



IBM System *i*™

DB2 UDB : Advanced Functionality Sampler

Alison Butterill
butteril@ca.ibm.com



i want stress-free IT.
i want control.
*i want an **i**.*

Simplify IT

© 2006 IBM Corporation

IBM System *i*

Acknowledgment and Disclaimer

Acknowledgment:

Many people contributed to this presentation. In particular (but in no particular order) thanks goes to:

Skip Marchesani - Lexel Custom Systems, New Jersey
Kent Milligan, IBM Corporation, Rochester, Minnesota

Disclaimer:

The information contained in this document has not been submitted to any formal IBM test and is distributed on an as is basis without any warranty either express or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customers' ability to evaluate and integrate them into the customers' operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will result elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

Reproduction:

The base presentation is the property of IBM Corporation. Permission must be obtained PRIOR to making copies of this material for any reason.

*i want an **i**.*

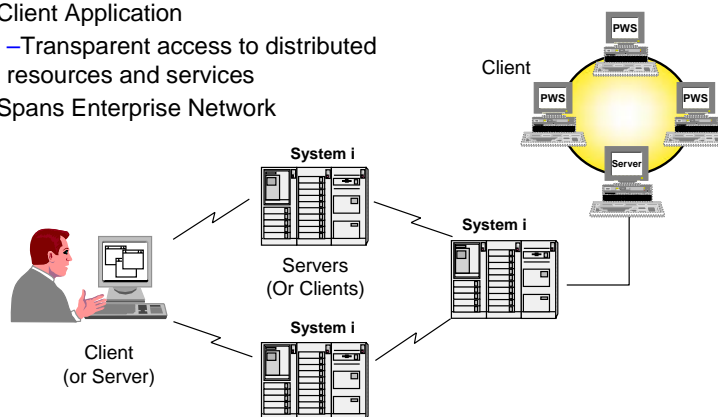
© 2006 IBM Corporation

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

Client/Server Applications

- Client Application
 - Transparent access to distributed resources and services
- Spans Enterprise Network




- **Optimize location of applications, data and other system resources**

Database Server Capability

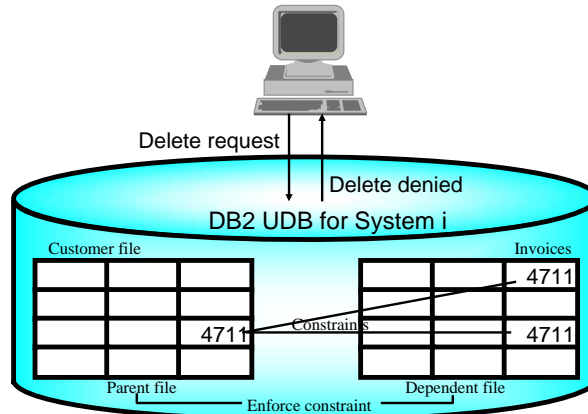
- Relational Database Function
- Stability
- Ease of Use
- Two Interfaces - Native and SQL
- Ease of Management
- Integration with System i Architecture
 - Journaling and commitment control
 - System managed access path protection
 - Database integrity
 - Security
- Most widely used multi-user RDB in the world

DB2 UDB for System i: Best Relational Database in the Marketplace

Agenda

- Introduction
-  • Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

What is Referential Integrity?



- A capability provided by the database management system to ensure
 - logical consistency of data values between files
 - validity of data relationships
 - robust enforcement of integrity constraints

RI Concepts

- Parent and Dependent files
- Unique key and Primary key constraint
- Parent key and Foreign key
- Referential constraint
- Referential integrity
- Referential constraint rules

Referential Integrity - Example

Customer Master file

Parent file

	1234
	2345
	4711
	3456

CUSTNO
(Parent key)

Invoice file

Dependent file

4711	
1234	
4711	
2345	
3456	

CUSTNUM
(Foreign key)

```

CL -->  ADDPFCST FILE(master)
        TYPE(*PRIKEY
        )
        CST(name)
        KEY(custno)

SQL -->  ALTER TABLE master
        ADD CONSTRAINT
        name
        PRIMARY KEY (custno)
    
```

```

ADDPFCST FILE(invoice)
        TYPE(*REFCST)
        CST(name2)
        KEY(custnum)
        PRNFILE(master) PRNKEY(custno)
        UPDRULE(*RESTRICT)
        DLTRULE(*CASCADE)

ALTER TABLE invoice
        ADD CONSTRAINT name2
        FOREIGN KEY (custnum)
        REFERENCES master
        (custno)
        ON DELETE CASCADE
        ON UPDATE RESTRICT
    
```

Referential Integrity - Example

Customer Master file

Parent file

	1234
	2345
	4711
	3456

CUSTNO
(Parent key)

Invoice file

Dependent file

4711	
1234	
4711	
2345	
3456	

CUSTNUM
(Foreign key)

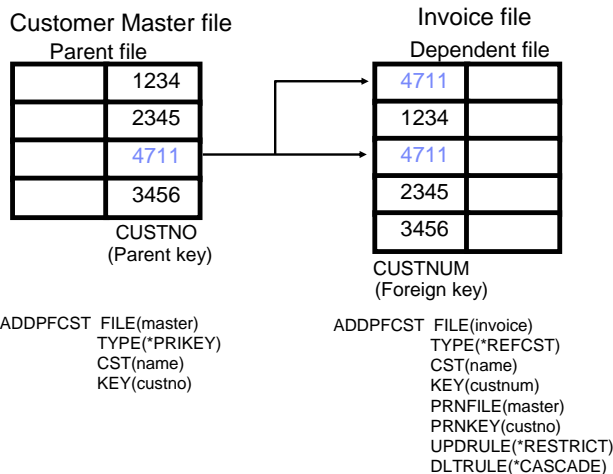
9000	
------	--

- Try to add this new row to the Invoice File
- What happens?

Referential Constraint Rules

- Insert rule:
 - no explicit rules, but insert operations into Dependent files are checked
- Delete rule:
 - RESTRICT
 - CASCADE
 - SET NULL
 - SET DEFAULT
 - NO ACTION
- Update rule:
 - restrict updates to a Parent key
 - RESTRICT
 - NO ACTION

Referential Integrity - Example



- Deleting customer 4711 will delete both related invoices
- Updating 4711 key value is prevented

Referential Integrity - Requirements

- Basic requirement is:
 - Parent key and foreign key must have matching field attributes
- On Parent file you can create
 - one primary key
 - multiple unique and parent keys
- Journal requirements
 - Restrict rule: journal not required
 - Any other rule:
 - Parent and Dependent files must be journaled to the same journal
 - Implicit commitment control is started

General Implementation Considerations

- Referential constraints not allowed for
 - System files
 - Source files
 - Program described files
 - Multi-member files
- Maximum of one member for parent or dependent file
- Referential constraints cannot span ASPs
- Maximum of 300 referential constraint relationships per file (parent or dependent)

Creating Referential Integrity Constraints

- Native interface:
 - Parent and dependent files can be created with SQL or DDS
 - Define constraints through a new Add Physical File Constraint (ADDPFCST) command
- SQL interface
 - Create Parent and Dependent files through CREATE TABLE
 - CREATE TABLE allows constraints definition
 - ALTER TABLE can be used to add constraints to existing tables
- Operations Navigator

Verifying Constraints

- DB2 UDB for System i implements "Declarative RI"
- Constraints will be verified:
 - at creation time
 - when the state goes to established/enabled
 - after a file has been restored

For Large Files, this may take some time!

- Constraints will be enforced
 - whenever an I/O operation occurs
 - delete from Parent file
 - insert into Dependent file
 - update on either file

What is Check Pending Status?

- Condition where some foreign key values do not match any parent key
- A check pending condition may occur:
 - when adding referential constraints to existing files
 - after abnormal system failures
 - after restoring files at different data levels
 - after applying/removing journal changes
- To get out of a Check Pending condition:
 - disable the constraint
 - remove the check pending condition by
 - inserting the required parent records
 - removing/changing the dependent records
 - enable the constraint
- WRKPCFCST - Work with Physical File Constraints →

Constraint Management Commands

Display Report			
Width	142		
Column	1		
Control			
Line	1	2	3
..	ORDER_NUMBER	CUSTOMER_NUMBER	ORDER_DATE
000001	02020	12312	02/03/1994
000002	02021	12312	04/13/1994
000003	02022	12312	04/25/1994
*****	* * * * * END OF DATA * * * * *		

- Display Check Pending Constraint - DSPCPCST
 - displays dependent records with no matching parent key
 - constraint must be in disabled state
- Change Physical File Constraint - CHGPCFCST
 - sets to enabled/disabled state

IBM System i

iSeries Navigator - RI Constraint Management (V5R3)

The following rows in table ISNAVCPTST.CHK1 do not meet the check condition. You may update or delete the following rows:

DEPARTMENT_ID	MANAGER_ID	DESCRIPTION
99	-1	<NULL>
99	-2	e
99	-2	b
99	-2	a
99	-2	a
99	-2	a
99	-2	a
99	-2	a
99	-2	a
99	-2	a
99	-2	a
99	-2	<NULL>
99	-2	a
99	-3	a
99	-3	a
99	-3	a
99	-3	a
99	-3	a
99	-3	a
99	-3	a
99	-3	a

Rows found: 22 Mode: Changes under commitment control

© 2006 IBM Corporation

IBM System i

Monitoring Exceptions in Applications

- Applications should handle new error messages and status codes
- System messages:
 - Notify: CPF502D, CPF502E, CPF523B
 - Escape: CPF523B, CPF523C
- OPM programs: Mapped to existing error codes
- ILE programs
 - ILE RPG
 - use indicators at first
 - new status code: 01222 and 01022
 - ILE COBOL
 - file status 9R
 - ILE C
 - mapped to existing error numbers
 - SQL
 - SQLCODEs -530, -531, -532
 - SQLSTATEs 23001, 23503, 23504

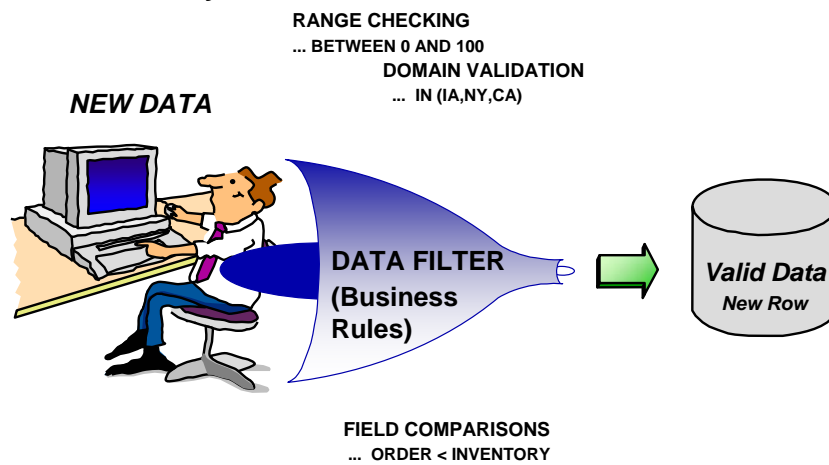
© 2006 IBM Corporation

Agenda

- Introduction
- Referential Integrity
- ➔ • Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

Check Constraints

What are they?



Defining Check Constraints

Native CL Interface:

```
ADDPFCST FILE(Employee) TYPE(*CHKCST)
CST(SalaryChk)
CHKCST(Salary<40000 AND Bonus<=Salary)
```

SQL Interface (Create Table or Alter Table):

```
ALTER TABLE Employee
ADD CONSTRAINT SalaryChk
CHECK(Salary<40000 AND Bonus<=Salary)
```

DB2 UDB for System i locks the table exclusively when adding and verifying a Check Constraint - verification time will NOT be measured in seconds for very large tables

- 4 million row table took 5 minutes to verify

Design Considerations

- Check Condition can contain any expressions or functions allowed on a SQL WHERE clause with the following exceptions:
 - Cannot reference columns in a different table
 - Cannot reference other rows in the table, meaning the following column functions are not allowed: SUM, AVERAGE, MIN, MAX, & COUNT
 - Subqueries are not allowed
- Check Condition clause can reference more than one column in the source table


Design Considerations

- Only single member files supported
- A table has a limit of 300 constraints per file (includes RI, Check, Unique and Primary Key constraints)
- Constraint name has to be unique across all constraint types that exist in the table's library
- Database does NOT prevent conflicting constraints from being defined
 - CHECK Constraint #1
ORDER_STATUS = 'Open'
 - CHECK Constraint #2
ORDER_STATUS <> 'Open'

Monitoring for Constraints

- Check Constraints will be enforced on:
 - Insert into table with constraints
 - Update on table with constraints
 - Delete operation on a Parent Table whose RI Constraint rule is Set Default or Set Null
- New System Message: **CPF502F**
 - **Message:** Check constraint violation on member TEST1.
- ILE Programs
 - ILE RPG: Use indicators, maps to status code 1022
 - ILE COBOL: File Status 9W
 - ILE C: Mapped to existing error number
- OPM Programs: Mapped to existing error codes
- SQL
 - SQLCODE -545
 - SQLSTATE 23513


Agenda

- Introduction
- Referential Integrity
- Check Constraints
-  • Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

Triggers: Introduction

- Triggers are user-written programs
 - associated with a physical file
 - activated by DB2 UDB for System i before or after a database change
 - independent from applications
 - can be developed with any supported compiler
- When do you need triggers?
 - to consistently enforce complex business rules
 - implement special application requirements
 - to monitor critical files
 - validate data
 - in a client-server environment

Triggers: An Example



- When a new order is inserted, a trigger is fired
- Trigger retrieves information about order and customer
- A confirmation fax is automatically sent

OS/400 Support for Triggers

- V4R5 and Prior Releases of OS/400
 - External triggers only (aka Native or System triggers)
 - Two commands
 - ADDPFTRG
 - RMVPFTRG
 - External triggers can be added or removed from a file using the Database function in Operations Navigator
 - A file can have a total of 6 triggers
 - No triggers on catalog files or tables
- V5R1 - Two Types of Triggers
 - External Triggers (see above)
 - First available with DB2/400 in V3R1
 - Originally referred to as Native Triggers
 - SQL Triggers
 - New in V5R1

V5R1 Trigger Enhancements

- SQL Triggers
 - Column level
 - Row level
 - Statement level
- More than 1 trigger per database event
 - maximum 300 per physical file or table
 - triggers for same event, fired in the order created
 - identified or qualified by Trigger name
- CHGPFTRG command
 - disable an active or enabled trigger
 - enable an inactive or disabled trigger
- 'Read only' Trigger
 - use carefully

Trigger Concepts

- Four main components for System or SQL Triggers
 1. base file
 2. trigger event
 - insert
 - delete
 - update
 3. trigger time
 - before
 - after
 4. trigger program
- Triggers have record level scope
- Will be activated whenever the event occurs
- Exception: Update triggers
 - *ALWAYS
 - *CHANGE

Defining Triggers

- Add or remove triggers by using new CL commands

–ADDPFTRG

–RMVPFTRG

```

Add Physical File Trigger (ADDPFTRG)
Type choices, press Enter.
Physical file . . . . . FILE          myfile
Library . . . . .
mylib
Trigger time . . . . . .TRGTIME      *BEFORE
Trigger event . . . . . TRGEVENT     *UPDATE
Program . . . . . PGM                mypgm
Library . . . . .
*LIBL
Replace trigger . . . . . RPLTRG     *NO
Trigger update condition . . . . . .TRGUPDCND *ALWAYS
    
```

- Trigger information will be stored in the file description
- Implications on recompiling, restoring, deleting and renaming trigger programs

Activating Triggers

- Triggers become part of the job that activates them

Opt	Request Level	Program or Procedure	Library	Statement	Instruction
		QCMD	QSYS	0351	
		QUICMENU	QSYS		00C1
1		QUIMNDRV	QSYS		0455
2		QUIMGFLW	QSYS		0483
3		QUICMD	QSYS		03E4
		QUOCCP	QPDA		0541
		QUOMAIN	QPDA		0FDD
4		QUOCMD	QSYS		0176
		T4249CINS	ORDENTLIB	136	00D9
		QSROUTE	QSYS		02F0
		QSQINS	QSYS		01C0
		QDBPUT	QSYS		0193
		T4249RADT	ORDENTLIB	.GET	021D


- Triggers and applications will share:
 - the library list
 - the QTEMP library

Trigger Buffer Fields

<u>Seq</u>	<u>Data Type</u>	<u>Len</u>	<u>Field Description</u>
1	Char	10	Physical file name
2	Char	10	Physical file Library
3	Char	10	Physical file member name
4	Char	1	Trigger event
5	Char	1	Trigger time
6	Char	1	Commit lock level
7	Char	3	Reserved
8	Binary	4	CCSID of data
9	Binary	4	Relative record number
10	Char	4	Reserved
11	Binary	4	Original record offset
12	Binary	4	Original record length
13	Binary	4	Original record null byte map offset
14	Binary	4	Original record null byte map length
15	Binary	4	New record offset
16	Binary	4	New record length
17	Binary	4	New record null byte map offset
18	Binary	4	New record null byte map length
19	Char	16	Reserved
20	Char	Var	Original Record
21	Char	Var	Original record null byte map
22	Char	Var	New record
23	Char	Var	New record null byte map

Triggers Feedback

- What happens if a trigger ends abnormally??
 - The unhandled exception will 'percolate'




- The originating change will fail

Triggers and Application Programs

- How can a trigger notify a 'logical failure' to the application?
- Triggers cannot pass parameters back
- An escape message must be sent back so I/O will fail
- Use the QMHSNDPM API
- The originating change will fail
- Application will get the CPF502B message
- Escape message is logged in the joblog

Example:



Guidelines for Trigger Programs

Operations that are NOT allowed

- Destructive data changes are not allowed if ALWREPCHG(*NO) is specified
 - CANNOT modify the record that fired the trigger, from within the trigger itself
 1. application updates existing record
 2. UPDATE trigger fires
 3. Trigger tries to update same record
 - conflicting operations NOT allowed on the same record
 1. application inserts new record
 2. INSERT trigger fires
 3. Trigger tries to delete new record
 - Contents of trigger buffer would be inconsistent with actual database record
- Commitment control operations
- SQL Connection statements

V5R1 SQL Triggers

- Component Definitions
- Examples

Trigger Components...

- SQL Trigger Components - V5R1+
 1. Base file or table
 2. Trigger name
 3. Trigger event
 4. Trigger time
 5. Trigger granularity
 6. Transition variables
 7. Transition tables
 8. Trigger mode
 9. Triggered action
- Same as system triggers {
- How to work with buffer/when to fire {
- What to do {

SQL Trigger Components...

5. Trigger Granularity

- Column level triggers
 - Extension of UPDATE trigger event
 - Columns listed as part of UPDATE trigger event
 - UPDATE OF column_name_1, column_name_2, ...
 - ✓ Only update of a listed column causes trigger to fire
 - ✓ If no columns listed, update to any column causes trigger to fire
- Row level triggers
 - FOR EACH ROW
 - ✓ Triggered action executed for each row satisfying trigger condition
 - ✓ If trigger condition never satisfied, triggered action never executed
- Statement level triggers
 - FOR EACH STATEMENT
 - ✓ Triggered action executed only once for the event causing the trigger to fire regardless of the number of rows processed
 - ✓ If trigger condition never satisfied, triggered action executed once at end of statement processing
 - Not valid with Before triggers or Trigger Mode of DB2ROW

SQL Trigger Components...

6. Transition Variables

- aka Correlation Variables
- Provides function similar to before and after images in trigger buffer for external triggers
- Qualification of column names for the single row image before and/or after the trigger event has completed
 - OLD ROW - Before image of row
 - NEW ROW - After image of row
 - REFERENCING OLD ROW AS oldrow REFERENCING NEW ROW AS newrow
 - ...newrow.salary > oldrow.salary + 10000...
- Not valid with Statement level triggers

SQL Trigger Components...

7. Transition Tables

- Provides function similar to before and after images in trigger buffer for external triggers
- A single SQL statement can process multiple rows
- Temporary tables that contain the image of all affected rows before and/or after the trigger event completes
 - OLD TABLE - Before image of all affected rows
 - NEW TABLE - After image of all affected rows
 - REFERENCING OLD TABLE AS oldtbl
 - ...(SELECT COUNT(*) FROM oldtbl)...
- Not valid with Before triggers or Trigger Mode of DB2ROW

SQL Trigger Components...

8. Trigger Mode

- MODE DB2ROW
 - Trigger fires after each row operation
 - Only valid with Row level triggers
 - Exclusive function of DB2 UDB for System i
 - ✓ Not available in other DB2 UDB implementations
- MODE DB2SQL
 - Trigger fires after all row operations are complete
 - If specified on a row level trigger, triggered action executed N times after all row operations completed
 - ✓ N = number of rows processed
 - Not as efficient as DB2ROW since each row is processed twice
 - Only valid with After triggers

SQL Trigger Components...

9. Triggered Action

- Analogous to trigger program in external triggers
- Three parts
 - SET OPTION
 - ✓ Specifies the options that will be used to create the trigger
 - WHEN
 - ✓ Search condition or execution criteria for Trigger Body
 - ✓ Specifies when the SQL statements in Trigger Body will be executed
 - SQL Trigger Body
 - ✓ Single SQL statement
 - ✓ Multiple SQL statements delineated with BEGIN and END

SQL Trigger Examples

Row Level Trigger with Simple Trigger Body

```
CREATE TRIGGER audit_spending
  AFTER UPDATE ON expenses
  REFERENCING NEW ROW AS nw
  FOR EACH ROW MODE DB2ROW
  WHEN (nw.total_amount > 10000)
  INSERT INTO travel_audit
    VALUES(nw.empno, nw.deptno, nw.total_amount,
           nw.end_date);
```

SQL Trigger Examples ...

Row Level Trigger with Complex Trigger Body

```
CREATE TRIGGER big_spenders
  AFTER INSERT ON expenses
  REFERENCING NEW ROW AS n
  FOR EACH ROW
  MODE DB2ROW
  WHEN (n.totalamount > 10000)
BEGIN
  DECLARE emplname CHAR(30);
  SET emplname = (SELECT lname FROM employee
                  WHERE empid = n.empno);
  INSERT INTO travel_audit
    VALUES(n.empno, emplname, n.deptno, n.totalamount, n.enddate);
END
```

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary



Column Level Security

What is it?

NAME	SALARY	DEPARTMENT
JOAN	20,000	SALES
FRED	18,000	SALES
LINDA	15,000	ACCOUNTING

Manager – needs
Update Access



Employee – NO
Update Access



Column Level Security

SQL statements required:

```
CREATE TABLE TestTable
  (Name CHAR(30), Salary DEC(8,2),
  Department CHAR (10))
GRANT SELECT,UPDATE(Name, Department) ON TABLE
  Payroll TO Clerk
GRANT SELECT, UPDATE(Name, Department, Salary) ON
  TABLE Payroll TO Manager
```

DSPOBJAUT Test Table

- To show Column Level Authorities
Operations Navigator


No support currently for CL interface:

- send cards and letters to IBM Rochester

Application Integration

- System Message Generated is CPD5035
 - Message: Not authorized to field COL1 of file TEST1
- SQL Code
 - Violations map to existing SQLCODE (-551) SQLSTATE (42501) for database authority violations
- ILE Programs & OPM Programs
 - Violations map to existing error code for database authority violations
- System messages when system is unable to enforce column level authorities due to unexpected problems with user profiles, database file object, or DB2 UDB for System i
 - CPF328F - Unsuccessful authority checking
 - CPD5030 - User Profile Destroyed
 - CPD5032 - User Profile Damage

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
-  • Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

Two-Phase COMMIT...

Applications using multiple relational databases


Single Phase COMMIT - Remote Unit of Work

Local Application

Remote Systems


TORONTO Database

- Application transaction #1
- CONNECT TO TORONTO
- SELECT * FROM EMP-HRS
- UPDATE EMP-MSTR
- UPDATE DEPT-MSTR
- COMMIT



ATLANTA Database

- Application transaction #2
- CONNECT TO ATLANTA
- SELECT * FROM EMP-HRS
- UPDATE EMP-MSTR
- UPDATE DEPT-MSTR
- COMMIT



Two-Phase COMMIT...

Applications using multiple relational databases

Two Phase Commit - Distributed Unit of Work; Application directed


Local Application

Remote Systems

TORONTO Database


Single Application transaction

- CONNECT TO TORONTO
- SELECT * FROM EMP-HRS
- UPDATE EMP-MSTR
- UPDATE DEPT-MSTR



ATLANTA Database

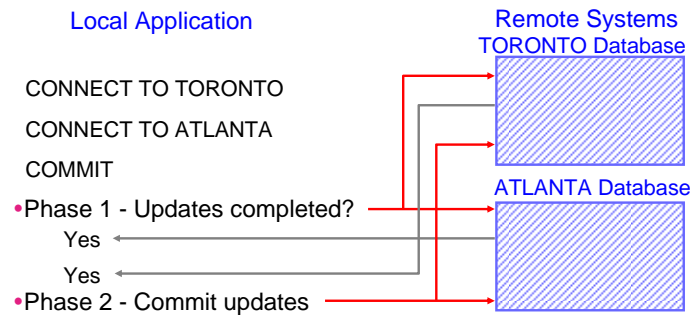
- CONNECT TO ATLANTA
- SELECT * FROM EMP-HRS
- UPDATE EMP-MSTR
- UPDATE DEPT-MSTR
- COMMIT



Two-Phase Commit - Concepts and Terminology

- The two phases of Two-Phase Commit
 - Phase 1: Are databases ready to commit? Have all database updates been completed?
 - Phase 2: Commit database updates for application transaction

Two Phase Commit - Distributed Unit of Work; Application Directed



Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- ➔ • Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

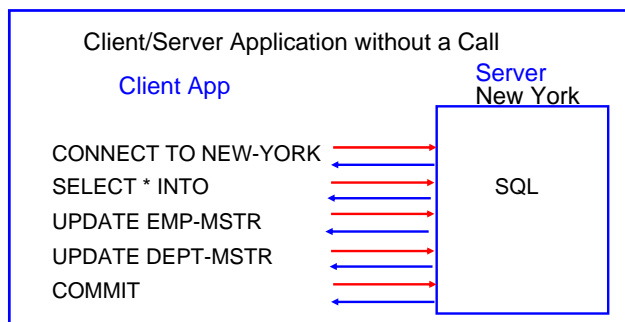
Stored Procedures... (SP)

What is a Stored Procedure?

- A program stored on a System i
 - Languages supported
 - SQL
 - Command Language (CL)
 - RPG
 - COBOL
 - C and C++
 - PL/1
 - JAVA
 - May contain embedded SQL statements
- Can be local or remote
- Supported via new CALL statement in SQL


Stored Procedures...

- Improved Client/Server (DRDA) Performance
 - Without CALL of stored procedure every executable SQL statement using DRDA causes a flow from client application to server
 - Each SQL statement is passed to server to perform SQL function to perform SQL function



Stored Procedures...

- Improved Client/Server (DRDA) Performance
 - CALL of stored procedure results in a significant reduction in conversation between client and server
 - SQL statements are embedded in (compiled) program on server resulting in more efficient performance of SQL



Stored Procedures


Using SQL for Stored Procedures

- Stored Procedures initially supported in V3R1
 - No SQL procedure language
 - Embedded SQL in a HLL programs like C, COBOL, RPG allowed.
- V4R2 + - Stored Procedure can be written purely in SQL
 - First DB2 Family member to deliver this support
- SQL is a programming language

```

CREATE PROCEDURE PROC1 (IN Emp# CHAR(4),IN NwLvl INT)
LANGUAGE SQL Proc1_Src:
BEGIN
  DECLARE CurLvl INT;
  SELECT edlevel FROM emptbl INTO CurLvl
  WHERE empno=Emp#;
  IF NwLvl > CurLvl THEN
    UPDATE emptbl SET edlevel=NwLvl,
    salary=salary + (salary*0.05) WHERE empno=Emp#;
  END IF;
END
  
```

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
-  • SQL Enhancements
- Miscellaneous Enhancements
- Summary

Basic Programming Constructs

Program Control Operations within SQL

- CASE
- IF
- FOR
- LOOP
- REPEAT
- Declaring Variables
- Error Handling

```
IF rating=1
  THEN SET price = price * 0.95;
ELSEIF rating = 2
  THEN SET price = price * 0.90;
ELSE SET price = price * 0.80;
END IF;
```

```
FOR loopvar AS
loopcursor CURSOR FOR
  SELECT firstname, middinit, lastname FROM emptbl
DO
  SET fullname=lastname||', ' || firstname||' ' || middinit;
  INSERT INTO namestbl VALUES( fullname );
END FOR;
```

SQL Logic Commands

Conditional Logic, Declaration of Variables, Etc

- BEGIN and END
- DECLARE (local variables)
- SET (local variables)
- Comments
- CASE (two forms), END CASE
- IF, THEN, ELSE, END IF
- FOR, END FOR, LOOP, END LOOP
- LEAVE (loop or block)
- REPEAT, END REPEAT
- WHILE, END WHILE
- GET DIAGNOSTICS (SQLCA-like information)
- CALL - SQL procedure, External Procedures like HLL
- Normal DDL and DML

FETCH FIRST n ROWS ONLY Clause

- FETCH FIRST n ROWS only clause added for SELECT statements
 - Enables Top N queries
 - Allows limiting the size of the result set - particularly useful for networked ODBC/JDBC clients
- After the Nth row has been fetched, end-of-file (SQLCODE=100) is returned to the user/application
- To guarantee which FIRST n ROWS are returned - an ORDER BY clause must be specified.
 - If ORDER BY is not specified the FIRST n ROWS will be unpredictable

FETCH FIRST n ROWS ONLY Clause

Examples

Query to return top 20 selling products:

```
SELECT product_item, COUNT(*)
  FROM orders
 GROUP BY product_item
 ORDER BY 2 DESC
 FETCH FIRST 20 ROWS ONLY
```

Query to return the top selling product:

```
SELECT product_item, COUNT(*)
  FROM orders
 GROUP BY product_item
 ORDER BY 2 DESC
 FETCH FIRST ROW ONLY
```

JOIN Enhancements

- V5R1 adds new join types for DB2 UDB for System i for more portable applications

```
RIGHT OUTER JOIN
RIGHT JOIN
```

- And removal of minor LEFT OUTER JOIN and INNER JOIN restrictions

- Allow ORs in the ON clause
- Allow the IS NULL, LIKE and BETWEEN predicate in the ON clause

- Examples:

```
SELECT t1.colA, t2.colB FROM t1 RIGHT OUTER JOIN t2 ON t1.c1 = t2.c1
```

```
SELECT * FROM t1 LEFT JOIN t2 ON t1.c1 > t2.c2 OR t1.c2 > t2.c2
```

```
SELECT * FROM t1 LEFT JOIN t2 ON t1.c1 LIKE 'A%'
```

JOIN Enhancements

Employee Table - EMP

NBR	NAM	CLS	SEX	DPT	SAL
20	Heikki	2	M	901	600
10	Ed	5	M	911	700
50	Marcela	3	F	911	750
40	Mike	4	M	977	650
30	John	5	M	977	320
60	Frank	2	M	990	650

Department Table - DEP

DPT	DNM
901	Accounts
977	Manufact
911	Sales
990	Spares

Left Outer Join	All from left
Left Inner Join (Inner Join)	Only records from left with matching record in right
Right Outer Join	All from right
Right Inner Join	Only records from right with matching record in left

SQL UNION

- UNIONS can be included in expressions and derived tables

```
CREATE VIEW total_sales AS
SELECT COUNT(*), SUM(total_sale) FROM Sales1999
WHERE product_id ='XYZ'
```


```
UNION
SELECT COUNT(*), SUM(total_sale) FROM Sales2000
WHERE product_id ='XYZ'
```

```
UNION
SELECT COUNT(*), SUM(total_sale) FROM Sales2001
WHERE product_id='XYZ'
```

Other SQL Enhancements (V5R3)

- SELECT and INSERT statements
 - blocked INSERT
INSERT INTO table1 VALUES (11,'TESTING'), (2,'ADMINISTRATION')
 - EXCEPT and INTERSECT
(SELECT cusnum FROM orders2003) EXCEPT DISTINCT
(SELECT cusnum FROM orders2004)
(return all rows that are in t1, but not t2)
- Enhancements to Stored Procedures
 - scrollable Stored Procedure result sets
 - result set consumer control
- Richer SQL function set
 - REPLACE, EXTRACT, INSERT, REPEAT, DAYNAME, MONTHNAME, RIGHT, TIMESTAMP_ISO
SELECT REPLACE('ABCXYZ','ABC','123') FROM t1
(returns '123XYZ')

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
-  • Miscellaneous Enhancements
- Summary

Predictive Query Governor

- Query Governor added in V3R1
- Limits amount of time for query execution
 - Message issued when time exceeded
 - CPA4259: Estimated query processing time &1 exceeds limit &2
- Decision is made prior to query starting
 - lower impact on system resources
- Change Query Attributes (CHGQRYA) command or via Operations Navigator
 - done on job level
- Query optimizer estimates query processing time
 - compares with time limit from CHGQRYA command
 - makes processing decision

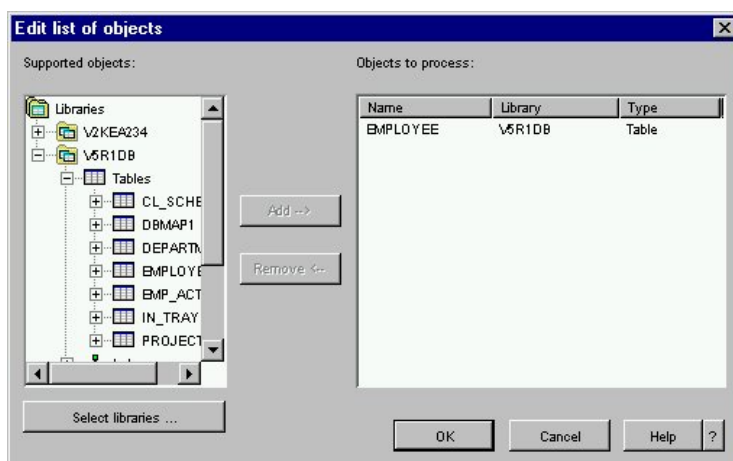
Predictive Query Governor

- Tip for estimating time limit:
 - CHGQRYA - set time limit to 0
 - STRDBG
 - run Query
 - Query will not execute
 - Low level messages will record optimizer estimate of time required to process
 - access plan creation
 - Then set time limit to reasonable value

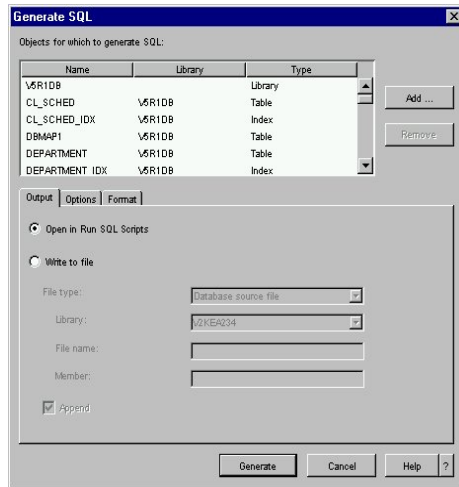
Generate SQL - Operations Navigator

- Reverse engineering of DDL for DB2 objects
 - Useful in converting object definitions from DDS to SQL
 - Not all DDS features can be converted, tool will convert as much as possible and generate warnings for unconvertible options (eg, EDTCDE)
 - Useful in generating SQL script for creating databases
- Object types supported:
 - Tables and physical files (including triggers and constraints)
 - Aliases
 - Functions & Procedures
 - Indexes, Views, & Logical files
 - Schemas (collections) and libraries
 - Distinct Types
- Can generate SQL for one or multiple objects
- Resulting script can be edited
- API interface - QSQGNDDL

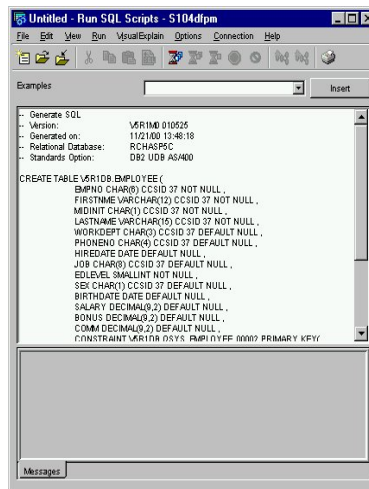
Generate SQL for Multiple Objects



Generate SQL for Multiple Objects



Generated SQL Output



IBM System i

Usability: Visual Explain Index Advisor

The screenshot shows the Visual Explain Index Advisor window. The main area displays a query plan diagram with nodes: Final Select, Nested Loop Join, Table Scan, Index Scan - Key Positioning, and Temporary Index. A red circle highlights the 'Statistics and Index Advisor' button in the toolbar. Another red circle highlights the 'Temporary Index' node in the diagram. On the right, a table shows performance metrics:

Attribute	Value
Time information (start time, tot...	
Timestamp for Creation of Monit...	2002-04-
Statement Start Timestamp	2002-04-
Statement End Timestamp	2002-04-
Optimization Time, in Milliseconds	13
ODP Open Time, in Milliseconds	250
Total Time, in Microseconds	325728
Statement Open Time, in Micros...	325728
Statement Fetch Time, in Micros...	0
Statement Close Time, in Micros...	0
Information about SQL stateme...	
Statement Number	14

Statement text: `SELECT * FROM PART WHERE P_PARTKEY IN (SELECT L_PARTKEY FROM LINEITEM WHERE L_DISCOUNT = 0)`

© 2006 IBM Corporation

IBM System i

Usability: Visual Explain Index Advisor

The screenshot shows the 'Statistics and Index Advisor - I400ws(1400ws)' window. The 'Index Advisor' tab is selected, showing a table of recommended indexes:

Create	Table Name	Library	Index Type	Columns
<input checked="" type="checkbox"/>	LINEITEM	TPCH	Binary Radix	L_PARTKEY ASCEND L_DISCOUNT ASCEND

A red arrow points to the 'Columns' column. Below this, the 'New Index on Table' dialog box is open, showing the following table of columns:

O...	Column Name	Type	Len...	Description
1	A... L_PARTKEY	INTEGER		
	L_SUPPKEY	INTEGER		
	L_LINENUMBER	INTEGER		
	L_QUANTITY	DECIMAL	12,2	
2	A... L_EXTENDEDPRICE	DECIMAL	12,2	
	L_DISCOUNT	DECIMAL	12,2	

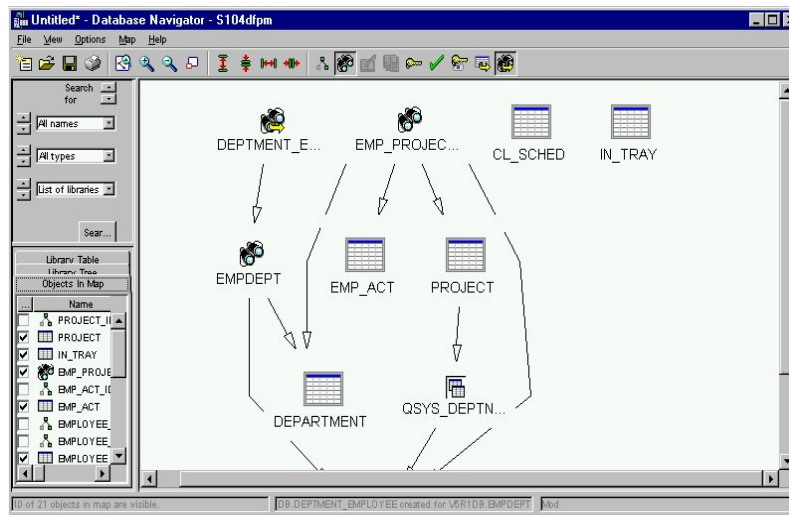
A red arrow points to the 'L_QUANTITY' row. The 'Index type' section has 'Unique' selected. The 'Number of distinct values' is set to 'Not specified'. Buttons for 'OK', 'Cancel', and 'Help' are visible.

© 2006 IBM Corporation

Database Navigator

- Provides a graphical view of a database and its relationships
- Gives a pictorial representation of a schema in order to:
 - Understand an existing complex database schema
 - Create a new schema
 - Manage objects and relationships in the schema
- Resulting picture is a Database Navigator Map (DNM)

Database Navigator Map



IBM System i

Usability: Enhanced Database Navigator

- Trigger Recognition
- Enhanced Graphics
- Improved Printing

The screenshot shows the 'Database Navigator - Asystem' window. On the left, a 'Finding Relations' panel displays a magnifying glass icon and a table of related objects:

Related Objects:			
8	Tables	11	Ft
5	Aliases	Omitted	Jc
1	Check constraints	7	Pr
12	Indexes	0	Ui
Omitted	Journals	20	Vi
1	Triggers		

The main window displays a complex network diagram of database objects. A 'Library Tree' on the left lists various objects like Tables, Views, and Indexes. The status bar at the bottom indicates '65 of 65 objects in map are visible'.

© 2006 IBM Corporation

IBM System i

Agenda

- Introduction
- Referential Integrity
- Check Constraints
- Database Triggers
 - System Triggers
 - SQL Triggers
- Column Level Security
- Two Phase Commit and DRDA Level 2
- Stored Procedures
- SQL Enhancements
- Miscellaneous Enhancements
- Summary

© 2006 IBM Corporation

Summary

- New functions of DB2 UDB for System i
- Very powerful
 - use where required
 - use them carefully
- Explore possible design concerns prior to implementation

Summary...

Bibliography

- SC41-4721 CL Programming
- SC41-4611 DB2 UDB for OS/400: SQL Programming
- SC41-4612 DB2 UDB for OS/400: SQL Reference
- SC41-3212 DB2 UDB for OS/400: Query MANAGER User's Guide
- S246-0100 DB2/400: The New AS/400 Database
 - Skip Marchesani - Midrange Computing
- Database Design and Programming for DB2/400
 - Paul Conte - Duke Communications Press
- SQL/400
 - Paul Conte and Mike Cravitz - Duke Communications Press
- DB2 publications on the Internet
 - ibm.com/series/db2/books.htm
- DB2 UDB for System i Web Page:
 - ibm.com/series/db2
- Redbooks ibm.com/series/db2/redbooks.htm