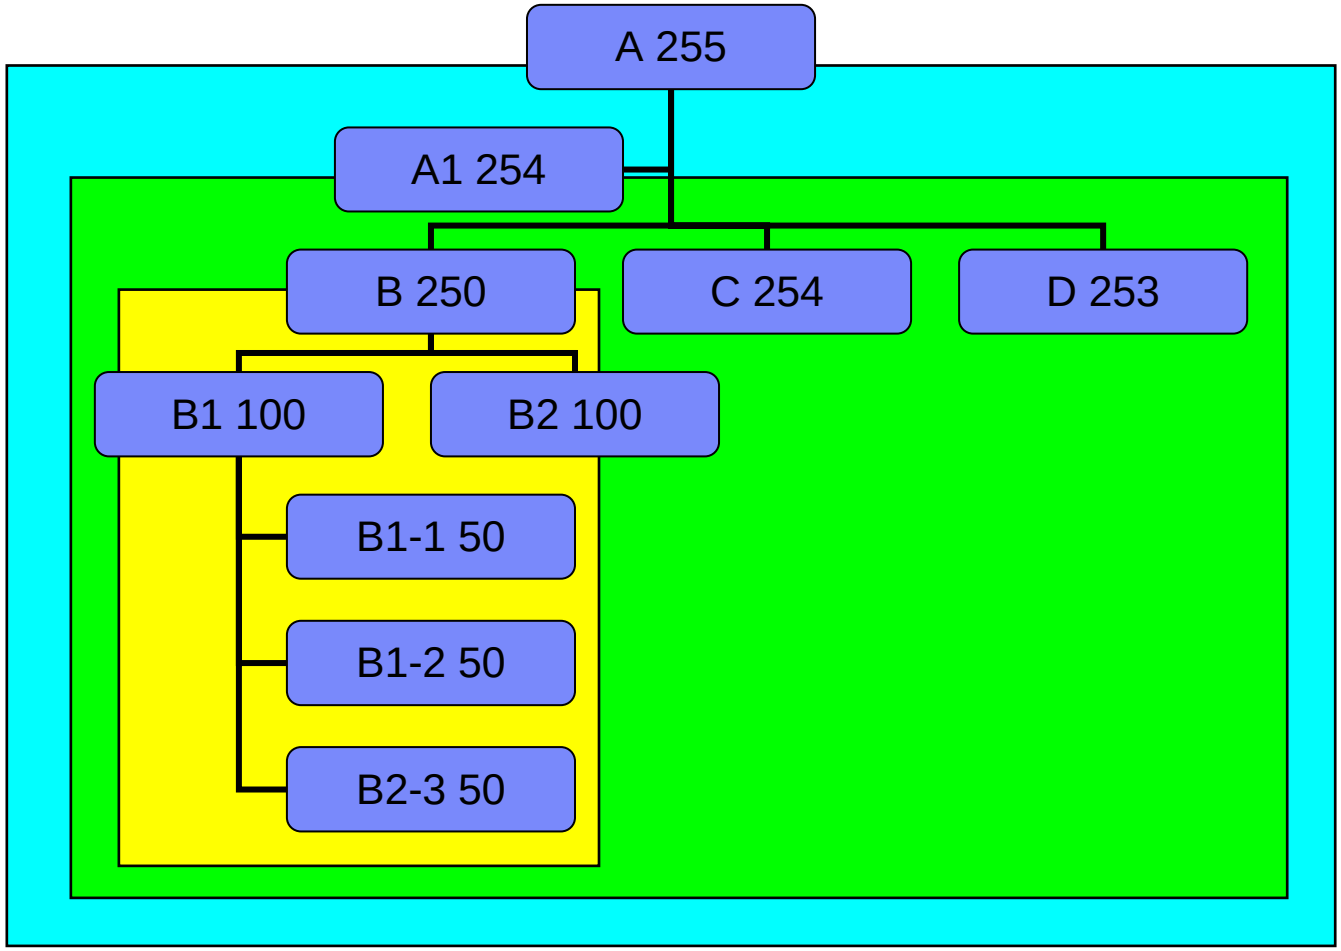

LBAC Query

```
SELECT * FROM EMP
  WHERE
    SALARY >= 50000
```

No LBAC	SEC=254	SEC=100	SEC=50	ID	SALARY
				255	60000
				100	50000
				50	70000
				50	45000
				60	30000
				250	56000
				102	82000
				100	54000
				75	33000
				253	46000
				90	83000
				200	78000

Take Control.

Lancio italiano di DB2 9 (Viper)

"L'annuncio più importante degli ultimi 20 anni"

24 Ottobre 2006
Roma – presso Auditorium della Tecnica, Via Umberto Tupini 65

8 Novembre 2006
Milano – IBM Forum

Perchè partecipare?

- E' il primo database a struttura multipla: DB2 9 è oggi l'unico sul mercato che contenga struttura relazionale e quella per dati XML in modalità nativa
- E' un cambio di generazione dei database, non è il semplice lancio di una nuova versione di prodotto

Per info e registrazione: http://www-5.ibm.com/it/events/db2v9/

IBM Software Group

DB2 9 Bootcamp
Technology Overview

SWIT – Information Management Technical Presales

Michele Benedetti

ON DEMAND BUSINESS