

IBM Open Partnership Center

DB2 Technical Workshop

DB2 v9 Overview

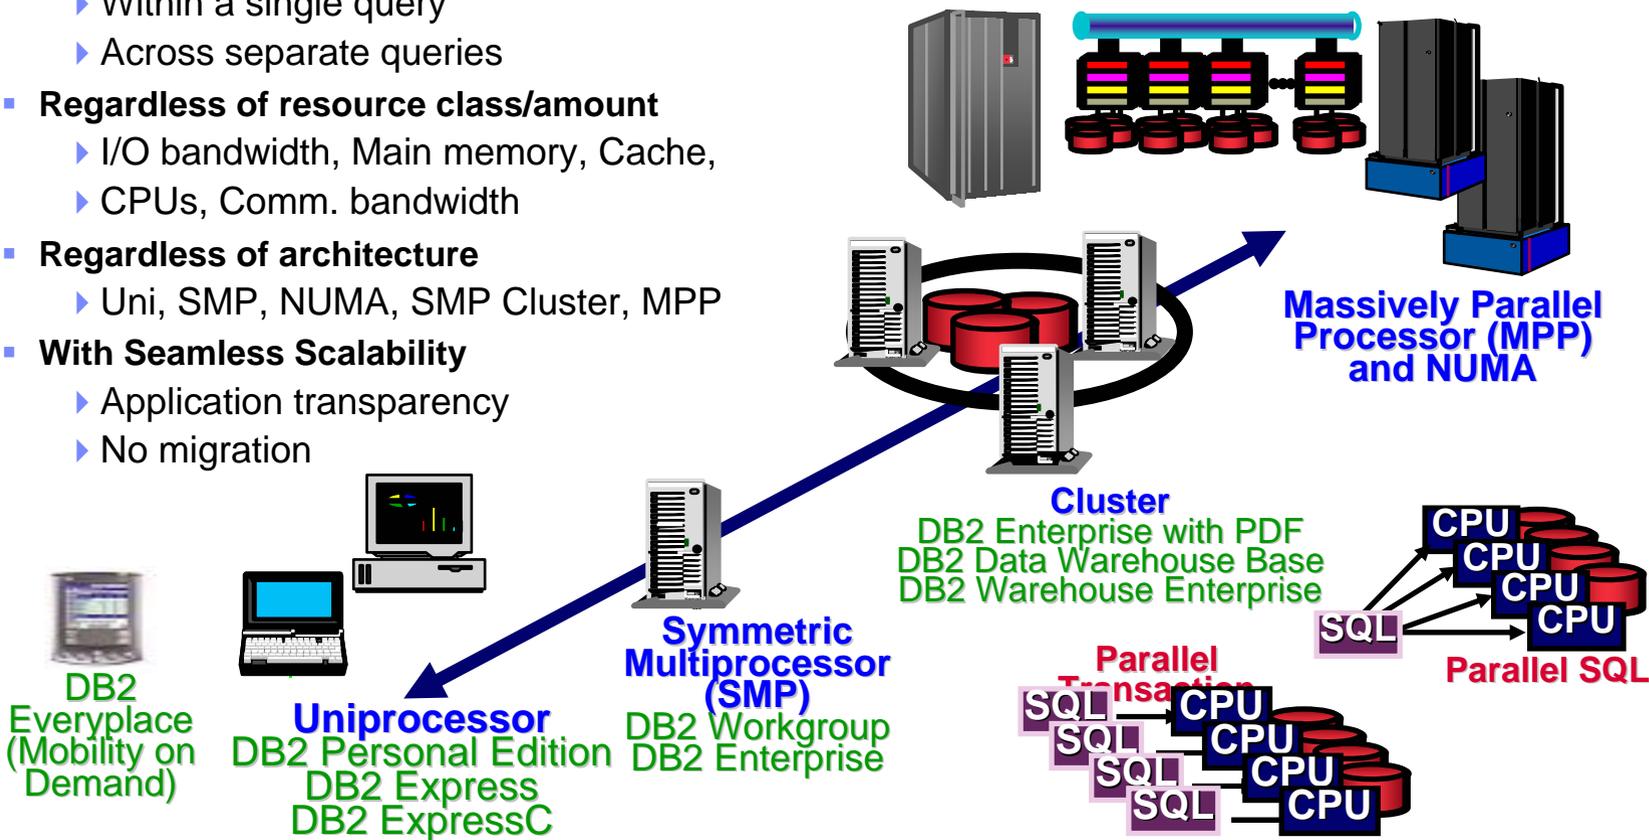
Shu Jian(shujian@cn.ibm.com)
China Linux Competency Center
IBM Open Partnership Center
2007/07/23



ON DEMAND BUSINESS™

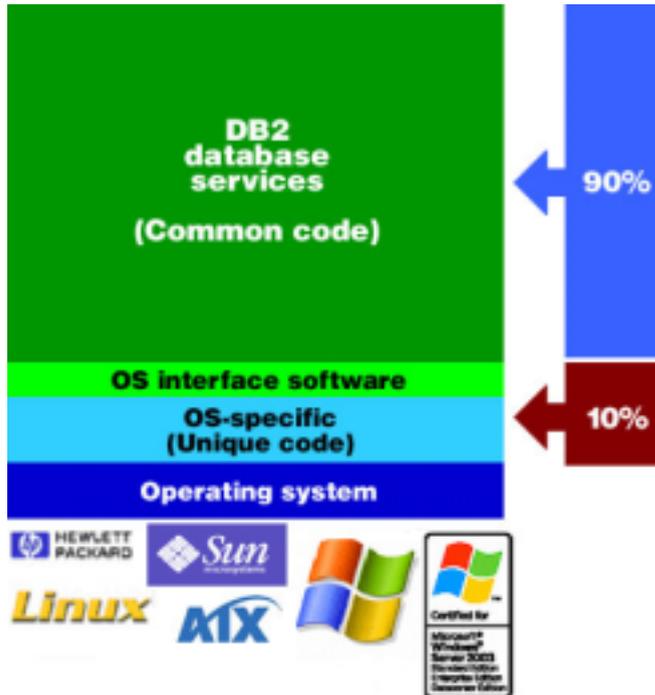
DB2 Product Family

- **Full exploitation of ALL available resources**
 - ▶ Within a single query
 - ▶ Across separate queries
- **Regardless of resource class/amount**
 - ▶ I/O bandwidth, Main memory, Cache,
 - ▶ CPUs, Comm. bandwidth
- **Regardless of architecture**
 - ▶ Uni, SMP, NUMA, SMP Cluster, MPP
- **With Seamless Scalability**
 - ▶ Application transparency
 - ▶ No migration



Same code base for all DB2 server editions

- The same application will work on all DB2 editions with NO modification required



- As business grows, a new DB2 edition may be needed
- Just install the new edition
 - ▶ Your databases will not be deleted (But a backup is always recommended)
 - ▶ Your application will not need to be modified

DB2 Express-C Edition



DB2 Express-C is:

- a version of DB2 Express for the community.
- **Free to develop, deploy, distribute* ...no limits!**

Production Ready

- No database size limit
- No limit on number of instances per server
- No user limit

Innovative Technology

- Same data server technology and programming support as DB2 Express in a smaller package
- DB2 Express-C 9 includes pureXML!

Developer Community

- Dedicated DB2 Express-C community team
- Free online community support
- Skills and applications applicable to all editions of DB2

* Free registration required

DB2 Express-C vs. other free databases

	DB2 Express-C	MySQL Pro	Oracle XE	SQL Server Express
Platforms	Linux, Windows	Linux, Windows and selected UNIX	Linux, Windows	Windows only
Download	390MB	35MB (no tools)	150-190MB	64MB (no tools)
Memory limit	4GB	None	1GB	1GB
32/64 bit	32/64 bit	32	32	32
CPU limit	2 CPU dual core	None	1	1
DB size limit	None	None	4GB	4GB
APIs	C, JDBC, .NET, XML, PHP, Python,etc	C, JDBC, .NET, PHP, Python	C, JDBC, .NET, PHP, XML	.NET
Development Tools	Eclipse, VS .Net, Toad, DB2 DC,...	None	HTML DB	VS .NET
Management	Extensive DB2 Suite (Control Center, Information Center, dozens of wizards, autonomic tools,...)	NONE (tools can be found in online repositories but no INTEGRATION provided by mySQL)	Custom Web Based Tools	SQL Express Manager
Replication	Yes	Yes	Yes	Yes
FreeProduction / Redistribution	YES / YES	NO / NO (commercial grade version)	YES / YES	YES / YES
License Fee	None	None for OSS (but mandatory for commercial grade version)	None	None
Support	Forum only	Forum and fee-based	Forum only	Only with upgrade

Application Development Freedom

- C/C++ (ODBC and Static SQL)
- JDBC and SQLJ
- Borland
- Python
- PHP
- Perl
- .NET languages
- OLE-DB
- ADO
- Web Services
- SQL
- MS Office: Excel, Access, Word



Technical Highlight

- XML Support
- Autonomic Enhancements and Administration Improvements
- Table Partitioning
- Larger Table spaces
- Label-based Access Control
- Row Compression
- SQL Enhancements



XML Integration

■ XML Technology

- ▶ XML = Extensible Markup Language

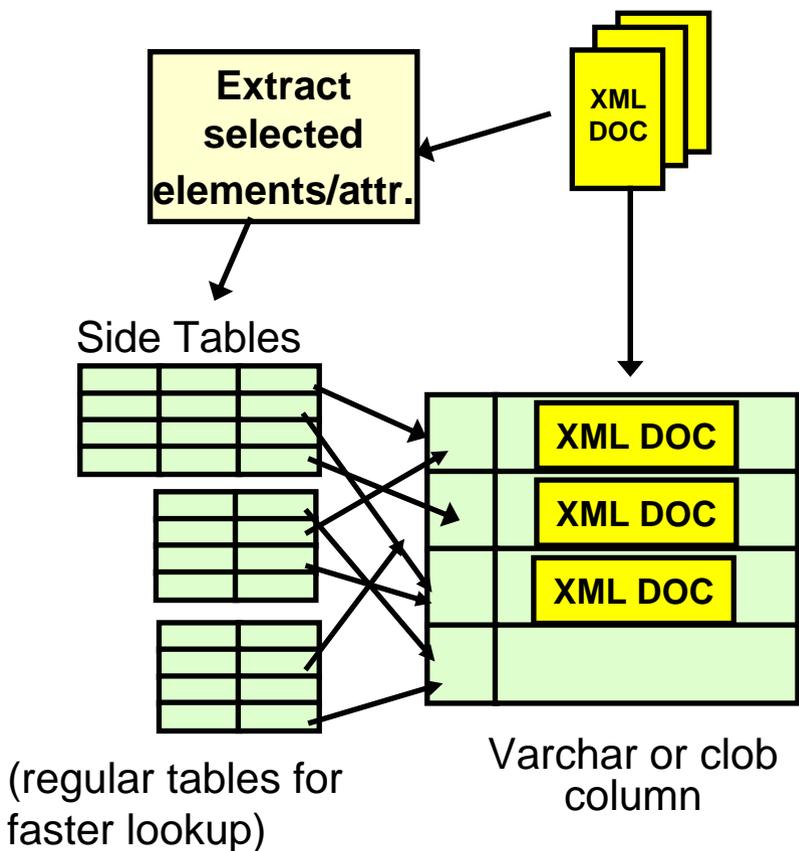
■ Benefits

- ▶ 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

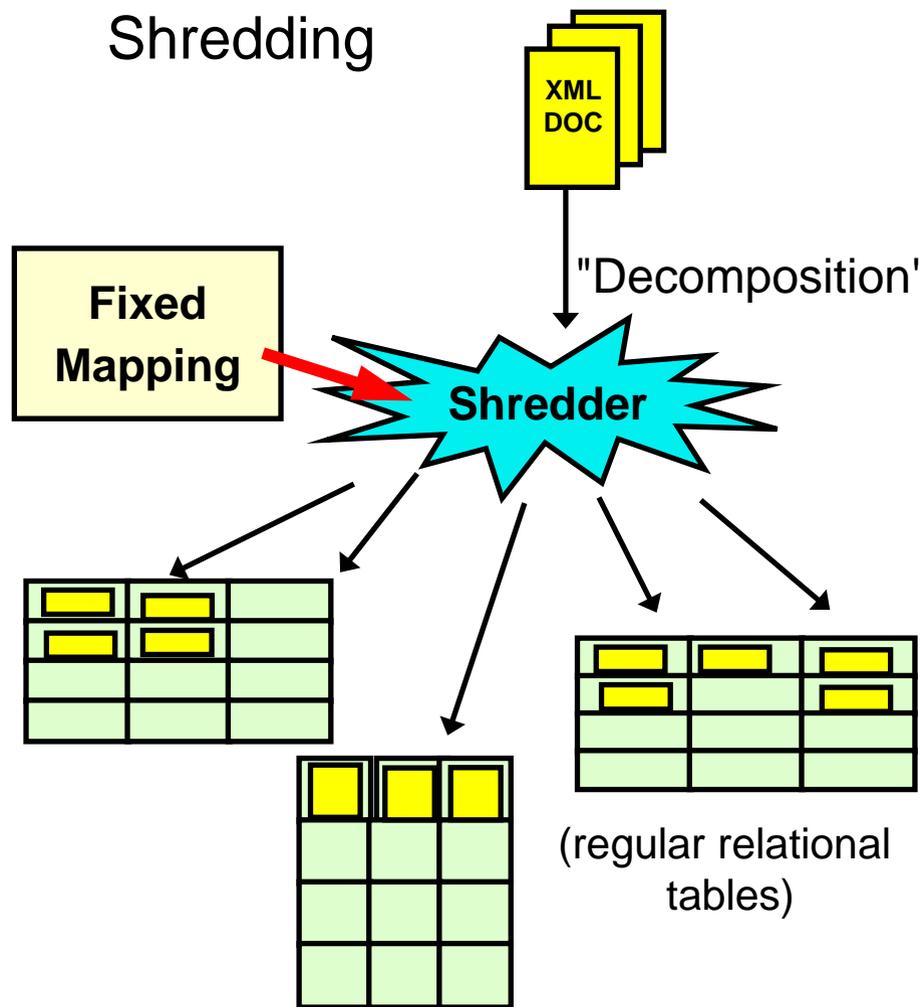
```
<? xml version="1.0" ?>
<purchaseOrder id='12345' secretKey='4x%$^'>
  <customer id="A6789">
    <name>John Smith Co</name>
    <address>
      <street>1234 W. Main St</street>
      <city>Toledo</city>
      <state>OH</state>
      <zip>95141</zip>
    </address>
  </customer>
  <itemList>
    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

XML-Enabled Databases: Two Main Options

CLOB/Varchar



Shredding



Make Changes Easily with DB2 9

```

<DEPARTMENT deptid="15" deptname="Sales">
  <EMPLOYEE>
    <EMPNO>10</EMPNO>
    <FIRSTNAME>CHRISTINE</FIRSTNAME>
    <LASTNAME>SMITH</LASTNAME>
    <PHONE>408-463-4963</PHONE>
    <PHONE>415-010-1234</PHONE>
    <SALARY>52750.00</SALARY>
  </EMPLOYEE>
  <EMPLOYEE>
    <EMPNO>27</EMPNO>
    <FIRSTNAME>MICHAEL</FIRSTNAME>
    <LASTNAME>THOMPSON</LASTNAME>
    <PHONE>406-463-1234</PHONE>
    <SALARY>41250.00</SALARY>
  </EMPLOYEE>
</DEPARTMENT>

```

Employee

DEPTID	EMPNO	FIRSTNAME	LASTNAME	PHONE	SALARY
15	27	MICHAEL	THOMPSON	406-463-1234	41250
15	10	CHRISTINE	SMITH	408-463-4963	52750

Requires:

- Normalization of existing data !
- Modification of the mapping
- Change of applications

Phone

EMPNO	PHONE
27	406-463-1234
10	415-010-1234
10	408-463-4963

Department

DEPTID	DEPTNAME
15	Sales

Costly!

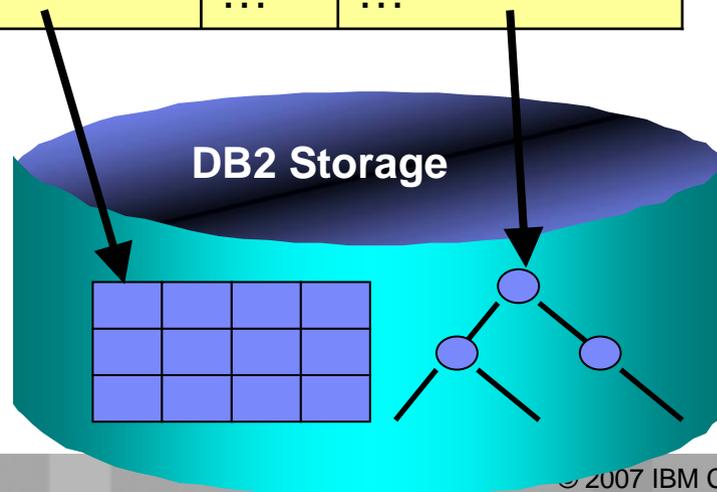
Native XML Storage

- DB2 stores XML in **parsed hierarchical** format (~DOM)

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

- Relational columns are stored in relational format (tables)
- XML columns are stored **natively**
- No XML parsing for query evaluation!**

deptID	...	deptdoc
"PR27"	...	<dept> ... <emp>...</emp> </dept>
...



Integration of XML & Relational Capabilities

**Table "PEOPLE"
Column "INFO"**

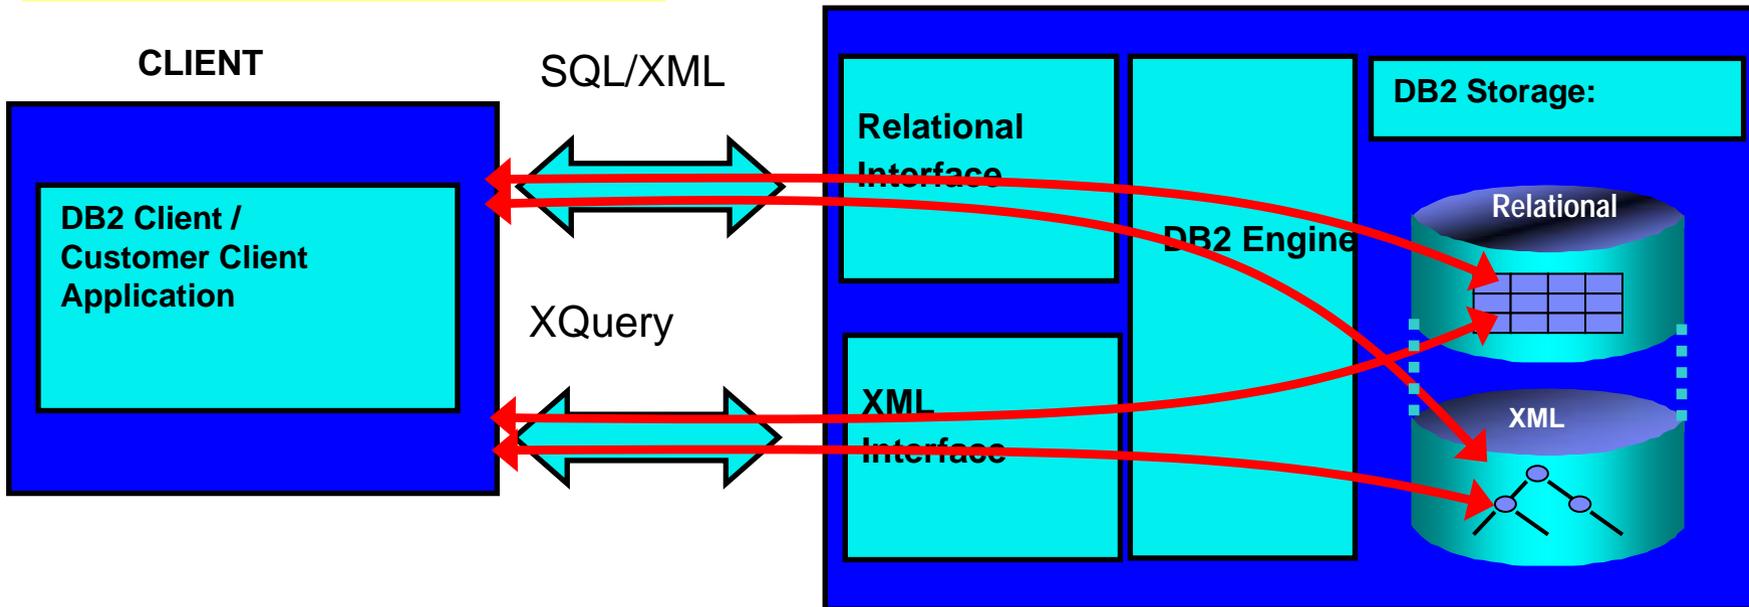
- Applications combine XML & relational data
- Native XML data type (server & client side)
- XML Capabilities in all DB2 components

```
select xmlquery
('$ColName/dept/employee[@id="901"]
/name/office' passing info as "ColName")
from dept;
```

```
<dept bldg="101">
  <employee id="901">
    <name>John Doe</name>
    <phone>408 555 1212</phone>
    <office>344</office>
  </employee>
  <employee id="902">
    <name>Peter Pan</name>
    <phone>408 555 9918</phone>
    <office>216</office>
  </employee>
</dept>
```

Answer

```
<office>344</office>
```

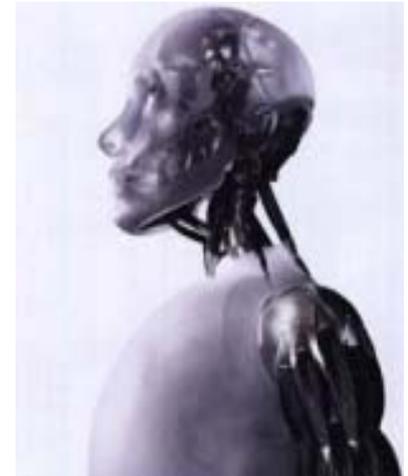


Improved Manageability

Less Skill, Less Resource, Better System Utilization

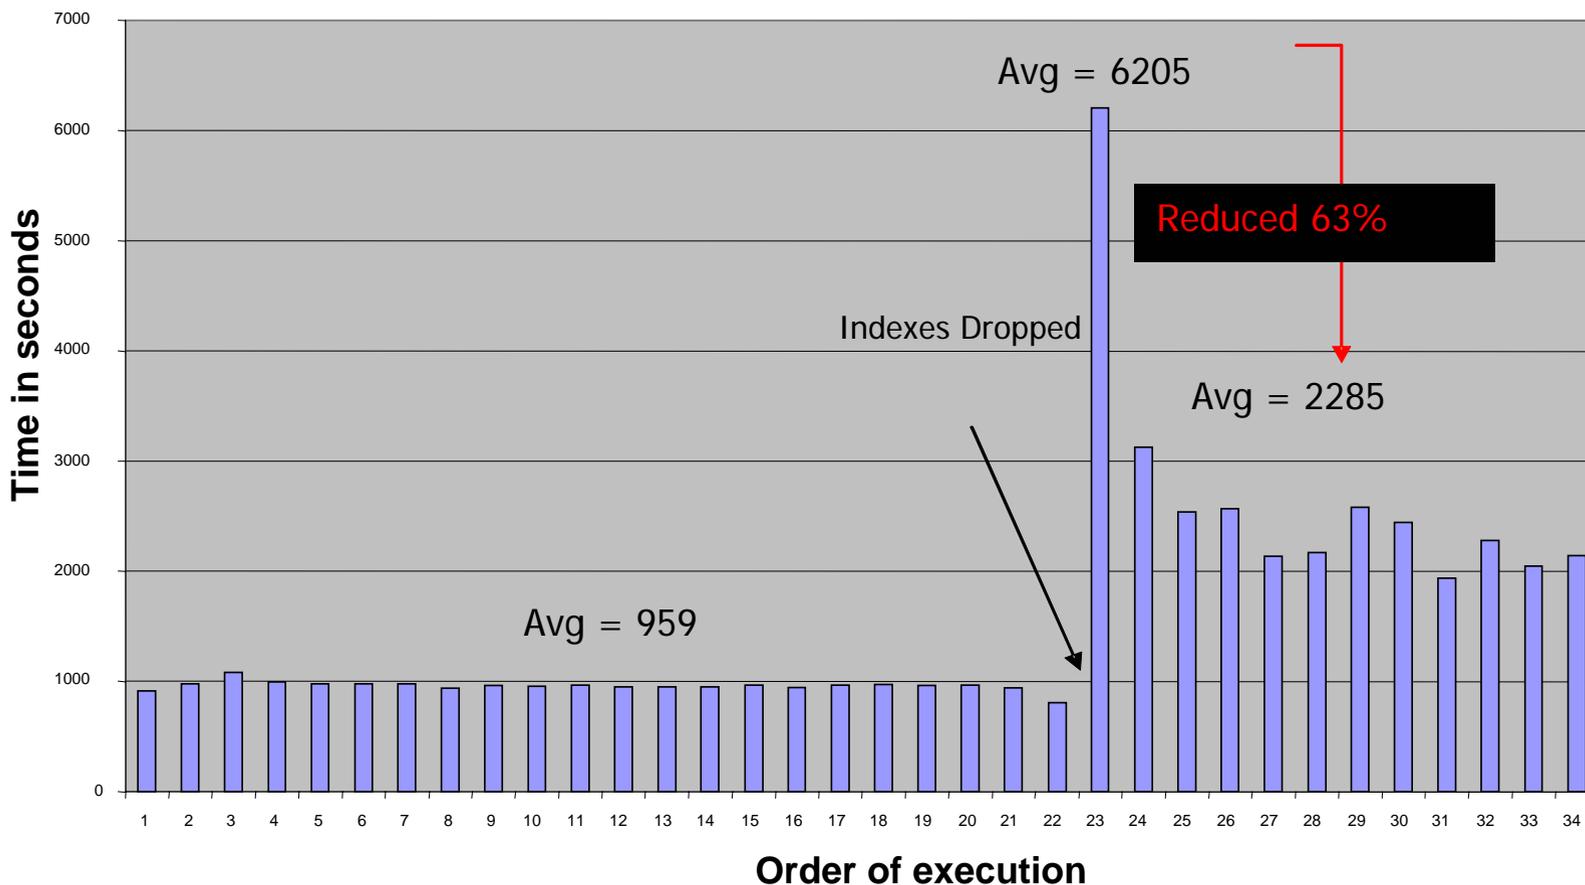
- Provide Automated Expert Advice
 - ▶ Configuration and Design Advisor
 - ▶ Health Monitor
- “Built-in” Automation
 - ▶ On-demand Storage
 - ▶ Automated Maintenance
 - ▶ Optimization, Parallelism and More
- Setup and Configuration
 - ▶ Configuration Assistant
 - ▶ Numerous Wizards

- Revolutionary memory tuning system called the Self Tuning Memory Manager (STMM)



STMM in Action – Dropping an Important Index

TPCH Query 21 - After drop index - Average times for the 10 streams



STMM in Action – Comparing Different Configurations

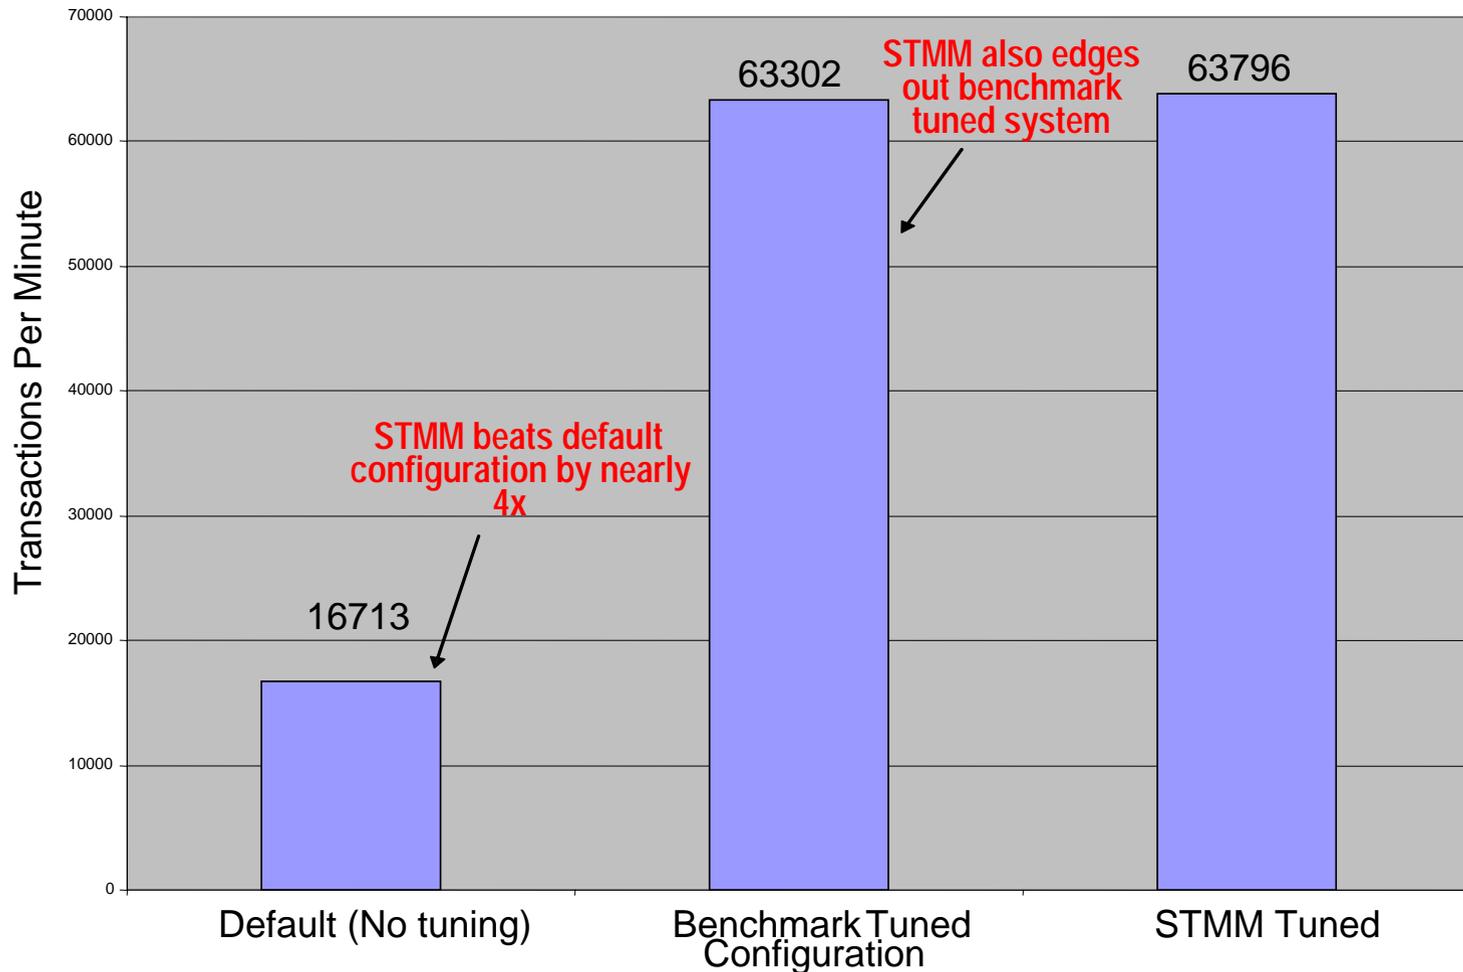


Table Partitioning

- What is Table (Range) Partitioning ?
 - ▶ Storing a table in more than one physical object, across one or more table spaces
- Why?
 - ▶ Increase table capacity limit
 - ▶ Increase large table manageability
 - ▶ Improve SQL performance through partition elimination
 - ▶ Provide fast & online data roll-in and roll-out

Table Partitioning

32K Partitions



64G

A-Z

Backup
Load
Recover



64G

A-C

Backup
Load
Recover



64G

D-M

Backup
Load
Recover



64G

N-Q

Backup
Load
Recover



64G

R-Z

Backup
Load
Recover



Backup
Load
Recover

Table partitioning example

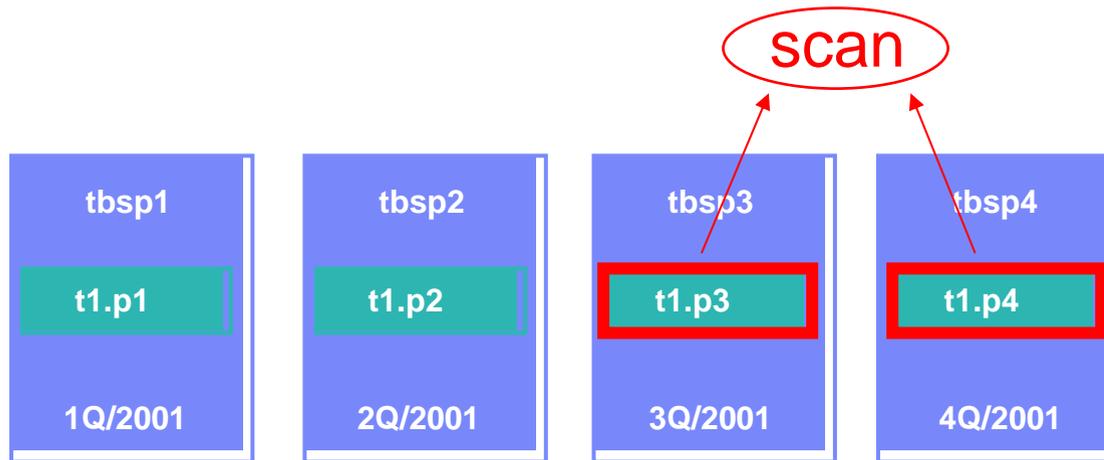
- Use **STARTING ... ENDING ...** to specify ranges
 - ▶ This example creates 4 ranges

```
CREATE TABLE sales(sale_date DATE, customer INT, ...)
PARTITION BY RANGE(sale_date)
(STARTING '1/1/2000' ENDING '3/31/2000',
STARTING '4/1/2000' ENDING '6/30/2000',
STARTING '7/1/2000' ENDING '9/30/2000',
STARTING '10/1/2000' ENDING '12/31/2000');
```

Partition Elimination : Table Scans

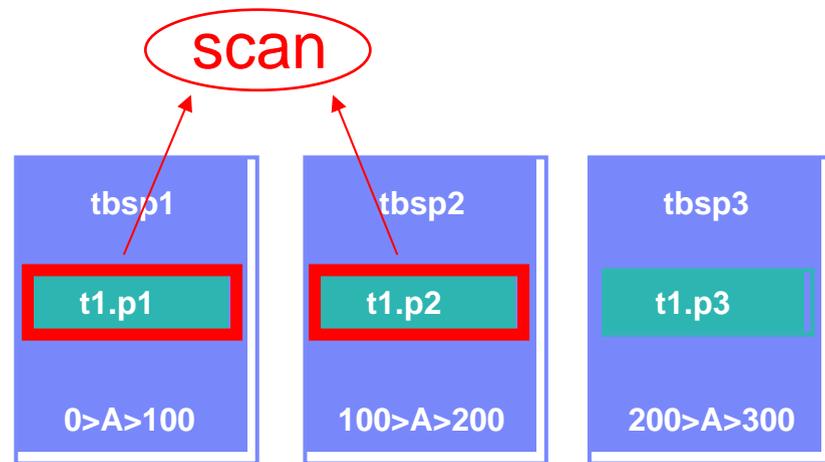
```
SELECT * FROM t1
WHERE
year = 2001 AND month > 7
```

- Will only access data in tablespace tbsp3 and tbsp4

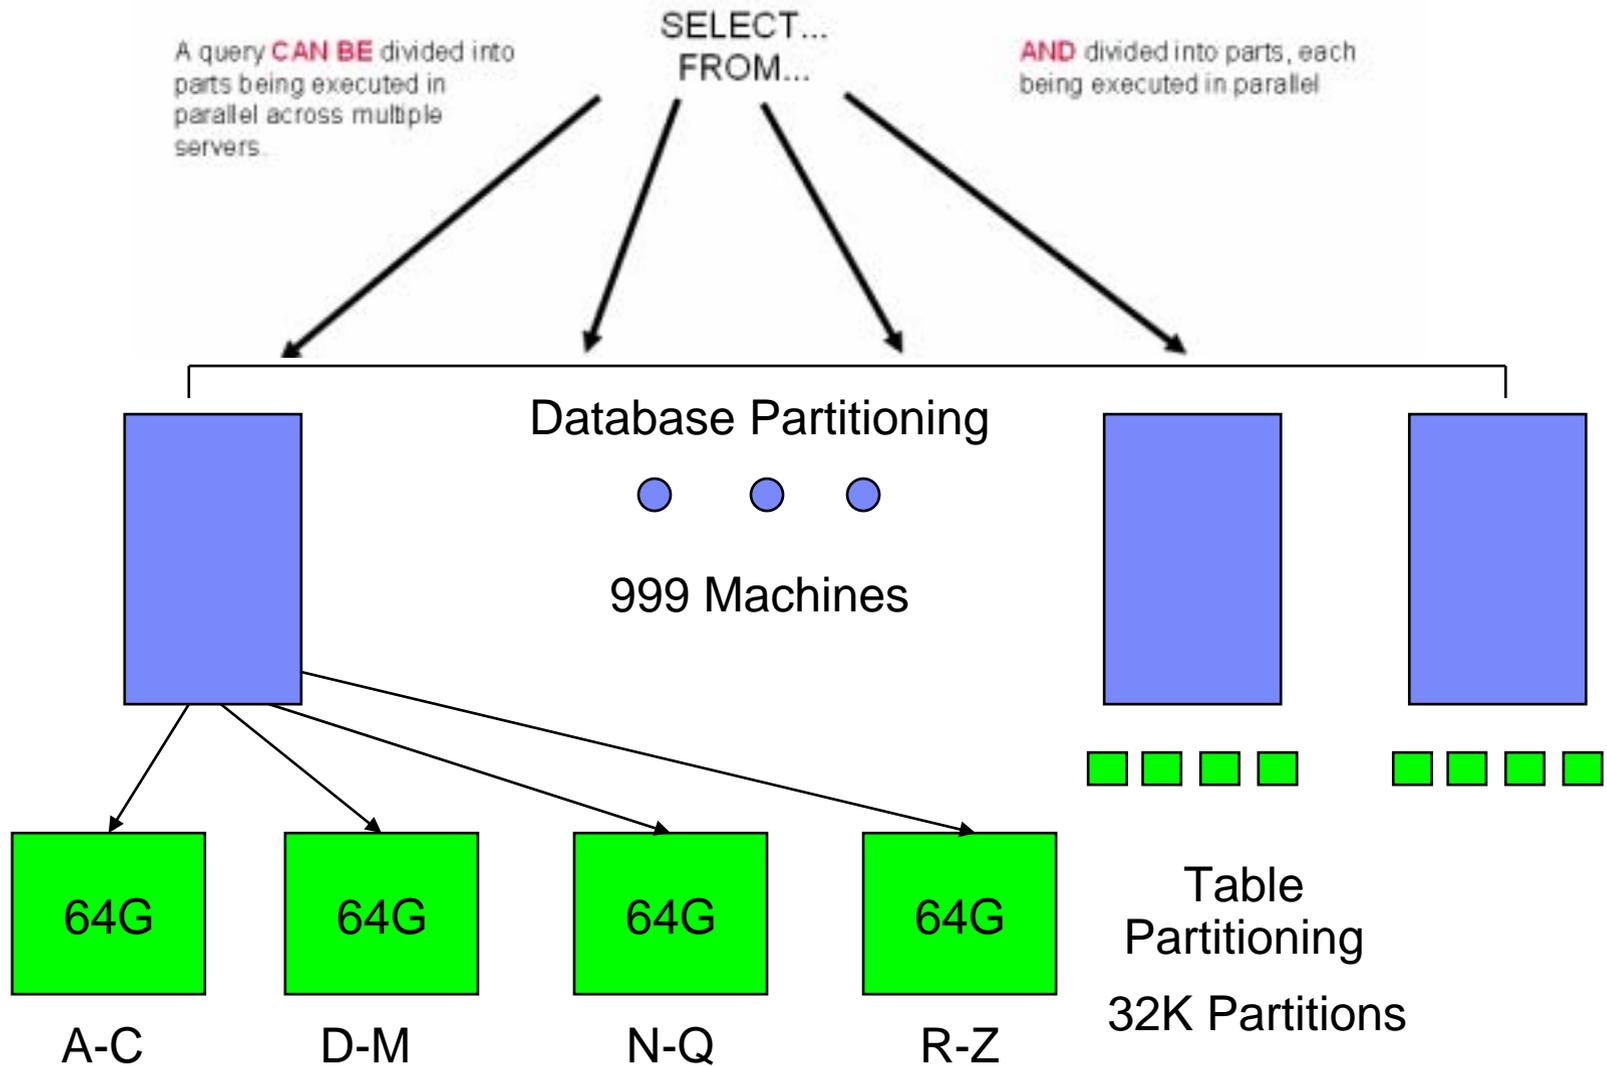


```
SELECT * FROM t1
WHERE
A>50 AND A<150
```

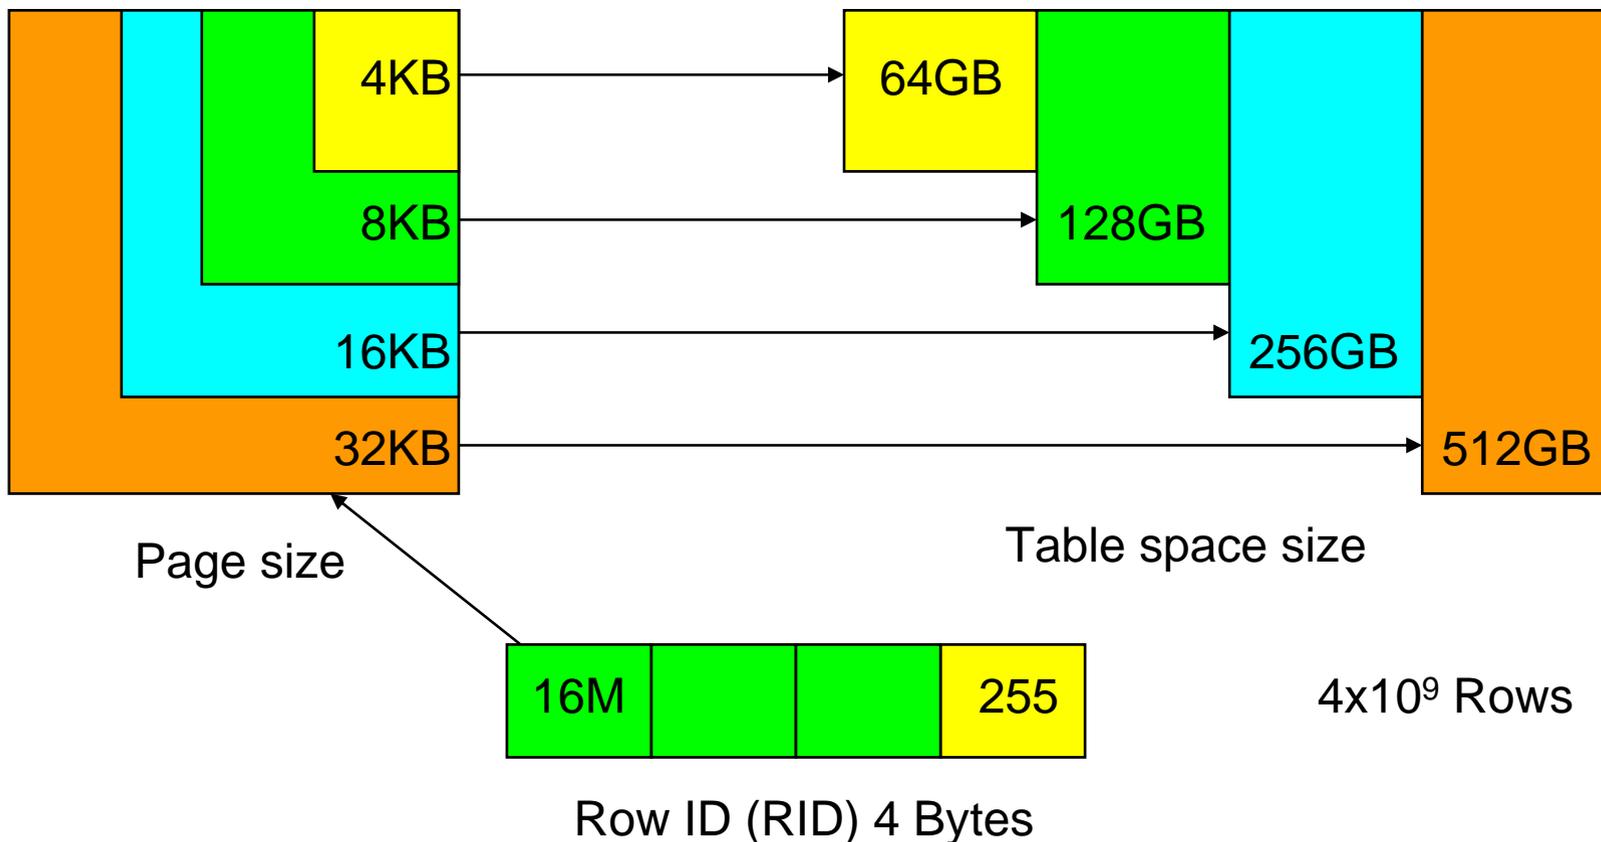
- Will only access data in tbsp1 and tbsp2



Hybrid Partitioning - Parallelism

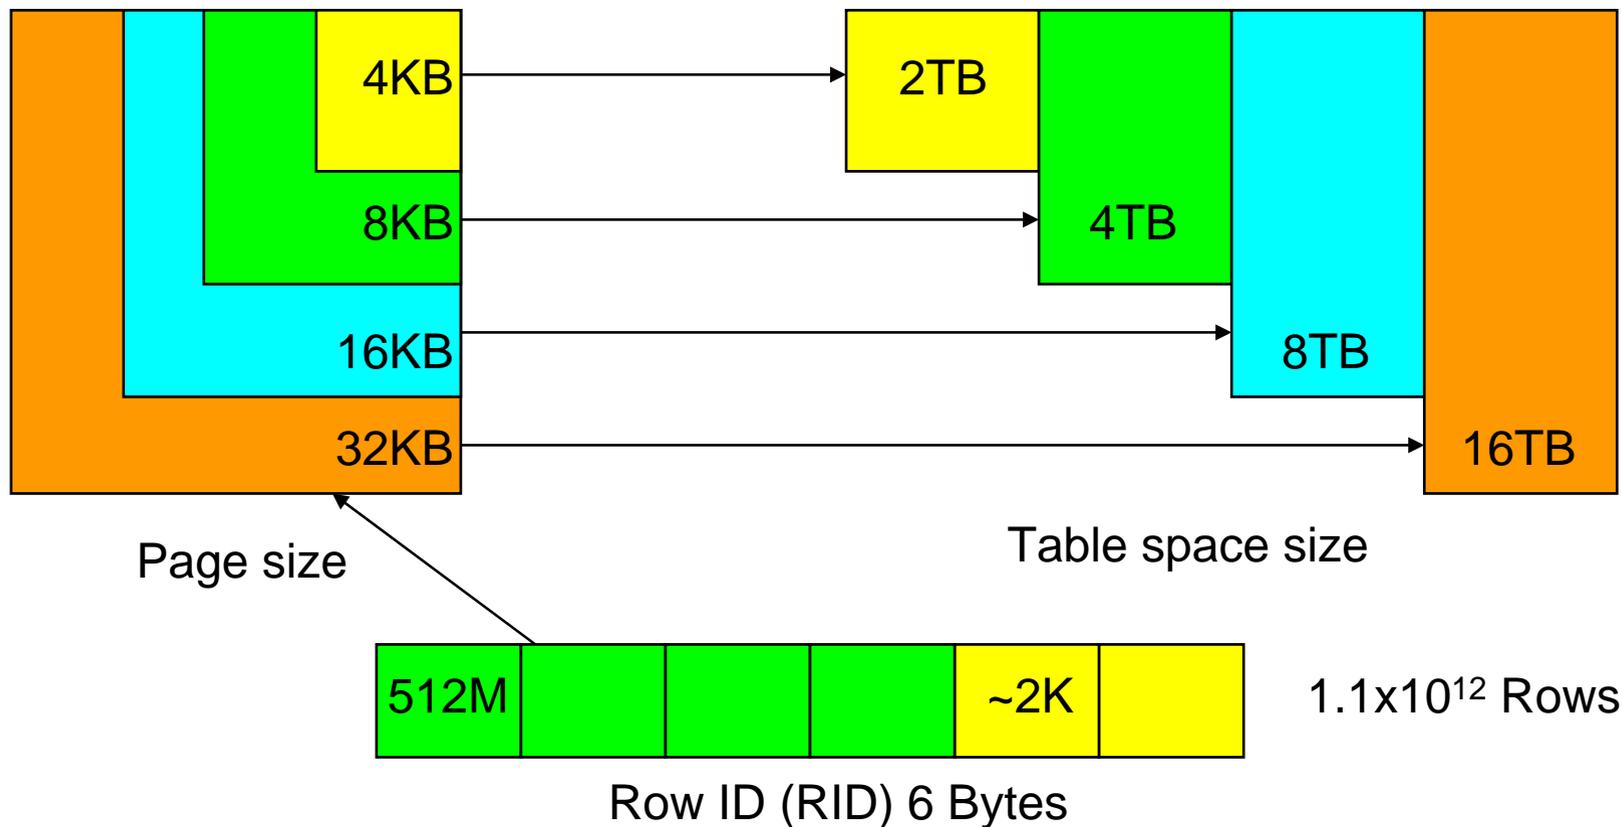


Current Table Space Design



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

New Large and Temporary Table Space Design



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

Larger Index Keys

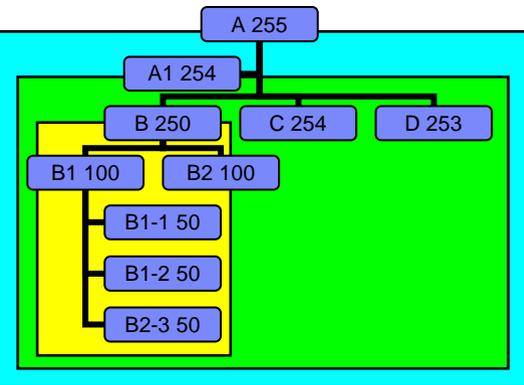
Version	Length of index key parts	# of columns in index key
Pre-Viper	1024	16
Post-Viper	1024 – 4K page	64
	2048 – 8K page	64
	4096 – 16 K page	64
	8192 – 32 k page	64

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 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

LBAC Query

```
SELECT * ... WHERE
SALARY >= 50000
```



No LBAC	SEC=255	SEC=254	SEC=100	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

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

- Repeating patterns within the data (and just within each row) is the key to good compression. Text data tends to compress well because of reoccurring strings as well as data with lots of repeating characters, leading or trailing blanks

Name	Dept	Salary	City	State	ZipCode
Fred	500	10000	Plano	TX	24355
John	500	20000	Plano	TX	24355

Fred	500	10000	Plano	TX	24355	John	500	20000	Plano	TX	24355	...
------	-----	-------	-------	----	-------	------	-----	-------	-------	----	-------	-----

Fred	(01)	10000	(02)	John	(01)	20000	(02)	...
------	------	-------	------	------	------	-------	------	-----

Dictionary

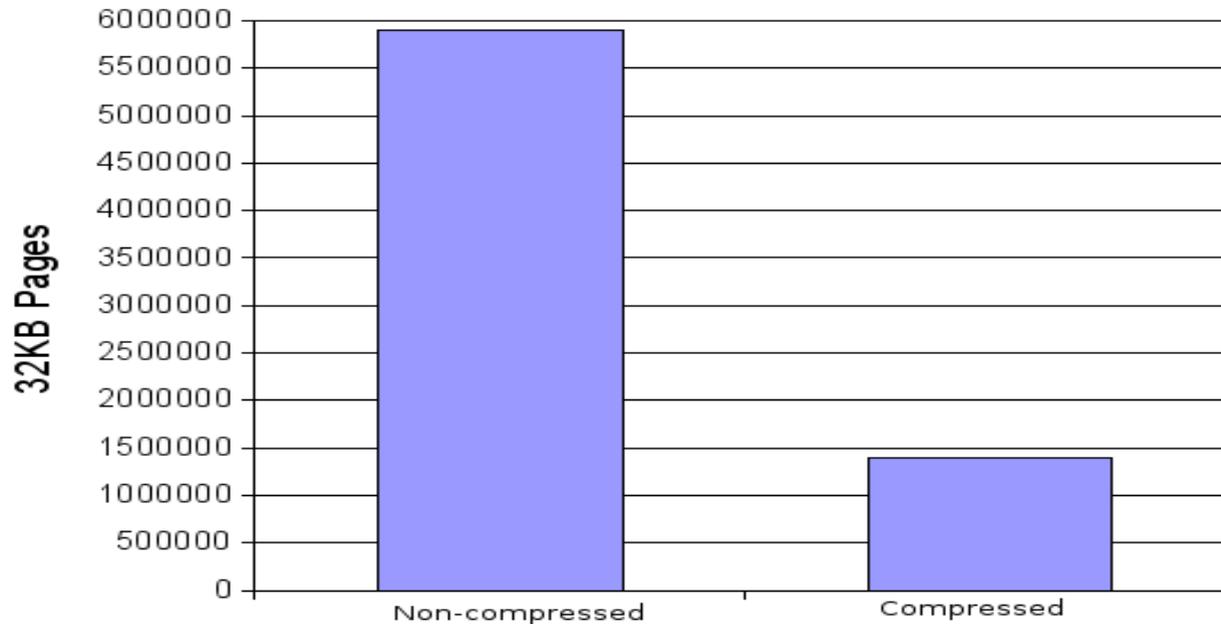
01	Dept 500
02	Plano, TX, 24355
...	...

Sample Compression Ratios (Customer Data)

<u>Compression Type</u>	<u>32KB Page Count</u>	<u>Space Required on Disk</u>
No compression	5893888	179.9GB
Row compression	1392446	42.5GB

% Pages Saved: 76.4%

T1 Compression - 179.9GB Initial Size



SQL Enhancements

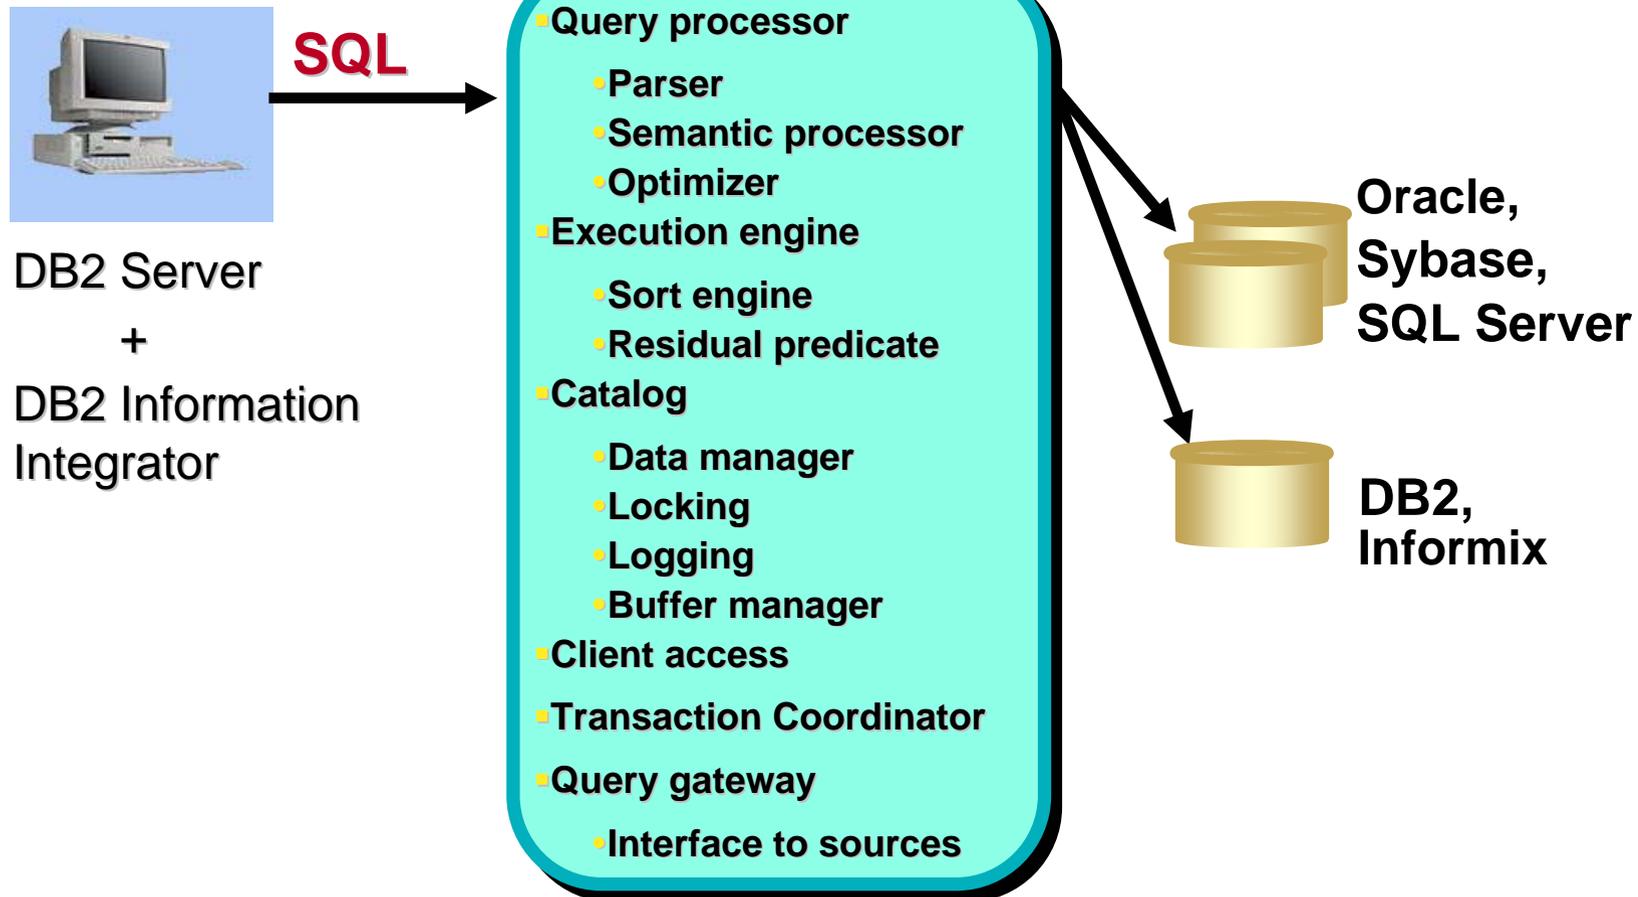
- Unicode Character Functions
- Administrative Functions
- CLP Limits
- SQL Error Message Function
- Alter Table Enhancements
- MQT Enhancements

Alter Table Command

- One of the highest administrative costs is making changes to tables such as dropping a column, changing column type, or changing column nullability
- The ALTER command will be modified to allow:
 - ▶ alter table drop column
 - ▶ alter table alter column type
 - ▶ alter table alter column nullability
- When an alter of any of these types is performed, the table is updated to the new design, but only SELECT scans will be allowed (no inserts or updates) until a REORG is completed

DB2 Information Integrator

Named changed to WebSphere Information Integrator



Important Resources

- IBM developerWorks - *IBM's resource for developers*
<http://www.ibm.com/developerworks/>
- DB2 developerWorks - *Technical resources for DB2 information management software*
<http://www.ibm.com/developerworks/db2/>
- Porting to DB2 UDB - *Technical resources and roadmap*
<http://www.ibm.com/db2/porting>
- DB2 Universal Database for Linux, UNIX and Windows - *DB2, DB2 Connect and DB2 Information Integrator Version 9 product manuals*
<http://www-306.ibm.com/software/data/db2/udb/support/manualsv9.html>
- DB2 Information Center – *DB2 V9 Help*
<http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp>

Important Resources (*continued*)

- DB2 Application Development
<http://www.ibm.com/software/data/db2/udb/ad/>
- DB2 Universal Database for Linux, UNIX and Windows - *Product support*
<http://www.ibm.com/software/data/db2/udb/support/>
 - ▶ Primary Support resources, client and fixpak downloads, newsgroups, FAQs
- DB2 Migrate Now! – *Facilitating migrations*
<http://www.ibm.com/software/data/db2/migration/>
- Migration Station – *“How to” Instructions for Migrating*
<http://www.ibm.com/developerworks/ondemand/migrate>
- IBM Virtual Innovation Center (VIC)
<http://www.developer.ibm.com/isv/welcome/vic.html>
 - ▶ Free Support! (DB2 Express, Websphere Express, etc...)
- DB2 Express 9
<http://www-306.ibm.com/software/data/db2/9/edition-express.html>

IBM Redbooks and Technical Articles

IBM Redbooks

- Oracle to DB2 UDB Conversion Guide
<http://www.redbooks.ibm.com/abstracts/sg247048.html>
- Microsoft SQL Server to IBM DB2 UDB Conversion Guide
<http://www.redbooks.ibm.com/abstracts/sg246672.html>
- MySQL to DB2 UDB Conversion Guide
<http://www.redbooks.ibm.com/abstracts/sg247093.html>
- Database Transition: Informix Dynamic Server to DB2 Universal Database
<http://www.redbooks.ibm.com/abstracts/sg246367.html>

IBM Technical Articles

- Porting to DB2 Universal Database Version 8.1 from Sybase Adaptive Server Enterprise
<http://www.ibm.com/developerworks/db2/library/techarticle/0307rada/0307rada.html>