IV Jornada Técnica de DB2 LUW
(19 Noviembre, 2009)

Information Management

DB2 Access Plans

Carlos Alejandro Ang
DB2 Certified Technical Specialist
calejandro@novalia.es

11/25/2009

© 2008 IBM Corporation

Fases de Compilación

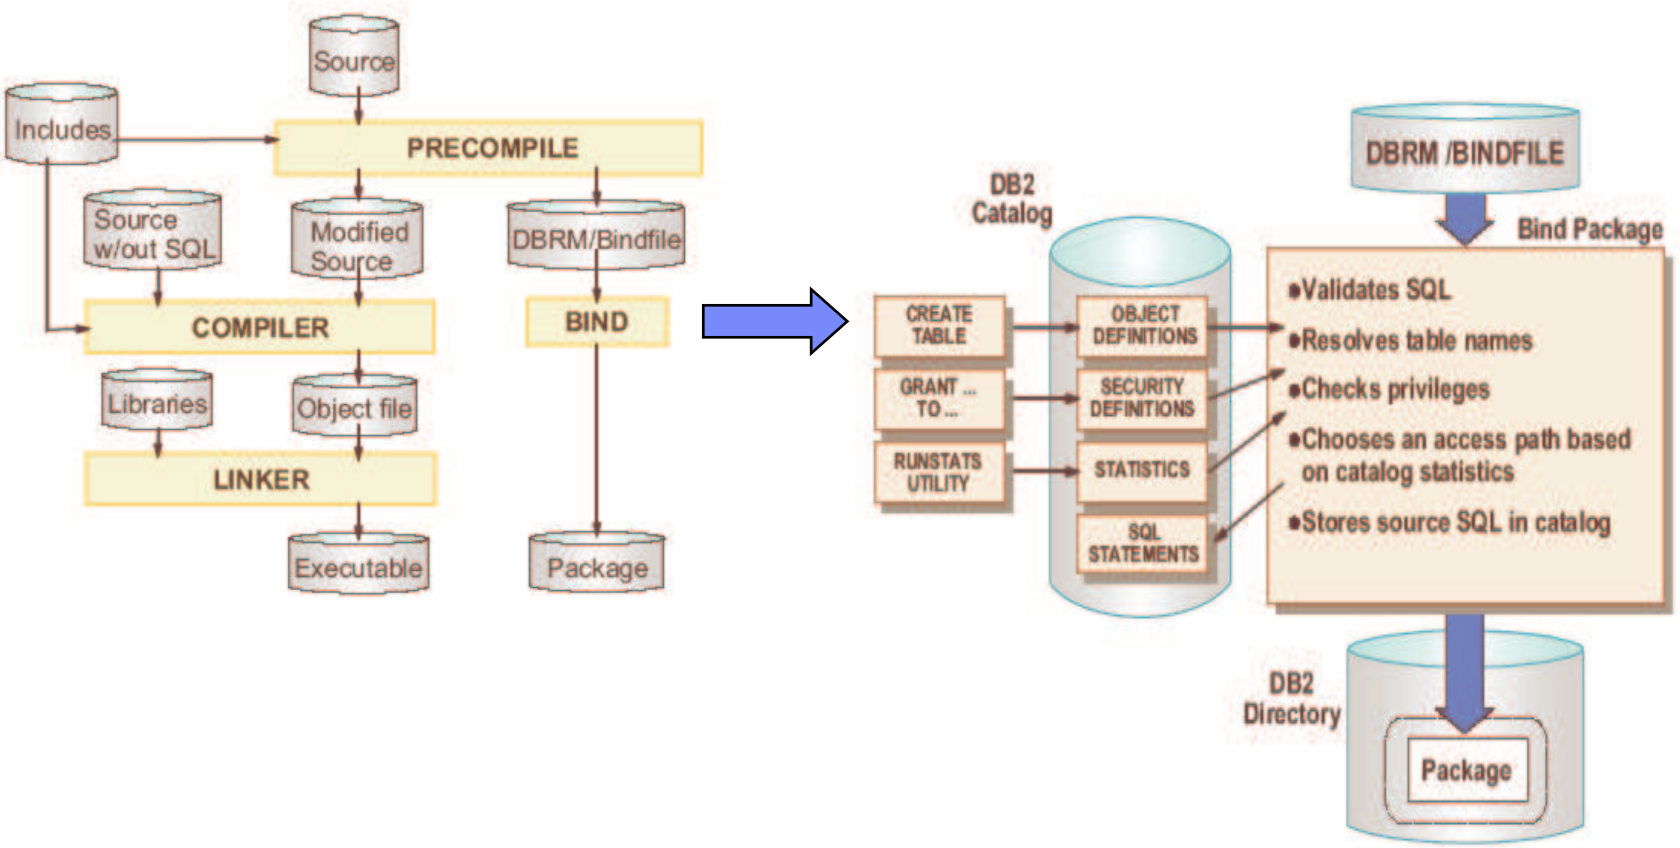

 © 2008 IBM Corporation

Generación Info de explain

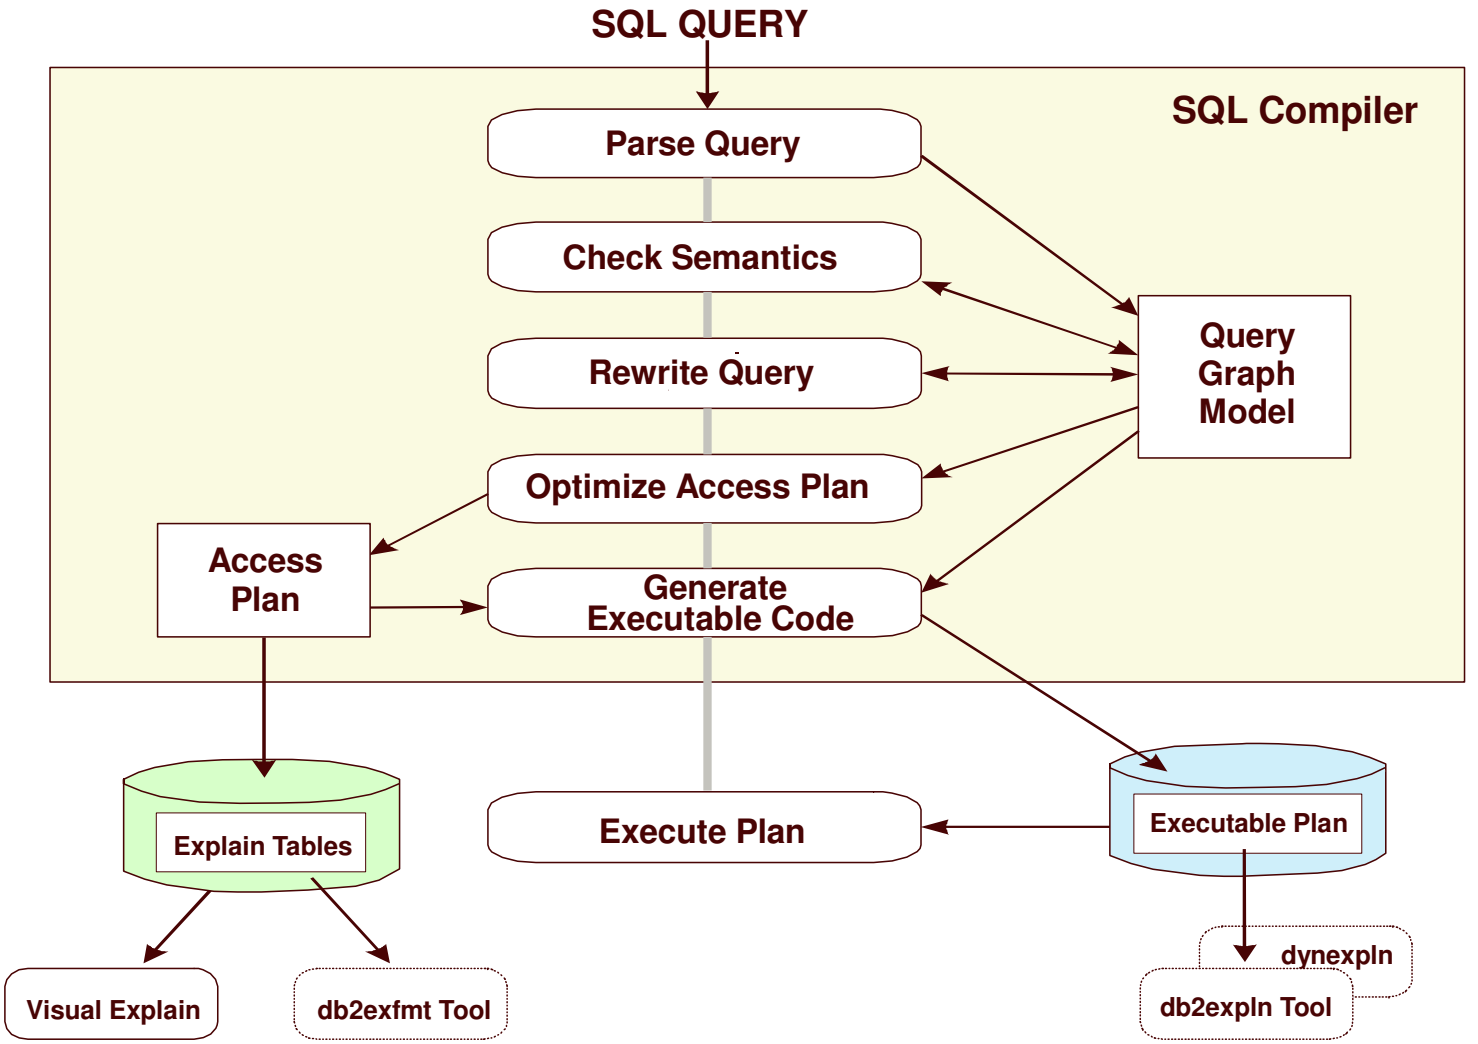

 © 2008 IBM Corporation

Centro de Control – Explain Gráfico

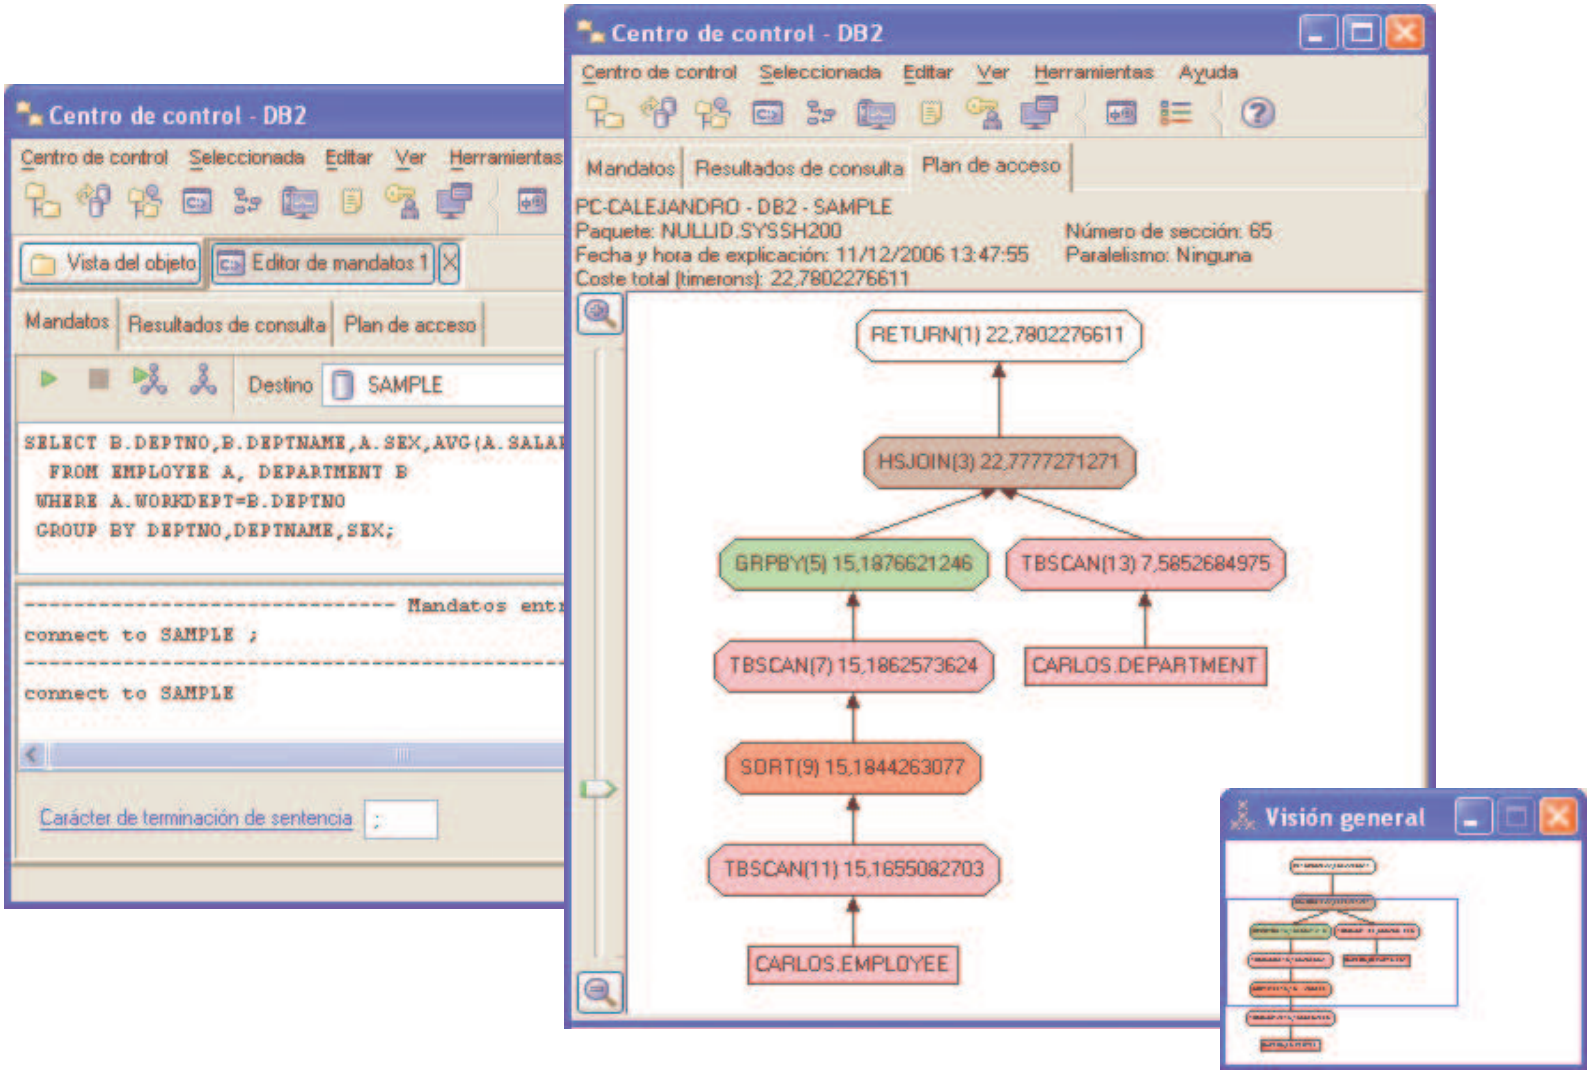

 © 2008 IBM Corporation

Generación Explain

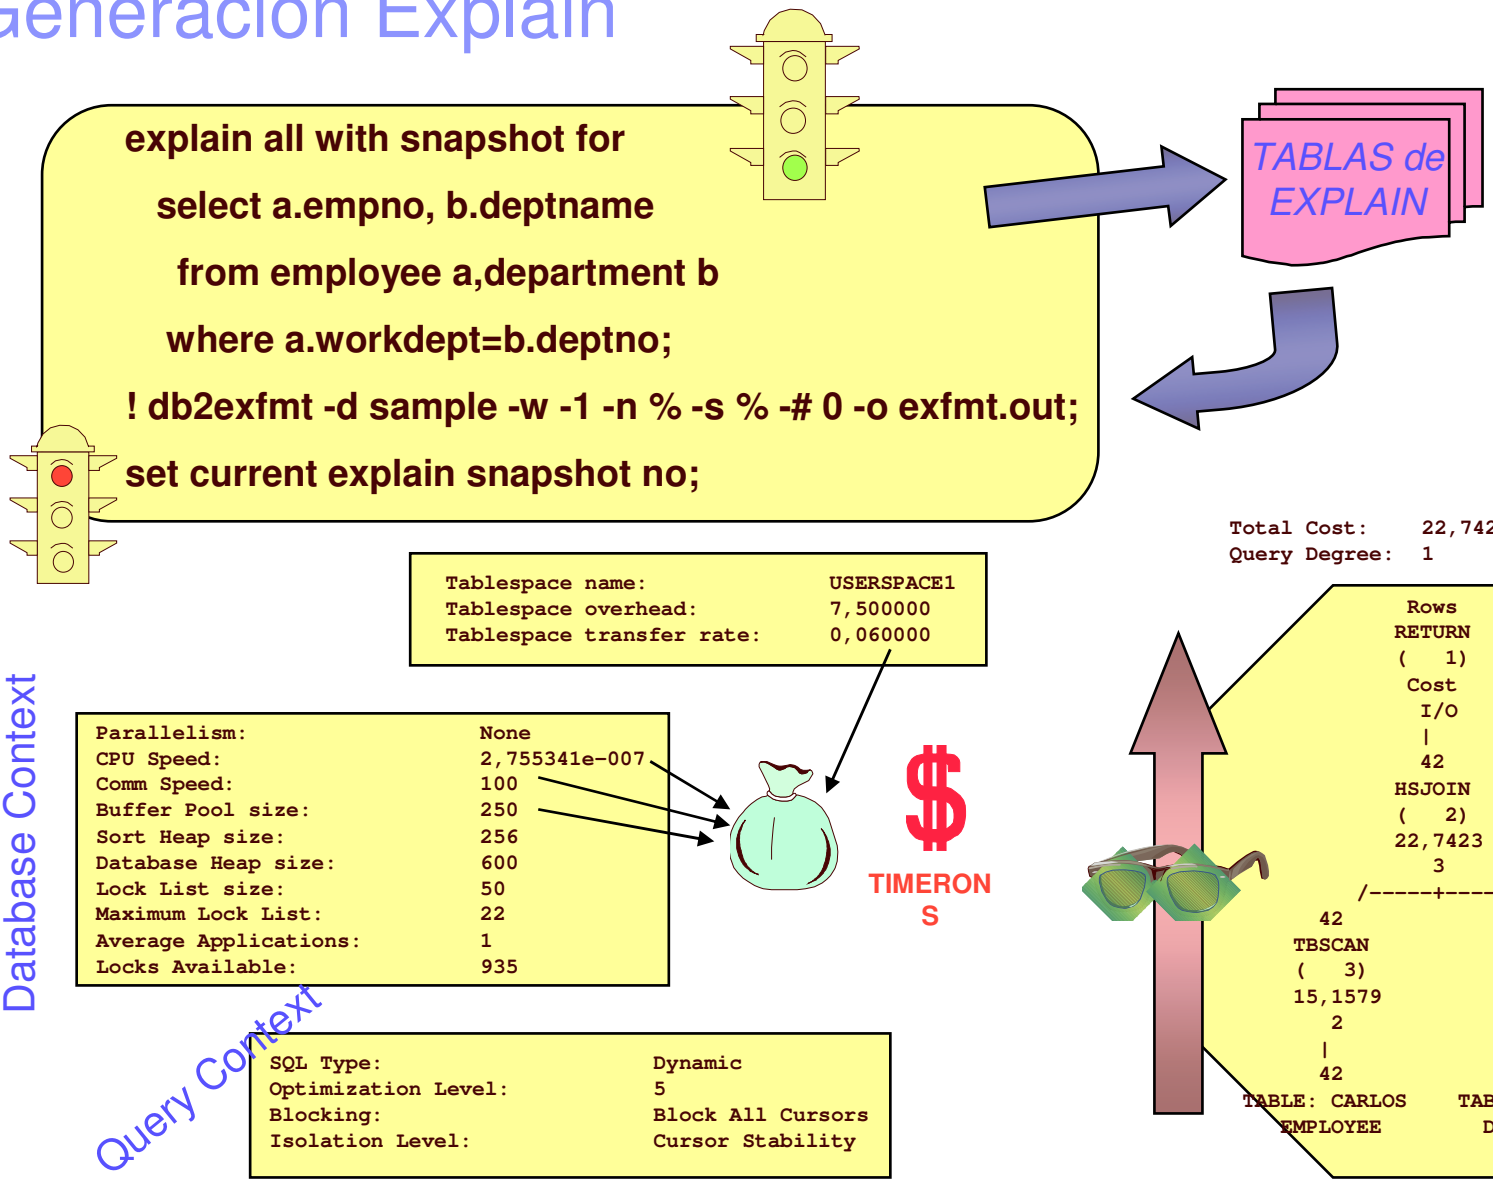

 © 2008 IBM Corporation

Nivel de Optimización

	0	1	2	3	5	7	9
non-uniform distribution statistics: - most frequent values - quantiles	N	N	Y	most frequent values quantiles-No	Y		
Rewrite Rules	basic	subset	- all, except comp. intensive rules	- most sub-query-to-join transform.	- all, except computationally intensive rules		all
complex dynamic SQL queries					reduced depend on machine size and number of predicates.	not reduced	
Search Strategies for Selecting Optimal Join	Greedy join if < 4 tables ==> greedy join = dynamic prog. join			Dynamic programming join			
				limited use of - composite inner tables - Cartesian products for star schemas involving "lookup" tables			no limits
Join Concepts	NLJoin Index Scan			MSJoin Index scan NLJoin Table scan			
Star Join Strategy	N			Y			
List Prefetch	N			Y			
Index ANDing	N			Y			
Appropriate to use with:	OLTP		OLAP/DSS	for queries with four or more joins.			

 © 2008 IBM Corporation

Esquema del Explain (Table Scan)

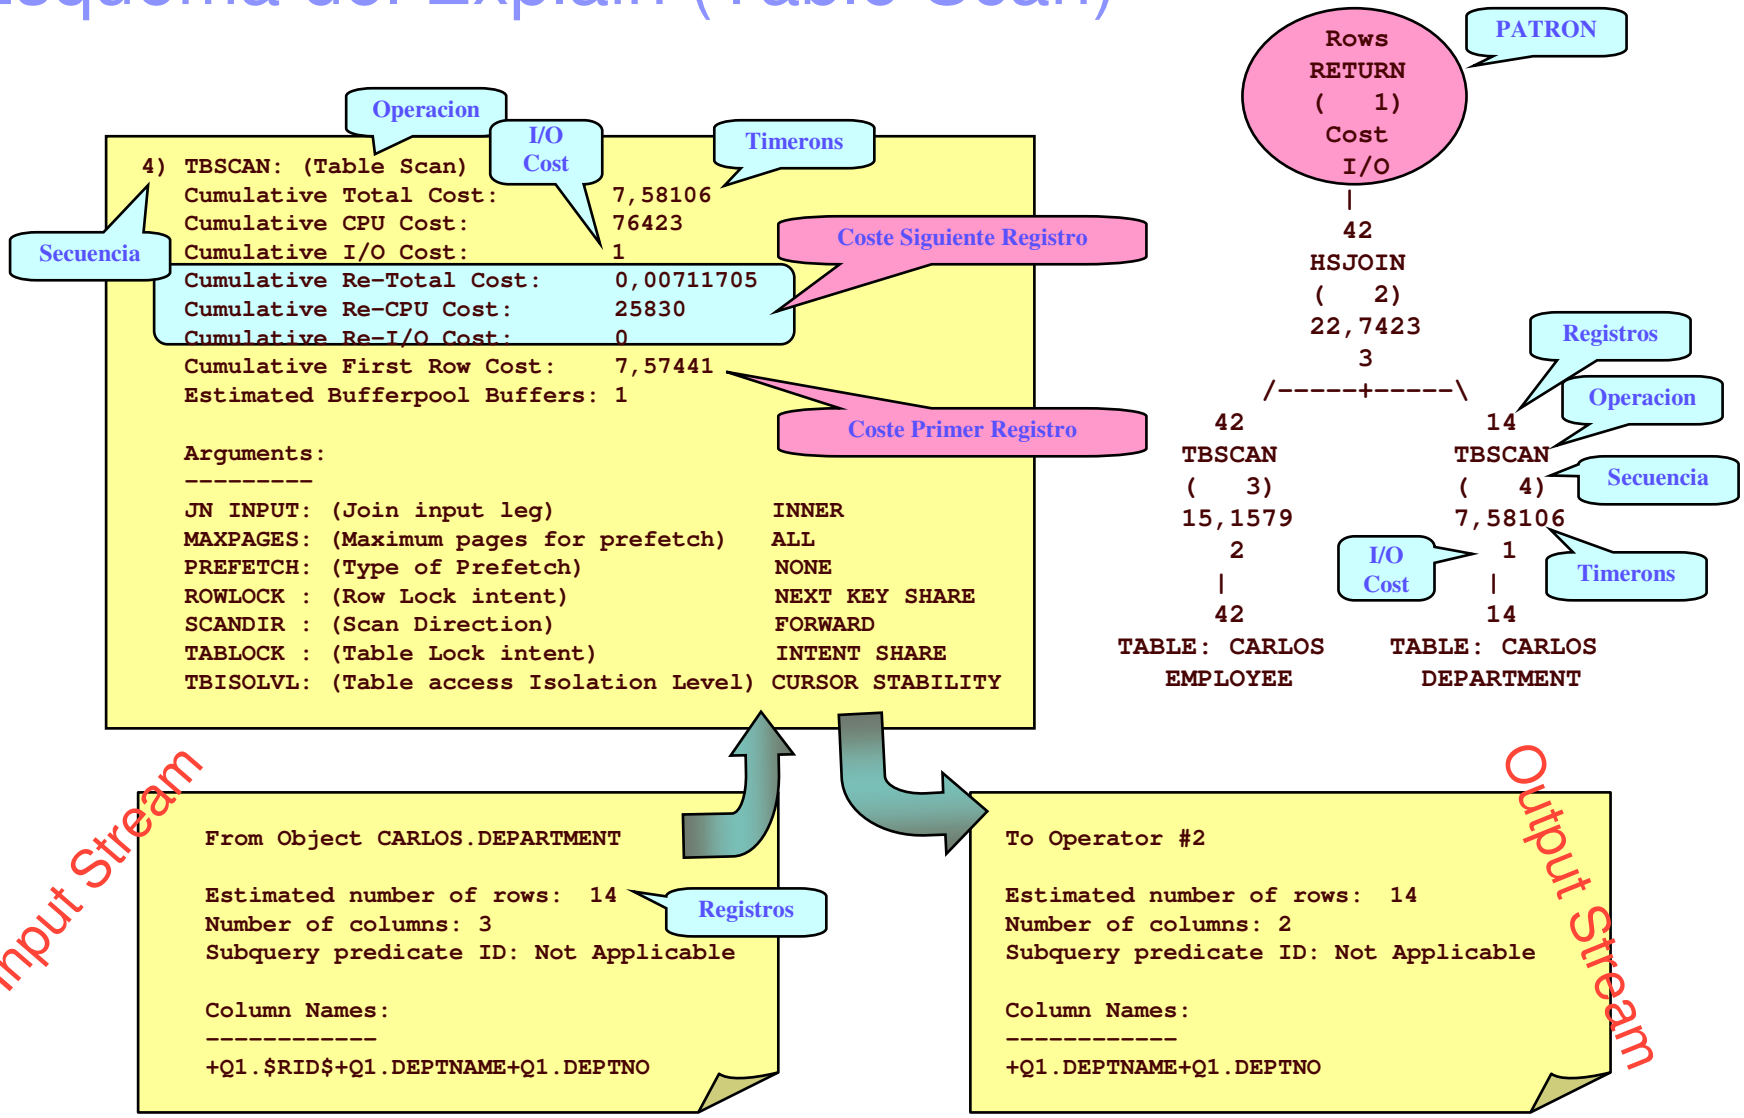

 © 2008 IBM Corporation

Uniones entre Tablas

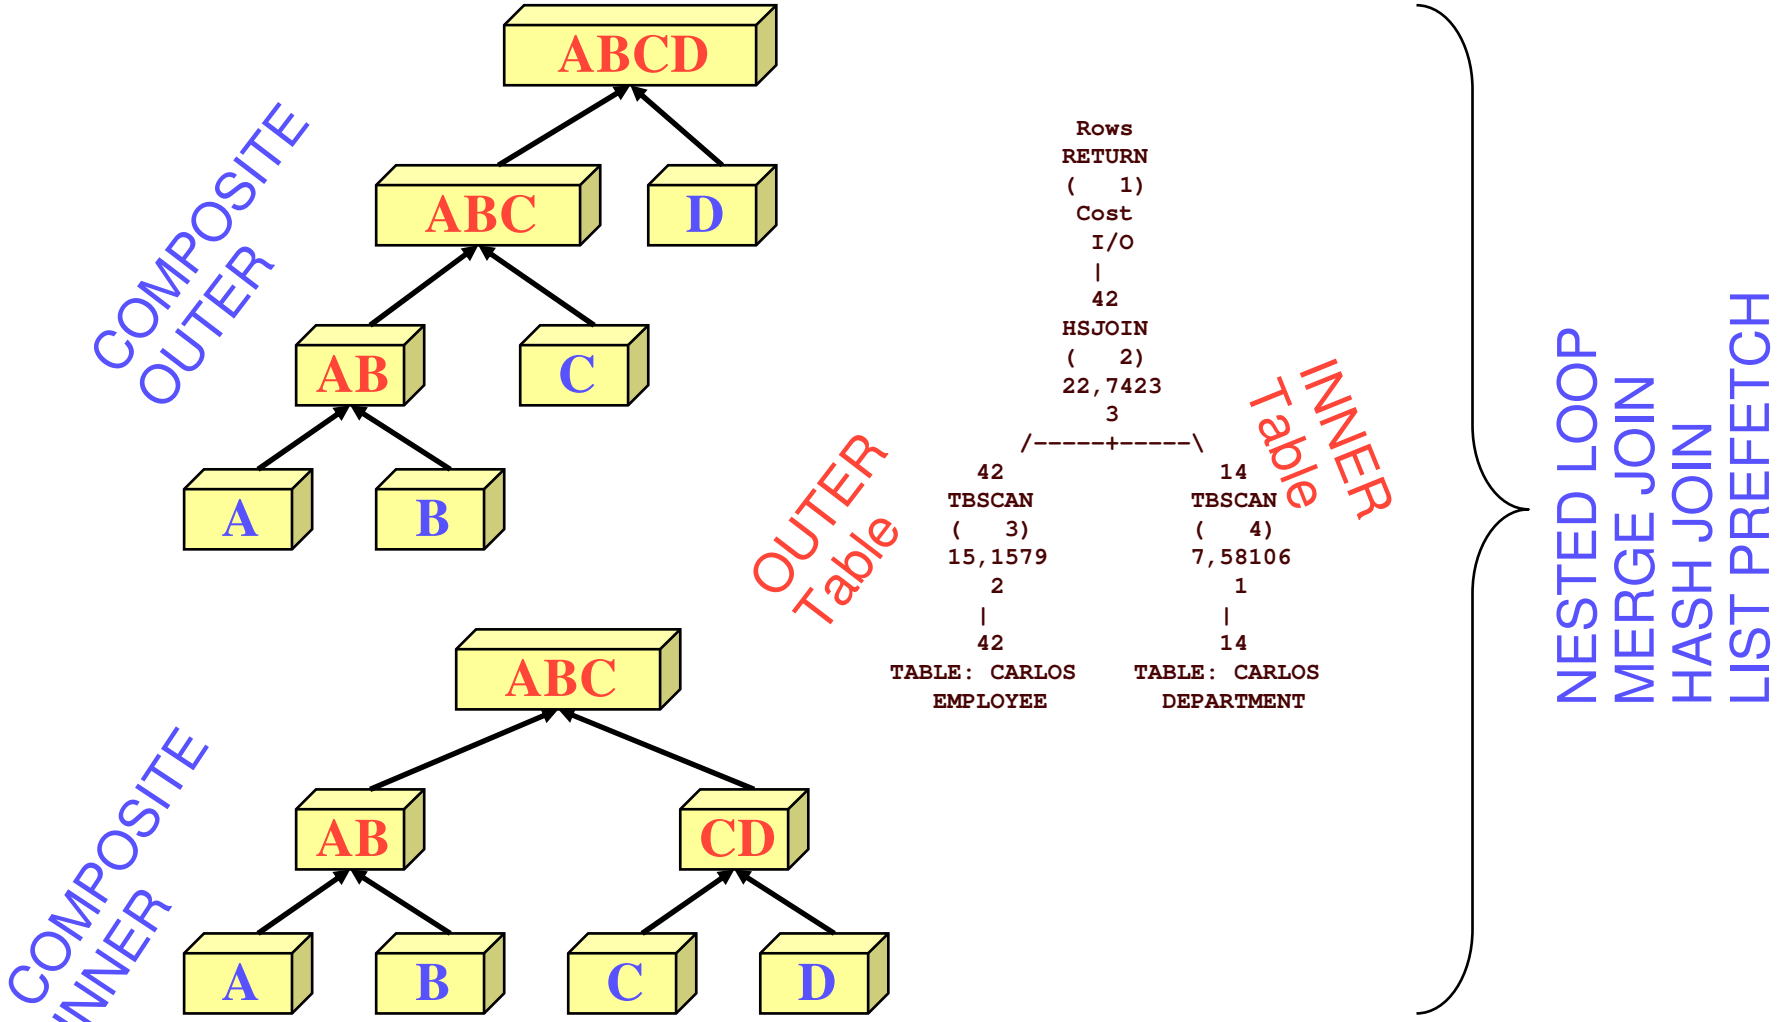

 © 2008 IBM Corporation

Operadores

```
SELECT Employee, Dept, Deptname
  FROM  Employee a, Department b
WHERE  a.Dept = b.Dept
```

NESTED LOOP

EMPLOYEE	DEPT
CHRISTINE HAAS	A00
MICHAEL THOMPSON	B01
VINCENZO LUCCHESSI	A00
JOHN GEYER	E01
IRVING STERN	D11
EVA PULASKI	D21
EILEEN HENDERSON	E11
THEODORE SPENSER	E21
SALLY KWAN	C01
SEAN O'CONNELL	A00
DELORES QUINTANA	C01

Outer Table

- *No existe indices*
- *No existe indice apropiado en Inner*
- *Puede utilizar índice en Outer*

Inner Table

DEPT	DEPTName
A00	SPIFFY COMPUTER SERVICE DIV.
B01	PLANNING
C01	INFORMATION CENTER
A00	SPIFFY COMPUTER SERVICE DIV.
D21	ADMINISTRATION SYSTEMS
E01	SUPPORT SERVICES
A00	SPIFFY COMPUTER SERVICE DIV.
B01	PLANNING

MERGE SCAN

EMPLOYEE	DEPT
CHRISTINE HAAS	A00
SEAN O'CONNELL	A00
MICHAEL THOMPSON	B01
DELORES QUINTANA	C01
SALLY KWAN	C01
IRVING STERN	D11
EVA PULASKI	D21
JOHN GEYER	E01
EILEEN HENDERSON	E11
THEODORE SPENSER	E21

- *Inner y Outer ordenados por unión*
- *Existe indice sobre columna unión*
- *Tablas Grandes*

Inner Table

DEPT	DEPTName
A00	SPIFFY COMPUTER SERVICE DIV.
A00	SPIFFY COMPUTER SERVICE DIV.
B01	PLANNING
C01	INFORMATION CENTER
D11	MANUFACTURING SYSTEMS

Puede decidir crear INDICE online

© 2008 IBM Corporation

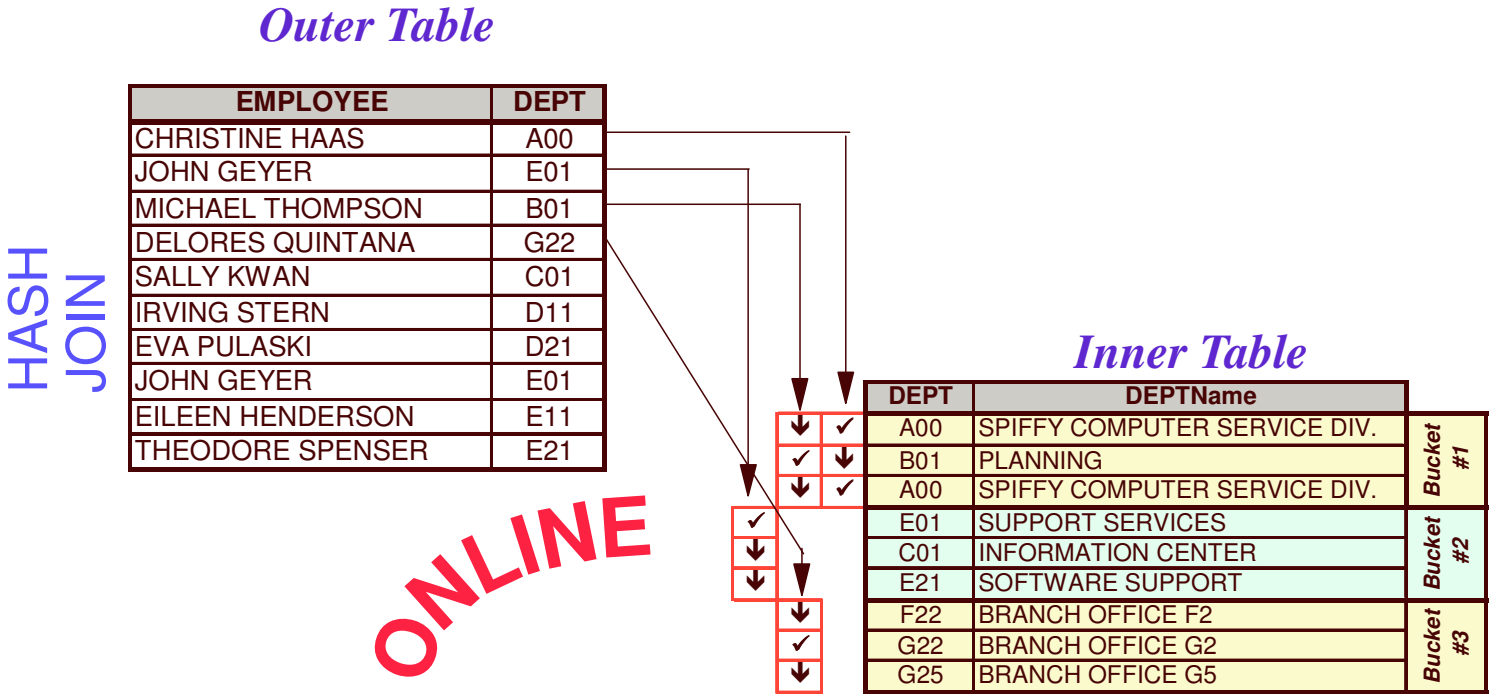

- *Inner subdividido en buckets mediante hash de la clave de unión.*
- *Outer hash por la clave de unión y nested loop contra el bucket de la Inner*
- *Inner bucket cabe en memoria, normalmente overflow a disco*

 © 2008 IBM Corporation

LIST PREFETCH

EMPLOYEE	Indice Employee DEPT	RID
CHRISTINE HAAS	A00	(1,1)
JOHN GEYER	E01	(2,1)
MICHAEL THOMPSON	B01	(1,2)
DELORES QUINTANA	G22	(3,2)
SALLY KWAN	C01	(2,2)
IRVING STERN	D11	(4,3)
EVA PULASKI	D21	(4,1)
JOHN GE	E01	(3,1)
EILEE	E11	(4,2)
TH	E21	(3,3)

Outer Table

I/O Random con respecto a Inner Table

Página

RID = (10 , 3)

Nº Registro

Convertir un índice NO clusterizada en CLUSTER de Tabla Inner!

ONLINE

Inner Table

IndiceCluster DEPT	RID
A00	(1,1)
B01	(1,2)
E01	(2,1)
C01	(2,2)
E01	(2,1)
G22	(3,2)
E21	(3,3)
D21	(4,1)
D11	
E11	(4,2)

IO Secuencial

DEPT	DEPTName	
A00	SPIFFY COMPUTER SERVICE DIV.	Página #1
B01	PLANNING	
A02	SPIFFY COMPUTER DIV.	
E01	SUPPORT SERVICES	Página #2
C01	INFORMATION CENTER	
E21	SOFTWARE SUPPORT	
F22	BRANCH OFFICE F2	Página #3
G22	BRANCH OFFICE G2	
G25	BRANCH OFFICE G5	

- *Ordena clave de unión por row id.*
- *Acceso secuencial al Inner con prefetch activado.*

© 2008 IBM Corporation

List Prefetch

select a.empno,a.workdept
from employee a
where a.empno = '100340';

```
Rows
RETURN
(   1)
Cost
 I/O
  |
 14,14
FETCH
(   2)
39,5495
32,7356
/---+---\
14,14        2828
RIDSCN    TABLE: CARLOS
(   3)       EMPLOYEE
7,58463
   1
  |
 14,14
SORT
(   4)
7,58427
   1
  |
 14,14
IXSCAN
(   5)
7,58261
   1
  |
 2828
INDEX: CARLOS
EMPNO
```

List Prefetch Cardinality

```
Schema: CARLOS
Name:   EMPNO
Type:   Index

Number of columns:              1
Number of rows:                 2828
Number of buffer pool pages:    64
Distinct row values:            No
Prefetch page count:            32
Container extent page count:    32
Index clustering statistic:     0,332489
Index leaf pages:               9
Index tree levels:              2
Index full key cardinality:     202
Index first key cardinality:    202
Index RID count:                2828
Index empty leaf pages:         0
```

1 valor distinto de 202 valores =

1/202 = 0,0049 = 0,005 (Factor de filtro)

0,005 * 2828 = 14,14 registros

Index Scan Cardinality

```
Schema: CARLOS
Name:   EMPLOYEE
Type:   Table

Number of columns:              14
Number of rows:                 2828
Number of buffer pool pages:    64
Distinct row values:            No
Prefetch page count:            32
Container extent page count:    32
Table overflow record count:    0
```

```
Start Key Predicate
Comparison Operator: Equal (=)
Filter Factor:       0,005

Predicate Text:
--------------
(Q1.EMPNO = '100340')
```

```
Stop Key Predicate
Comparison Operator: Equal (=)
Filter Factor:       0,005

Predicate Text:
--------------
(Q1.EMPNO = '100340')
```

 © 2008 IBM Corporation

Total Cost: 39,5495

```
        Rows
       RETURN
       (   1)
        Cost
         I/O
         |
       14,14
       FETCH
       (   2)
       39,5495
       32,7356
      /---+---\
   14,14        2828
  RIDSCN    TABLE: CARLOS
  (   3)       EMPLOYEE
  7,58463
     1
     |
   14,14
  SORT
  (   4)
  7,58427
     1
     |
   14,14
  IXSCAN
  (   5)
  7,58261
     1
     |
   2828
INDEX: CARLOS
     EMPNO
```

db2 runstats on table carlos.employee
with distribution on all columns
and detailed indexes all

Total Cost: 65,6494

```
        Rows
       RETURN
       (   1)
        Cost
         I/O
         |
         92
       TBSCAN
       (   2)
       65,6494
         64
         |
        2828
    TABLE: CARLOS
       EMPLOYEE
```

92 valores distintos de 2828 registros =

92/2828 = 0,0325 (Factor de filtro)

SYSCAT.COLDIST

COLVALUE	VALCOUNT
'100340'	92

```
Start Key Predicate
Comparison Operator: Equal (=)
Filter Factor:       0,0325

Predicate Text:
--------------
(Q1.EMPNO = '100340')

Stop Key Predicate
Comparison Operator: Equal (=)
Filter Factor:       0,0325

Predicate Text:
--------------
(Q1.EMPNO = '100340')
```

db2 reorg table carlos.employee
index carlos.empno
db2 runstats on table carlos.employee
with distribution on all columns
and detailed indexes all

Total Cost: 23,3904

```
        Rows
       RETURN
       (   1)
        Cost
         I/O
         |
         92
       FETCH
       (   2)
       23,3904
       3,08204
      /---+---\
    92          2828
  IXSCAN    TABLE: CARLOS
  (   3)       EMPLOYEE
  7,61944
     1
     |
   2828
INDEX: CARLOS
     EMPNO
```


© 2008 IBM Corporation

Proximamente …

- **Constraints**
- **Informational Constraints**
- **Tablas Volatiles**

Uniform Data Distribution

CURVA de I/O
por Bufferpool size
para INDEXSCAN

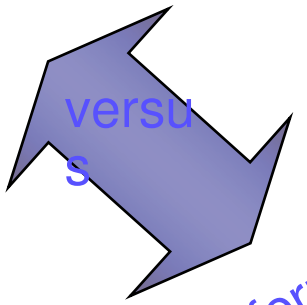

Non - Uniform Data Distribution

Tratamiento de
Correlated
Columns

- **Estadisticas FRECUENCIA**
- **Estadisticas QUANTILES**
- **Estadisticas Multi-column GROUP**
- **Stadistical Views**

QUEDA MUCHO POR VER
! ! !

 © 2008 IBM Corporation

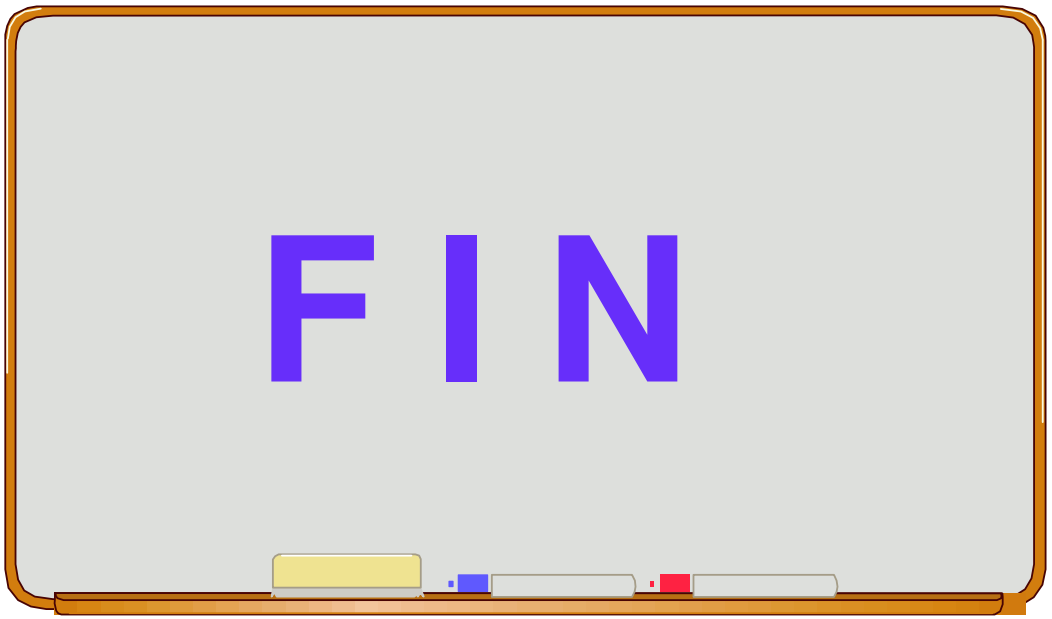

 © 2008 IBM Corporation